

# Embodied Representation Alignment with Mirror Neurons

Wentao Zhu<sup>1,2</sup> Zhining Zhang<sup>1</sup> Yuwei Ren<sup>3</sup> Yin Huang<sup>3</sup> Hao Xu<sup>3</sup> Yizhou Wang<sup>1,4</sup>

<sup>1</sup> Center on Frontiers of Computing Studies, School of Computer Science, Peking University

<sup>2</sup> Eastern Institute of Technology, Ningbo <sup>3</sup> Qualcomm AI Research

<sup>4</sup> Inst. for Artificial Intelligence, Peking University

## Abstract

Mirror neurons are a class of neurons that activate both when an individual observes an action and when they perform the same action. This mechanism reveals a fundamental interplay between action understanding and embodied execution, suggesting that these two abilities are inherently connected. Nonetheless, existing machine learning methods largely overlook this interplay, treating these abilities as separate tasks. In this study, we provide a unified perspective in modeling them through the lens of representation learning. We first observe that their intermediate representations spontaneously align. Inspired by mirror neurons, we further introduce an approach that explicitly aligns the representations of observed and executed actions. Specifically, we employ two linear layers to map the representations to a shared latent space, where contrastive learning enforces the alignment of corresponding representations, effectively maximizing their mutual information. Experiments demonstrate that this simple approach fosters mutual synergy between the two tasks, effectively improving representation quality and generalization.

## 1. Introduction

“The body is our general medium for having a world.”

— Maurice Merleau-Ponty

Neuroscience research has uncovered a fascinating mechanism behind multiple cognitive abilities: *mirror neurons*. First identified in macaque monkeys, these neurons fire both when an individual observes an action and when they perform it themselves [4, 9, 32]. In essence, the neural representations of observed and executed actions are inherently *aligned*. Subsequent studies confirm similar systems in the human brain, where observing others’ actions activates corresponding motor regions, as if the observer were performing them [7, 23]. This suggests that action under-




Figure 1. **Conceptual demonstration of mirror neurons.** These neurons activate both when observing an action and when performing it oneself, illustrating the shared neural representations that link perceptual and motor systems.

standing arises from neural simulations of observed behaviors, not merely abstract reasoning [33]. By mapping external movements onto its own motor repertoire, the brain internally “experiences” the action, enabling an intuitive grasp of others’ intentions [21].

This neural mechanism highlights the tight bond between two fundamental cognitive abilities: action understanding and embodied execution. Concretely, action understanding enables agents to interpret the meaning and intent behind others’ actions, while embodied execution allows them to physically interact with the environment to achieve goals. Crucially, these two abilities are deeply interconnected: action understanding supports embodied execution by guiding imitation learning and skill acquisition, while in turn, embodied execution provides firsthand sensorimotor experience that refines and deepens action understanding [6, 26]. The biological foundation of this interplay is exemplified by mirror neurons, which also demonstrates the core idea of *embodied cognition* [1, 3, 8, 36, 40]—the notion that cognitive processes are not merely functions of the brain or abstract activities of the mind, but are deeply rooted in the body’s sensorimotor interactions with the world. Despite this biological synergy, current machine learning approaches typically address action understanding and embodied execution independently, overlooking their potential to inform and enhance each other [12, 41, 46]. This separation impedes the learning of generalizable and comprehensive action representations, which in turn limits

performance on downstream tasks.

In light of this, this work proposes to unify action understanding and embodied execution through the lens of representation learning. We begin by examining the relationship between the neural representations learned by models for these two tasks. Specifically, we investigate whether the representations align when observing and executing the same actions, how this alignment evolves during training, and its correlation with task success (Sec. 2). Building on these insights, we propose a paradigm that jointly trains both models, bridging the gap between action understanding and embodied execution. The core idea is to project the agent’s representations of observed and executed actions into a shared latent space, where contrastive learning enforces their alignment (Sec. 3). This approach explicitly aligns the representations of observed and executed actions, which is exactly what biological mirror neurons do. From the viewpoint of information theory, it is equivalent to maximizing the mutual information between the neural representations of the same semantic actions in action understanding and embodied execution. We evaluate our method on action recognition and multi-task object manipulation benchmarks, demonstrating that the proposed framework enables the two tasks to reinforce each other (Sec. 4). Our results show that the learned representations are more disentangled and robust, leading to improved generalization. Furthermore, we investigate the impact of different alignment strategies, exploring the appropriate granularity for alignment. We summarize the key contributions of this paper below:

- We conceptualize action understanding and embodied execution as a unified system, grounded in neuroscience and cognitive insights.
- We discover that these two models exhibit representation alignment spontaneously, which correlates to some extent with task success.
- We introduce a representation learning approach inspired by mirror neurons, which directly aligns the neural representations of observing and executing corresponding tasks, thereby effectively maximizing their mutual information.
- Experiments demonstrate that our method is simple and effective, enhancing the generalization and representation learning quality for both tasks.

## 2. Probing Representation Alignment

First, we aim to investigate whether action understanding and embodied execution models, when trained separately, exhibit neural representation alignment similar to mirror neurons. If such alignment exists, we further seek to understand the factors that influence it.

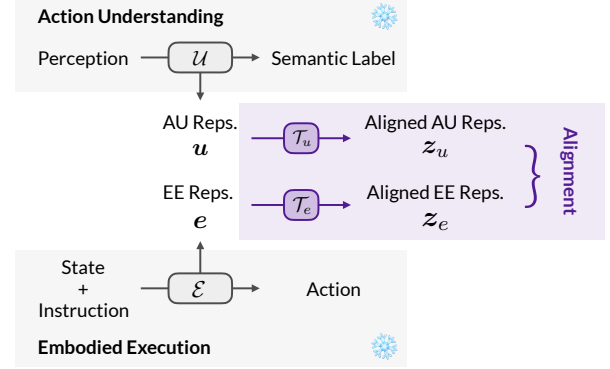


Figure 2. **Alignment probing of off-the-shelf model representations.** We extract internal neural representations from an Action Understanding (AU) model and an Embodied Execution (EE) model and train two linear transformations,  $\mathcal{T}_u$  and  $\mathcal{T}_e$ , to align them. Both models are pretrained separately and remain frozen during this process.

### 2.1. Problem Formulation

Given an action understanding model  $\mathcal{U}(\cdot; \theta_u)$  and an embodied execution model  $\mathcal{E}(\cdot; \theta_e)$ , we aim to extract their intermediate neural representations and study their alignment. Specifically, the action understanding model  $\mathcal{U}$  takes a perceptual sequence  $V$ , such as a video, and produces an internal representation  $u$ , from which it predicts the semantic label  $\hat{y} = \mathcal{U}(V)$ . Meanwhile, the embodied execution model  $\mathcal{E}$  takes the environment state  $S$  and an instruction  $I$  (e.g., a natural language command) as input, generating the representation  $e$ , which is then decoded to the next action  $\hat{a} = \mathcal{E}(S, I)$ .

The neural representations of observed and executed actions,  $u$  and  $e$ , originally reside in different high-dimensional representation spaces. In principle, they share some common information, such as the semantic concepts and spatial relationships of the scene. Meanwhile, they would also capture unique aspects specific to the task. For instance,  $u$  may emphasize high-level action semantics and visual patterns essential for understanding, whereas  $e$  is more attuned to physical constraints and feasibility, which are crucial for execution.

Therefore, to effectively measure their alignment, we aim to learn a pair of linear transformations,  $\mathcal{T}_u(\cdot; \phi_u)$  and  $\mathcal{T}_e(\cdot; \phi_e)$ , that map  $u$  and  $e$  onto a shared latent space  $\mathbb{Z}$  and align them. Formally:

$$z_u = \mathcal{T}_u(u) \in \mathbb{Z}, \quad z_e = \mathcal{T}_e(e) \in \mathbb{Z}, \quad (1)$$

where  $z_u$  and  $z_e$  denote the aligned action understanding and embodied execution representations, respectively.

### 2.2. Probing

We employ alignment probing [43] to align the representations and evaluate the alignment. Figure 2 provides an

overview of the pipeline. Given off-the-shelf  $\mathcal{U}(\cdot)$  and  $\mathcal{E}(\cdot)$ , we use contrastive learning to align the corresponding action pairs in the shared latent space  $\mathbb{Z}$ . Specifically, we optimize the following bidirectional InfoNCE loss:

$$\mathcal{L}_{\text{align}} = -\frac{1}{2B} \sum_{i=1}^B \left[ \log \frac{\exp(\text{sim}(\mathbf{z}_u^{(i)}, \mathbf{z}_e^{(i)})/\tau)}{\sum_{j=1}^B \exp(\text{sim}(\mathbf{z}_u^{(i)}, \mathbf{z}_e^{(j)})/\tau)} + \log \frac{\exp(\text{sim}(\mathbf{z}_e^{(i)}, \mathbf{z}_u^{(i)})/\tau)}{\sum_{j=1}^B \exp(\text{sim}(\mathbf{z}_e^{(i)}, \mathbf{z}_u^{(j)})/\tau)} \right]. \quad (2)$$

where  $B$  is the batch size, and  $\tau$  is a temperature scaling parameter. Each pair  $(\mathbf{z}_u^{(i)}, \mathbf{z}_e^{(i)})$  consists of feature representations derived from observing and executing actions that share the same intent. The similarity function  $\text{sim}(\cdot, \cdot)$  is defined as the cosine similarity:

$$\text{sim}(\mathbf{z}_u, \mathbf{z}_e) = \frac{\mathbf{z}_u^\top \mathbf{z}_e}{\|\mathbf{z}_u\| \|\mathbf{z}_e\|}. \quad (3)$$

We optimize the parameters of the linear transformations to minimize the alignment loss:

$$(\phi_u^*, \phi_e^*) = \arg \min_{\phi_u, \phi_e} \mathcal{L}_{\text{align}}. \quad (4)$$

From an information-theoretic perspective, this optimization objective is equivalent to optimizing a lower bound to estimate the mutual information between the action understanding representation  $\mathbf{u}$  and the embodied execution representation  $\mathbf{e}$ . The theoretical derivation can be found in Appendix A and Theorem 1. In practice, we adopt ViCLIP [41], a video-text representation learning model, as the action understanding model  $\mathcal{U}$ . We take the output feature of the video encoder as  $\mathbf{u}$ . Furthermore, we utilize ARP [44], a pretrained language-conditioned robotic manipulation model, as the embodied execution model  $\mathcal{E}$ . We take the output feature from the last block of the policy network (a chunking causal transformer) as  $\mathbf{e}$ . For more implementation details, please refer to Sec. 4.1 and Appendix C. To measure the degree of alignment, we compute the average Recall@1 of bidirectional nearest neighbor retrieval on a held-out test set after training  $\mathcal{T}_u(\cdot; \phi_u)$  and  $\mathcal{T}_e(\cdot; \phi_e)$ .

### 2.3. Observations

We explore whether models of action understanding and embodied execution, though trained in isolation, converge spontaneously—like dancers finding harmony without rehearsal—in their neural representations. To this end, we measure the representation alignment of the two models at different stages of the training process. Figure 3 (L) shows that the alignment between the two models increases rapidly in the early stages of training and reaches a high level. In




Figure 3. **Alignment probing results.** Left: Alignment scores tracked over the model training progress. Right: Alignment scores computed separately for task success and failure subsets from two pretrained models. We use the average retrieval accuracy (%) on the test set as a measure of representation alignment. Alignment is trained and tested under four different settings while ensuring equal sample sizes.

particular, training only two linear transformations achieves more than 60% accuracy in determining whether the representations correspond to the same underlying action. This phenomenon may arise because the action understanding model and the embodied execution model, despite being trained for different tasks, both require effective modeling and abstraction of the underlying principles and structured patterns of object interactions. This observation also aligns with the *Platonic Representation Hypothesis* [15], which states that neural networks, trained with different objectives on different datasets and modalities, converge to a shared statistical model of reality in their representation spaces.

**Observation #1:** Independently trained action understanding and embodied execution models exhibit a swift emergence of meaningful neural alignment, suggesting a convergence towards representations of the common underlying reality.

Furthermore, we aim to investigate whether there exists a relationship between the degree of representation alignment and task success. To this end, we conduct two sets of experiments: (1) We compute the representation alignment for correctly and incorrectly recognized samples in action understanding with respect to embodied execution. (2) Similarly, we compute the alignment for successfully and unsuccessfully completed tasks in embodied execution with respect to action understanding. In all experiments, we control for the same number of training and testing samples in alignment probing and train separate probes for each setting.

As shown in Figure 3 (R), we observe that the alignment score for the subset of task-successful samples is significantly higher than that for the subset of task-failed samples. We hypothesize that this may be because task-successful samples are associated with higher-quality representations

of the underlying reality, leading to more robust alignment.

**Observation #2:** Task-successful samples exhibit significantly higher representation alignment, potentially due to their higher-quality representations of the underlying reality.

This observation further motivates us to consider whether a causal relationship exists. Specifically, is representation alignment merely a byproduct of better representations, or can it also actively contribute to their formation? Could promoting neural representation alignment improve representation quality and, in turn, benefit the task performance of both models?

### 3. Aligning Representations with Mirror Neurons

Inspired by the observations above, we further explore explicitly aligning the neural representations of action understanding and embodied execution during model training. Structurally, this approach establishes a shared representational space by directly linking perceptual and motor pathways, reflecting the anatomical organization of sensorimotor circuits in biological systems. Functionally, these structural principles promote bidirectional information flow, enabling seamless sensorimotor integration, akin to the role of mirror neurons in action recognition and execution.

To achieve this, we simultaneously train the action understanding model  $\mathcal{U}(\cdot; \theta_u)$ , and the embodied execution model  $\mathcal{E}(\cdot; \theta_e)$ , in a coupled structure, as shown in Fig. 4. We adopt a minimal modification approach, where, on top of the original training objectives of both models, we only introduce two linear transformations  $\mathcal{T}_u(\cdot; \phi_u)$  and  $\mathcal{T}_e(\cdot; \phi_e)$  to align their internal representations throughout this process. Specifically, we denote the training objective of the action understanding model as  $\mathcal{L}_{AU}(\theta_u)$  and that of the embodied execution model as  $\mathcal{L}_{EE}(\theta_e)$ . We define these two task-specific training objectives following previous work [41, 44], as they are not the primary focus of this study. Notably, we simply introduce a bidirectional contrastive loss  $\mathcal{L}_{align}$ , as defined in Eq. (2). The first term encourages each sample’s representation in the action understanding model to align closely with its paired counterpart in the embodied execution model while distinguishing it from other samples in the batch. The second term symmetrically enforces the same constraint for the embodied execution model relative to the action understanding model. This bidirectional alignment objective helps regularize the latent spaces robustly, effectively bringing the representations of observing and executing similar actions closer while pushing apart those of different actions.

Our final training objective can be formulated as:

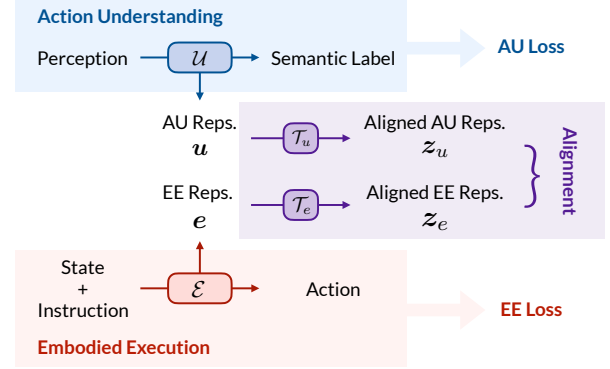


Figure 4. **Aligning action understanding and embodied execution with mirror neurons.** We employ two linear layers to align intermediate representations between the Action Understanding (AU) model and the Embodied Execution (EE) model. The two models are jointly trained to align their representations while simultaneously optimizing for their respective task objectives.

$$\mathcal{L}_{final} = \mathcal{L}_{AU} + \lambda_{EE} \mathcal{L}_{EE} + \lambda_{align} \mathcal{L}_{align}, \quad (5)$$

where  $\lambda_{EE}$  and  $\lambda_{align}$  are hyperparameters to balance the loss scales. We jointly optimize parameters of all models end-to-end:

$$(\phi_u^*, \phi_e^*, \theta_u^*, \theta_e^*) = \arg \min_{\phi_u, \phi_e, \theta_u, \theta_e} \mathcal{L}_{final}. \quad (6)$$

In this process, the optimization objective of representation alignment effectively regularizes the training of  $\mathcal{U}(\cdot; \theta_u)$  and  $\mathcal{E}(\cdot; \theta_e)$  simultaneously. This objective is consistent with the mechanism of mirror neurons, which map observed and executed similar actions to shared neural representations. From an information-theoretic perspective, this optimization objective is equivalent to optimizing a lower bound to maximize the mutual information between  $u$  and  $e$ . Please refer to Appendix A and Theorem 2 for the theoretical derivation.

## 4. Experiments

In this section, we aim to investigate the impact of the proposed mirror framework on model training. Specifically, we design experiments to address the following questions: (1) How does it affect the performance of models for action understanding and embodied execution, respectively? (2) What influence does it have on the representations of these two tasks? (3) Which representations should be aligned, and what are the effects of this strategy?

### 4.1. Implementation details

**Action Understanding** We adopt ViCLIP [41], a generic video-language model, consisting of a video encoder (a



Method	Close Jar	Drag Stick	Insert Peg	Meat off Grill	Open Drawer	Place Cup	Place Wine	Push Btn.	Put in Cpd.	Put in Drawer	Put in Safe	Screw Bulb	Slide Block	Sort Shape	Stack Block	Stack Cup	Sweep Dust	Turn Tap	Avg.
ViCLIP [41]†	0.0	0.0	0.0	48.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	16.0	8.0	4.0	52.0	60.0	10.4
ViCLIP [41]	69.3	<b>96.0</b>	21.3	90.7	<b>100.0</b>	88.0	<b>56.0</b>	<b>76.0</b>	14.7	<b>100.0</b>	42.7	78.7	80.0	77.3	<b>38.7</b>	69.3	<b>100.0</b>	89.3	71.6
MN (Ours)	<b>72.0</b>	<b>96.0</b>	<b>33.3</b>	<b>96.0</b>	<b>100.0</b>	<b>89.3</b>	<b>56.0</b>	<b>76.0</b>	<b>17.3</b>	<b>100.0</b>	<b>49.3</b>	<b>84.0</b>	<b>82.7</b>	<b>85.3</b>	37.3	<b>80.0</b>	<b>100.0</b>	<b>93.3</b>	<b>74.9</b>

Table 1. Performance comparison of action recognition across 18 diverse tasks. † indicates results obtained without task-specific fine-tuning (zero-shot).

standard ViT [5] with spatiotemporal attention) and a text encoder (identical to that of CLIP [31]). We employ the model weights pretrained on InternVid [41] and fine-tune it in the object manipulation domain. Specifically, we use demonstration trajectories from RLBench [18, 37], where RGB videos rendered from a front-facing perspective are paired with language instructions, to fine-tune both the video encoder and the text encoder. The learning rate of the video encoder is  $1 \times 10^{-5}$ , and that of the text encoder is 5% of it. Subsequently, we evaluate action recognition accuracy on a test set by selecting the language instruction with the highest similarity for each test video and computing the accuracy accordingly.

**Embodied Execution** We use an Autoregressive Policy (ARP) [44] with a Multi-View Transformer (MVT) [12, 13] backbone. The model takes multi-view RGBD images (processed into point clouds) and language instructions as input, predicting the next target end-effector pose and gripper states. We conduct experiments on a standard multi-task manipulation benchmark from RLBench [18]. The benchmark includes 18 tasks, each defined by a language description and featuring 2 to 60 variations, such as different object colors or locations. A Franka Panda robot with a parallel jaw gripper is tasked with execution, simulated via CoppeliaSim [34]. We train and test on the same dataset as prior works [12, 13, 37, 44], using 100 demonstrations per task for training and 25 unseen demonstrations for testing. Training settings and hyperparameters remain identical to those in the baseline [44].

**Representation Alignment** We take the output features of the video encoder and the last block of the policy network and map them to  $\mathbb{Z} \subset \mathbb{R}^{1024}$ , using two separate linear layers. For contrastive learning training, we construct positive sample pairs based on language instruction consistency. That is, observing and executing actions with the same language instruction (e.g., “take the steak off the grill”) are encouraged to align in representation, even if they do not come from the exact same episode (e.g., differing in object layout or action sequence). We set the temperature parameter to  $\tau = 0.1$  and the learning rate to  $1 \times 10^{-4}$ . The loss weights are set as  $\mathcal{L}_{EE} = 1$  and  $\mathcal{L}_{align} = 0.5$ .

## 4.2. Action Understanding

First, we investigate the effect of the proposed mirror neuron alignment framework on the action understanding model. To this end, we compare the top-1 action recognition accuracy with regard to language instruction among the following models: the original ViCLIP model pretrained on InternVid [41], tested in a zero-shot manner; the same model fine-tuned on object interaction data; and one trained jointly with the proposed mirror neuron alignment objective. For the latter two, all other training conditions remain the same.

Table 1 shows that ViCLIP exhibits some level of zero-shot action recognition capability, considering the total number of 200+ possible fine-grained variants. However, its performance on most action classes is suboptimal, as fine-grained classification requires nuanced spatial (e.g., “stack the wine bottle to the right of the rack”), temporal (e.g., “push the maroon button, then push the green button”), and quantitative reasoning (e.g., “stack 2 maroon blocks”), which is challenging and relatively scarce in large-scale pretraining data. Fine-tuning on the corresponding dataset significantly enhances the model performance. Notably, our approach extensively outperforms the baseline in action recognition accuracy.

We attribute this improvement to the intrinsic connection and complementary nature between the embodied execution task and action recognition. For example, action recognition must distinguish between commands such as “put the ring on the azure spoke” and “put the ring on the yellow spoke”, recognizing their nuances in relation to the corresponding goal. Meanwhile, embodied execution not only requires identifying the correct 3D interaction locations (affordances) but also generating the appropriate motion trajectory to complete the action. The feature interaction and alignment mechanisms in the proposed framework facilitate learning a more comprehensive task representation, which in turn enhances the generalization ability of action recognition. These results demonstrate the synergistic role of embodied execution in improving action understanding, highlighting the effectiveness of our approach.

## 4.3. Embodied Execution

We investigate the impact of the proposed framework on embodied execution tasks, specifically language-conditioned multi-task object manipulation. Following

Method	Close Jar	Drag Stick	Insert Peg	Meat off Grill	Open Drawer	Place Cup	Place Wine	Push Btn.	Put in Cpd.	Put in Drawer	Put in Safe	Screw Bulb	Slide Block	Sort Shape	Stack Block	Stack Cup	Sweep Dust	Turn Tap	Avg.
BC-Z CNN [20]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.0	1.3
BC-Z ViT [20]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	16.0	1.3
C2F-ARM [19]	24.0	24.0	4.0	20.0	20.0	0.0	8.0	72.0	0.0	4.0	12.0	8.0	16.0	8.0	0.0	0.0	0.0	68.0	20.1
HiveFormer [14]	52.0	76.0	0.0	100.0	52.0	0.0	80.0	84.0	32.0	68.0	76.0	8.0	64.0	8.0	8.0	0.0	28.0	80.0	45.3
PolarNet [2]	36.0	92.0	4.0	100.0	84.0	0.0	40.0	96.0	12.0	32.0	84.0	44.0	56.0	12.0	4.0	8.0	52.0	80.0	46.4
PerAct [37]	55.2 $\pm 4.7$	89.6 $\pm 4.1$	5.6 $\pm 4.1$	70.4 $\pm 2.0$	88.0 $\pm 5.7$	2.4 $\pm 3.2$	44.8 $\pm 7.8$	92.8 $\pm 3.0$	28.0 $\pm 4.4$	51.2 $\pm 4.7$	84.0 $\pm 3.6$	17.6 $\pm 2.0$	74.0 $\pm 13.0$	16.8 $\pm 4.7$	26.4 $\pm 3.2$	2.4 $\pm 2.0$	52.0 $\pm 0.0$	88.0 $\pm 4.4$	49.4
Act3D [11]	92.0	92.0	27.0	94.0	93.0	3.0	80.0	99.0	51.0	90.0	95.0	47.0	93.0	8.0	12.0	9.0	92.0	94.0	65.0
RVT [12]	52.0 $\pm 2.5$	99.2 $\pm 1.6$	11.2 $\pm 3.0$	88.0 $\pm 2.5$	71.2 $\pm 6.9$	4.0 $\pm 2.5$	91.0 $\pm 5.2$	100.0 $\pm 0.0$	49.6 $\pm 3.2$	88.0 $\pm 5.7$	91.2 $\pm 3.0$	48.0 $\pm 5.7$	81.6 $\pm 5.4$	36.0 $\pm 2.5$	28.8 $\pm 3.9$	26.4 $\pm 8.2$	72.0 $\pm 0.0$	93.6 $\pm 4.1$	62.9
RVT-2 [13]	<b>100.0</b> $\pm 0.0$	99.0 $\pm 1.7$	40.0 $\pm 0.0$	<b>99.0</b> $\pm 1.7$	74.0 $\pm 11.8$	38.0 $\pm 4.5$	95.0 $\pm 3.3$	<b>100.0</b> $\pm 0.0$	66.0 $\pm 4.5$	96.0 $\pm 0.0$	<b>96.0</b> $\pm 2.8$	88.0 $\pm 4.9$	<b>92.0</b> $\pm 2.8$	35.0 $\pm 7.1$	<b>80.0</b> $\pm 2.8$	69.0 $\pm 5.9$	<b>100.0</b> $\pm 0.0$	99.0 $\pm 1.7$	81.4
ARP [44]	<b>100.0</b> $\pm 0.0$	<b>100.0</b> $\pm 0.0$	93.3 $\pm 2.3$	92.0 $\pm 0.0$	90.7 $\pm 2.3$	49.3 $\pm 6.1$	93.3 $\pm 8.3$	<b>100.0</b> $\pm 0.0$	66.7 $\pm 6.1$	<b>100.0</b> $\pm 0.0$	88.0 $\pm 4.0$	<b>92.0</b> $\pm 4.0$	86.7 $\pm 6.1$	49.3 $\pm 6.1$	56.0 $\pm 4.0$	82.7 $\pm 6.1$	98.7 $\pm 2.3$	97.3 $\pm 2.3$	85.3
MN (Ours)	<b>100.0</b> $\pm 0.0$	<b>100.0</b> $\pm 0.0$	<b>94.7</b> $\pm 4.6$	93.3 $\pm 2.3$	<b>93.3</b> $\pm 2.3$	<b>53.3</b> $\pm 12.2$	<b>97.3</b> $\pm 2.3$	<b>100.0</b> $\pm 0.0$	<b>70.7</b> $\pm 2.3$	<b>100.0</b> $\pm 0.0$	93.3 $\pm 2.3$	88.0 $\pm 0.0$	82.7 $\pm 2.3$	<b>66.7</b> $\pm 10.1$	72.0 $\pm 6.9$	<b>93.3</b> $\pm 4.6$	<b>100.0</b> $\pm 0.0$	<b>100.0</b> $\pm 0.0$	<b>88.8</b>

Table 2. Performance comparison of embodied execution across 18 diverse tasks. A single model is evaluated across all tasks, with success criteria defined according to RL Bench [18].

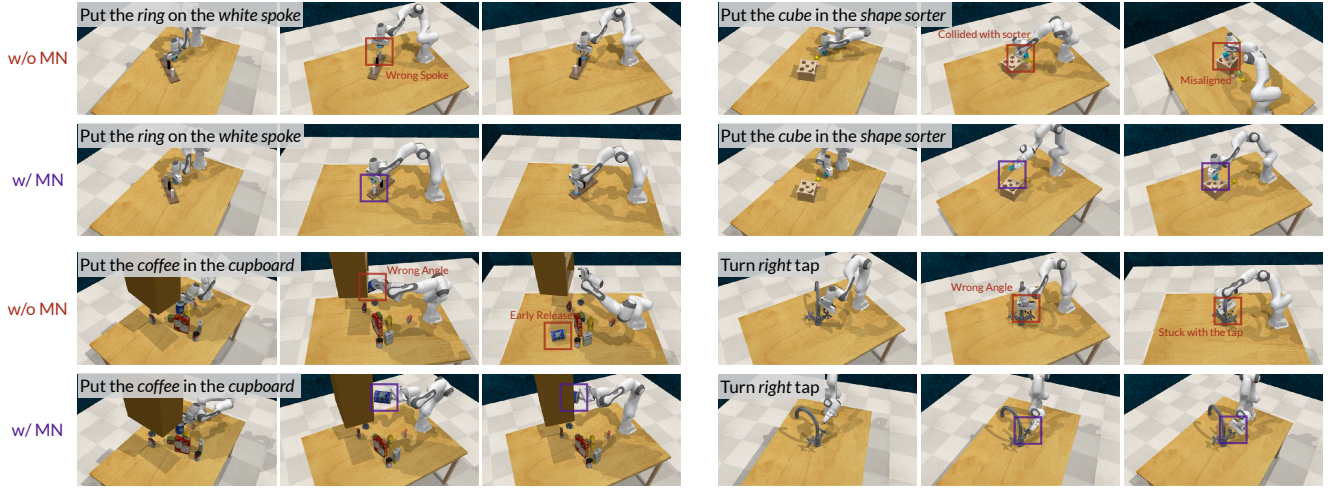


Figure 5. Visualization comparison of embodied execution results. We compare our approach with a baseline model without the MN design (*i.e.*, ARP [44]). The visualization sequences are captured from rotating camera perspectives, with key details highlighted in boxes for clarity. The proposed MN design helps learning affordances and fine-grained operations, leading to improved performance.

prior works, we evaluate the trained policy by performing rollouts in unseen environments from RL Bench [18] and measuring the success rates for various tasks. We compare our approach with various baselines, including simple image-to-action behavioral cloning baselines [20], 3D object manipulation methods [12–14, 18, 37], as well as the most directly comparable baseline, ARP [44]. All approaches are trained and tested using input images of size  $128 \times 128$ , except for Act3D, which uses images of size  $256 \times 256$ .

As shown in Table 2, our method achieves notable improvements over state-of-the-art approaches in most tasks, with an average success rate increase of 3.5% compared to the direct baseline. Furthermore, it demonstrates significant gains in tasks requiring fine-grained affordance reasoning, such as *Sort Shape* and *Stack Cup*. To better understand the reasons behind the improvements, we compare the rollout trajectories of models with the MN design against those without MN (which degenerates to ARP [44]). Fig-

ure 5 presents several comparison results. Our proposed MN method demonstrates enhanced understanding of affordances, such as how to interact with different objects. It also improves target accuracy and refines the precision of fine-grained interactions. Additional details can be found in the supplementary videos.

Several factors may contribute to these improvements. First, the representation alignment with the action recognition model enables learning a more comprehensive and disentangled task representation, which helps improve robust generalization. Second, action recognition may provide higher-quality appearance and geometric representations, which, through alignment, enhance the representation learning of embodied execution and improve its corresponding task performance.

#### 4.4. Representation Analysis

We aim to further analyze the impact of the proposed MN module on learned neural representations to gain deeper in-

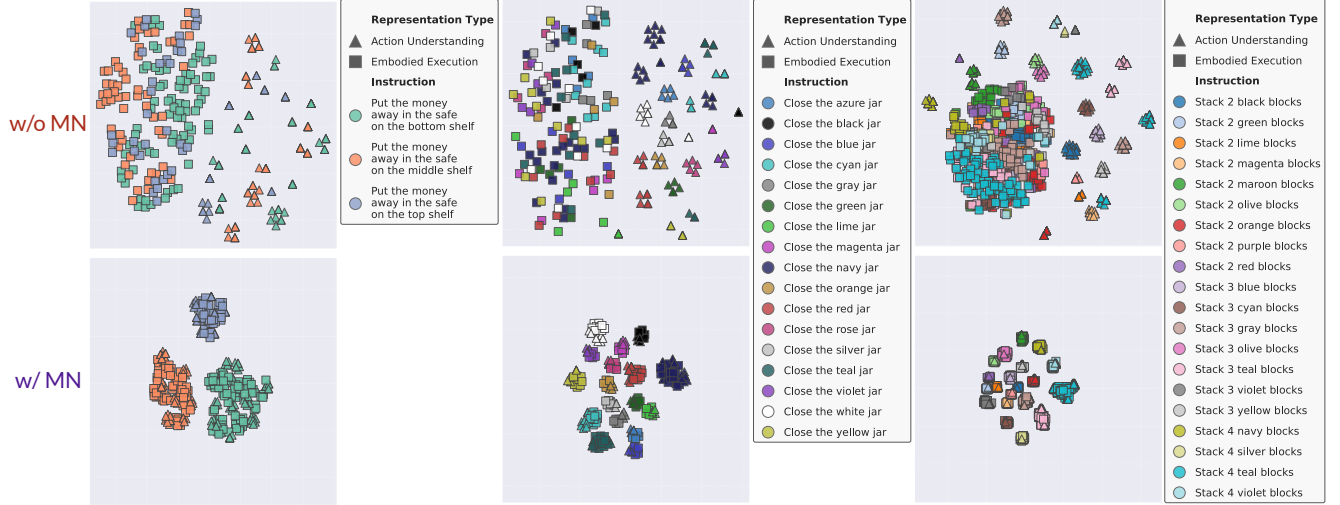


Figure 6. Visualization comparison of latent representations. We align the representations of two models with the MN design and two models without it into their respective shared latent spaces, followed by dimensionality reduction for visualization. Different shapes represent representations from different models, while different colors indicate the corresponding language instructions.

sight. We align the representations of action understanding and embodied execution to a shared latent space  $\mathbb{Z}$  using alignment probing and then visualize them via t-SNE [39] dimensionality reduction. Specifically, we compare baseline models that are trained separately without the MN module to our proposed MN approach, which couples the training of both models.

Figure 6 shows that our proposed method not only facilitates the alignment of representations between action understanding and embodied execution, as indicated by shapes of the same color clustering together, but also enhances the ability to distinguish fine-grained nuances in instructions, with different colors forming distinct clusters. We believe that this representation alignment and disentanglement may be the cause of the improved generalization ability. Additionally, we find that the MN module exhibits a similar disentanglement effect on representations before the linear transformation (in higher-dimensional space). We also illustrate the evolution of representations throughout the training process; please see Appendix D for details.

#### 4.5. Ablation Study

Additionally, we study how the implementation of the proposed mirror neuron alignment module influences model performance. We primarily explore two questions: (1) What kind of action representations should be aligned? (2) How strict should the alignment criterion be? These questions correspond to fundamental aspects of contrastive learning: positive sample construction and temperature.

In our experiments, we explore three different strategies for constructing positive samples:

1. **By Episode:** Positive samples are drawn from the same episode. This means that the action understanding and

Task	Action Understanding			Embodied Execution		
	0.02	0.1	0.2	0.02	0.1	0.2
By Episode	74.0	72.9	74.0	86.4	88.1	87.8
By Instruction	74.9	74.9	77.1	86.7	88.8	87.0
By Class	73.6	71.6	72.2	86.9	85.7	85.6

Table 3. Ablation study results. We analyze the impact of alignment strategies and contrastive learning temperatures ( $\tau$ ). Model performance is evaluated on both Action Understanding and Embodied Execution tasks.

embodied execution models are aligned when learning from the exact same episode, including object placement, initial positions, and action progression, although their input and prediction modalities are different. In practice, we sample paired data from the same episode for both models to construct positive sample pairs.

2. **By Instruction:** Positive samples need not originate from the same episode. Even if scene layouts and action sequences differ, as long as the underlying goal corresponds to the same language instruction (e.g., “open the top drawer”), their representations are aligned. In implementation, for each sample, we randomly sample a paired positive sample from the subset corresponding to the same instruction. This is the default setting used in our other experiments.
3. **By Task:** This is the most relaxed criterion. Positive samples do not even need to correspond to the same language instruction (e.g., “open the top drawer” and “open the bottom drawer”). As long as they belong to the same action class (*Open Drawer*), we align their representations accordingly.

Since the choice of a reasonable temperature parameter is closely related to the definition of positive samples, we

evaluate each strategy using three different temperature values:  $\tau = 0.02, 0.1, 0.2$ . Based on the results in Table 3, we derive the following observations:

- All strategies generally outperform the baseline, demonstrating the effectiveness of aligning representations of the two modules akin to MN.
- The optimal temperature varies depending on the positive sample construction strategy. For example, *by episode* benefits from a higher temperature, whereas *by class* performs better with a lower temperature. This aligns with the relationship between positive sample definitions and intra-/inter-class distances.
- Overall, *by instruction* proves to be a well-balanced strategy, maintaining a good trade-off between variation and semantic consistency, leading to the best generalization performance. In contrast, *by class* exhibits slightly lower performance, possibly because the two tasks still require distinguishing between fine-grained language descriptions.

## 5. Related Work

### 5.1. Mirror Neurons

The discovery of mirror neurons (MNs) in macaques represents a pivotal advancement in understanding the neural mechanisms underlying action perception and execution. These neurons, primarily located in area F5 of the premotor cortex and in the inferior parietal lobule (IPL), activate both when an individual performs a specific action and when they observe another individual executing the same or a similar action [4, 9, 32]. In humans, neuroimaging studies provide evidence for a homologous mirror neuron system, including motor-related regions such as the precentral gyrus and the inferior frontal gyrus [10, 16, 23]. These areas are activated during both action execution and observation, suggesting their role in forming shared neural representations of motor and sensory experiences. This shared representation is thought to facilitate action understanding, imitation, and potentially social cognition [17, 38]. Inspired by the functionality of MNs, several studies propose corresponding computational models [25, 30, 35, 45]. Nonetheless, existing approaches have yet to fully establish a unified representation that integrates embodied action execution with action understanding. Furthermore, the contribution of mirroring to these cognitive processes remains underexplored.

### 5.2. Representation Alignment

Representation alignment seeks to bridge the gap between feature representations across different models, modalities, or domains, emerging as a vital topic in machine learning. It has been shown to serve as a meaningful objective for improving model performance, enhancing training efficiency, and enabling generalization across diverse

tasks. For instance, prior studies have explored the similarity of neural network representations, revealing that models with different architectures and initializations tend to exhibit a certain degree of alignment [24, 29]. In the context of cross-modal alignment, researchers achieve robust task generalization by aligning vision and language representations [31], scaling this process with noisy text supervision [22], or leveraging pretrained unimodal models [43]. Moreover, recent studies have found that internal representations tend to align even within separately trained unimodal models [15, 28, 43]. Specifically, *The Platonic Representation Hypothesis* [15] suggests that this convergence reflects the emergence of a shared statistical model of reality, reminiscent of Plato’s concept of an idealized world. In generative modeling, aligning with pretrained discriminative representations has also been shown to significantly enhance image generation quality and accelerate training convergence [27, 42]. Our approach explores representation alignment from a novel perspective, integrating insights from mirror neurons. Specifically, we investigate how neural representations for action understanding and embodied execution align at both the task and functional levels in embodied agents.

## 6. Conclusion

In this paper, we present a novel framework that unifies action understanding and embodied execution through representation learning, inspired by the biological mechanism of mirror neurons. We first discover that models separately trained for these two tasks exhibit spontaneous representation alignment, which is associated with task success. Building on this insight, we introduce an approach that explicitly aligns the representations of observed and executed actions within a shared latent space using contrastive learning. Experiments on action recognition and multi-task object manipulation benchmarks show that this simple method promotes synergy between the two tasks, enhancing representation quality and generalization.

We hope our work offers a novel perspective by treating action understanding and embodied execution as intertwined, rather than modeling them as isolated cognitive processes. On a broader scale, it reflects how cognition emerges from sensorimotor engagement with the environment, as emphasized by embodied cognition. Future work could also adopt more sophisticated representation learning strategies, such as hierarchical alignment, and incorporate multisensory integration to better handle complex real-world tasks. Finally, exploring aspects of social cognition could further enrich the framework by capturing interactive and cooperative dynamics.



## Acknowledgment

This work was supported by National Science and Technology Major Project (2022ZD0114904) and NSFC-6247070125.

## References

- [1] Lawrence W Barsalou. Grounded cognition. *Annu. Rev. Psychol.*, 59(1):617–645, 2008. 1
- [2] Shizhe Chen, Ricardo Garcia, Cordelia Schmid, and Ivan Laptev. Polarnet: 3d point clouds for language-guided robotic manipulation. 2023. 6
- [3] Andy Clark. *Being there: Putting brain, body, and world together again*. MIT press, 1998. 1
- [4] Giuseppe Di Pellegrino, Luciano Fadiga, Leonardo Fogassi, Vittorio Gallese, and Giacomo Rizzolatti. Understanding motor events: a neurophysiological study. *Experimental brain research*, 91:176–180, 1992. 1, 8
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 5, 2
- [6] Wolfram Erlhagen, Albert Mukovskiy, and Estela Bicho. A dynamic model for action understanding and goal-directed imitation. *Brain research*, 1083(1):174–188, 2006. 1
- [7] Luciano Fadiga, Leonardo Fogassi, Giovanni Pavesi, and Giacomo Rizzolatti. Motor facilitation during action observation: a magnetic stimulation study. *Journal of neurophysiology*, 73(6):2608–2611, 1995. 1
- [8] Vittorio Gallese and George Lakoff. The brain’s concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive neuropsychology*, 22(3-4):455–479, 2005. 1
- [9] Vittorio Gallese, Luciano Fadiga, Leonardo Fogassi, and Giacomo Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119(2):593–609, 1996. 1, 8
- [10] Valeria Gazzola and Christian Keysers. The observation and execution of actions share motor and somatosensory voxels in all tested subjects: single-subject analyses of unsmoothed fmri data. *Cerebral cortex*, 19(6):1239–1255, 2009. 8
- [11] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: 3d feature field transformers for multi-task robotic manipulation. In *7th Annual Conference on Robot Learning*, 2023. 6
- [12] Ankit Goyal, Jie Xu, Yijie Guo, Valts Blukis, Yu-Wei Chao, and Dieter Fox. Rvt: Robotic view transformer for 3d object manipulation. *CoRL*, 2023. 1, 5, 6, 2, 3
- [13] Ankit Goyal, Valts Blukis, Jie Xu, Yijie Guo, Yu-Wei Chao, and Dieter Fox. Rvt2: Learning precise manipulation from few demonstrations. *RSS*, 2024. 5, 6, 2, 3
- [14] Pierre-Louis Guhur, Shizhe Chen, Ricardo Garcia Pinel, Makarand Tapaswi, Ivan Laptev, and Cordelia Schmid. Instruction-driven history-aware policies for robotic manipulations. In *6th Annual Conference on Robot Learning*, 2022. 6
- [15] Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. Position: The platonic representation hypothesis. In *Forty-first International Conference on Machine Learning*, 2024. 3, 8
- [16] Marco Iacoboni, Roger P Woods, Marcel Brass, Harold Bekkering, John C Mazziotta, and Giacomo Rizzolatti. Cortical mechanisms of human imitation. *Science*, 286(5449):2526–2528, 1999. 8
- [17] Marco Iacoboni, Istvan Molnar-Szakacs, Vittorio Gallese, Giovanni Buccino, John C Mazziotta, and Giacomo Rizzolatti. Grasping the intentions of others with one’s own mirror neuron system. *PLoS biology*, 3(3):e79, 2005. 8
- [18] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 2020. 5, 6, 2
- [19] Stephen James, Kentaro Wada, Tristan Laidlow, and Andrew J. Davison. Coarse-to-fine q-attention: Efficient learning for visual robotic manipulation via discretisation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13739–13748, 2022. 6
- [20] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederick Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. BC-z: Zero-shot task generalization with robotic imitation learning. In *5th Annual Conference on Robot Learning*, 2021. 6
- [21] Marc Jeannerod. Neural simulation of action: a unifying mechanism for motor cognition. *Neuroimage*, 14(1):S103–S109, 2001. 1
- [22] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, pages 4904–4916. PMLR, 2021. 8
- [23] Christian Keysers and Valeria Gazzola. Expanding the mirror: vicarious activity for actions, emotions, and sensations. *Current opinion in neurobiology*, 19(6):666–671, 2009. 1, 8
- [24] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pages 3519–3529. PMLR, 2019. 8
- [25] Hangxin Liu, Chi Zhang, Yixin Zhu, Chenfanfu Jiang, and Song-Chun Zhu. Mirroring without overimitation: Learning functionally equivalent manipulation actions. In *AAAI*, 2019. 8
- [26] Jeff Loucks and Jessica A Sommerville. The role of motor experience in understanding action function: The case of the precision grasp. *Child development*, 83(3):801–809, 2012. 1
- [27] Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai yu, Liang Zhao, Yisong Wang, Jiaying Liu, and Chong Ruan. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation, 2024. 8
- [28] Mayug Maniparambil, Raiymbek Akshulakov, Yasser Abdelaziz Dahou Djilali, Mohamed El Amine Seddik, Sanath

- Narayan, Karttikeya Mangalam, and Noel E O'Connor. Do vision and language encoders represent the world similarly? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14334–14343, 2024. 8
- [29] Ari Morcos, Maithra Raghu, and Samy Bengio. Insights on representational similarity in neural networks with canonical correlation. *Advances in neural information processing systems*, 31, 2018. 8
- [30] Erhan Oztop and Michael A. Arbib. Schema design and implementation of the grasp-related mirror neuron system. *Biological Cybernetics*, 87(2):116–140, 2002. 8
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 5, 8, 2
- [32] Giacomo Rizzolatti and Laila Craighero. The mirror-neuron system. *Annu. Rev. Neurosci.*, 27(1):169–192, 2004. 1, 8
- [33] Giacomo Rizzolatti and Corrado Sinigaglia. *Mirrors in the brain: How our minds share actions and emotions*. Oxford University Press, USA, 2008. 1
- [34] E. Rohmer, S. P. N. Singh, and M. Freese. Coppeliasim (formerly v-rep): a versatile and scalable robot simulation framework. In *Proc. of The International Conference on Intelligent Robots and Systems (IROS)*, 2013. 5
- [35] M Yunus Seker, Alper Ahmetoglu, Yukie Nagai, Minoru Asada, Erhan Oztop, and Emre Ugur. Imitation and mirror systems in robots through deep modality blending networks. *Neural Networks*, 146:22–35, 2022. 8
- [36] Lawrence Shapiro and Shannon Spaulding. Embodied Cognition. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Fall 2024 edition, 2024. 1
- [37] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022. 5, 6, 2, 3
- [38] Maria Alessandra Umiltà, Evelyne Kohler, Vittorio Gallese, Leonardo Fogassi, Luciano Fadiga, Christian Keysers, and Giacomo Rizzolatti. I know what you are doing: A neurophysiological study. *Neuron*, 31(1):155–165, 2001. 8
- [39] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9 (86):2579–2605, 2008. 7, 3
- [40] Francisco J Varela, Evan Thompson, and Eleanor Rosch. *The embodied mind, revised edition: Cognitive science and human experience*. MIT press, 2017. 1
- [41] Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan Chen, Yaohui Wang, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *arXiv preprint arXiv:2307.06942*, 2023. 1, 3, 4, 5, 2
- [42] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. Representation alignment for generation: Training diffusion transformers is easier than you think. In *The Thirteenth International Conference on Learning Representations*, 2025. 8
- [43] Le Zhang, Qian Yang, and Aishwarya Agrawal. Assessing and learning alignment of unimodal vision and language models. *arXiv preprint arXiv:2412.04616*, 2024. 2, 8
- [44] Xinyu Zhang, Yuhao Liu, Haonan Chang, Liam Schramm, and Abdeslam Boularias. Autoregressive action sequence learning for robotic manipulation. *arXiv preprint arXiv:2410.03132*, 2024. 3, 4, 5, 6, 2
- [45] Junpei Zhong, Cornelius Weber, and Stefan Wermter. Robot trajectory prediction and recognition based on a computational mirror neurons model. In *International Conference on Artificial Neural Networks*, pages 333–340, 2011. 8
- [46] Wentao Zhu, Xiaoxuan Ma, Zhaoyang Liu, Libin Liu, Wayne Wu, and Yizhou Wang. Motionbert: A unified perspective on learning human motion representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 1