

FUNCanon: Learning Pose-Aware Action Primitives via Functional Object Canonicalization for Generalizable Robotic Manipulation

Hongli Xu^{2,3*}, Lei Zhang^{1,3*†}, Xiaoyue Hu^{2*}, Boyang Zhong², Kaixin Bai^{1,3},
 Zoltán-Csaba Márton³, Zhenshan Bing², Zhaopeng Chen³, Alois Christian Knoll², Jianwei Zhang¹




Fig. 1: We present **FunCanon**, a framework that transforms long-horizon manipulation into action chunks defined by actor, verb, and object. Through **functional object canonicalization** using affordance cues from vision–language models, we **automatically generate manipulation trajectories** based on **functional alignment** and train an object-centric and action-centric diffusion policies **FuncDiffuser** that generalize across object instances, categories, and tasks, enabling action primitives for generalizable robotic manipulation and robust sim-to-real transfer.

Abstract—General-purpose robotic skills from end-to-end demonstrations often leads to task-specific policies that fail to generalize beyond the training distribution. Therefore, we introduce FunCanon, a framework that converts long-horizon manipulation tasks into sequences of action chunks, each defined by an actor, verb, and object. These chunks focus policy learning on the actions themselves, rather than isolated tasks, enabling compositionality and reuse. To make policies pose-aware and category-general, we perform functional object canonicalization for functional alignment and automatic manipulation trajectory transfer, mapping objects into shared functional frames using affordance cues from large vision–language models. An object-centric and action-centric diffusion policy FuncDiffuser trained on this aligned data naturally respects object affordances and poses, simplifying learning and improving generalization ability. Experiments on simulated and real-world benchmarks demonstrate category-level generalization, cross-task behavior reuse, and robust sim-to-real deployment, showing that functional

canonicalization provides a strong inductive bias for scalable imitation learning in complex manipulation domains. Details of the demo and supplemental material are available on our project website <https://sites.google.com/view/funcanon>.

I. INTRODUCTION

As robots transition from controlled laboratory settings to unstructured real-world environments, developing robust and generalizable manipulation policies becomes increasingly critical. A fundamental challenge is enabling agents to generalize across unseen objects, diverse poses, and varying tasks — a capability that remains elusive for current imitation learning approaches.

Imitation learning methods based on RGB images [1] or point clouds [2] often suffer from limited precision and generalization due to viewpoint sensitivity, noisy observations, and redundant scene encodings. In contrast, 3D scene representations have shown promise in improving generalization [3]. To further address these challenges, object-centric representations, which focus on structured, object-level information such as 6D poses [4, 5] and scene flow [6], have gained significant attention. SPOT [4] demonstrates that SE(3) pose diffusion policy can improve cross-embodiment generalization, even

* The first three authors contribute equally to this paper.

†Corresponding author. zhanglei.cn.de@gmail.com, lei.zhang-1@studium.uni-hamburg.de

¹TAMS (Technical Aspects of Multimodal Systems), Department of Informatics, University of Hamburg, Hamburg, Germany.

²Technical University of Munich, Germany.

³Agile Robots SE, Munich, Germany.

This work is supported by New Generation Artificial Intelligence-National Science and Technology Major Project (2025ZD0122903).

when trained solely on passive human videos. Existing object-centric approaches [4] often depend on instance-specific, goal-conditioned trajectories, limiting generalization across categories and tasks. We attribute this to viewing manipulation as monolithic programs rather than modular, reusable behaviors. In the field of computer vision, prior work, such as UAD [7] and Object Canonicalization [8], mainly targets improving visual representations and semantic understanding. Related efforts have also investigated category-level affordance pose estimation [9]. However, these methods have not explored how such representations can be leveraged for improving category-level alignment of robotic manipulation data, or augmenting manipulation trajectories.

Key open questions remain: *how can generalized representations be leveraged to synthesize diverse manipulation data, model long-horizon tasks, and train robust, generalizable manipulation policies?*

To address these challenges, we propose **FunCanon**, a framework that models manipulation as compositions of reusable *action primitives*—such as pouring, grasping, or inserting—defined over functionally aligned bi-object interactions. By leveraging affordance cues from large vision-language models (VLMs), FunCanon canonicalizes semantically related objects (e.g., kettles and pitchers) into shared functional frames. This functional alignment enables the automatic manipulation trajectory transfer and the training of *pose-aware, object-centric diffusion policies* that focus solely on interaction dynamics, decoupled from specific object identities, camera viewpoints, or task semantics. To achieve robust and generalizable manipulation, we first decompose long-horizon tasks into meaningful *action chunks*, each specifying an actor, an action, and an object. This segmentation is performed by a large multimodal language model (MLLM), such as GPT-4o, in combination with large vision models (LVM) that extract affordance cues to guide chunking based on functional relevance. Next, we integrate these affordance cues with precise object pose estimates to perform *functional alignment*, canonicalizing objects into shared functional frames. This process identifies action-related affordance regions and aligns bi-object poses, producing a semantically grounded representation of manipulation interactions. Leveraging this functional alignment, automatic trajectory transfer method is proposed to augment training data on RL Bench base tasks to increase data diversity and functional coverage. During policy training and inference, our object-centric diffusion policy receives inputs encoding both affordances, poses of bi-object pairs, point clouds and action verb and estimate actions. Our main contributions are:

- Introducing **FunCanon**, which decomposes complex manipulation tasks into reusable *action primitives* grounded in functionally aligned object pairs. We explicitly incorporate the *grasping phase* within action primitives, addressing a key gap in prior approaches.
- Leveraging large vision-language models for **affordance-driven functional canonicalization** for **functional alignment** and **automatic trajectory transfer**, enabling pose-

aware and category-generalizable policy learning.

- Developing an object-centric diffusion policy trained on functionally aligned data, achieving both instance-level and category-level generalization and robust sim-to-real transfer.

II. RELATED WORK

A. Generalizable Robotic Manipulation

End-to-end visuomotor policies, trained via imitation [1, 2] or reinforcement learning [10], have laid the groundwork for robotic manipulation [11]. While effective on single, well-defined tasks, these approaches [1, 2] often face difficulties when scaling to long-horizon tasks, multimodal action distributions. To overcome these challenges, recent works incorporate structured priors [12] or generative models [3] to improve robustness and generalization. For example, diffusion-based methods [3, 13] have shown promising results in multi-task skill learning. The 3D Diffuser Actor (3DA) [3] further advances this line by conditioning policies on 3D scene representations, achieving better viewpoint robustness and sim-to-real transfer. Prior works typically achieve sim-to-real transfer through domain randomization or domain adaptation techniques to mitigate the domain gap between simulation and the real world [14, 15, 16, 17]. Despite these advances, these methods still tend to treat manipulation as monolithic, task-specific trajectory generation, limiting their flexibility and composability.

In contrast, object-centric representations focus on disentangling robot embodiment from task semantics by grounding policies in object frames. Techniques like CLIPort [18] and VIMA [19] leverage language and vision to localize manipulation actions relative to object-centric coordinates. Category-level representations, such as Neural Descriptor Fields [20] and keypoints [21], enhance generalization by capturing shared object properties. However, many object-centric approaches [4] remain tightly coupled to task-specific goals and trajectories, hindering reusability and broad generalization.

B. Affordance Reasoning through LLM and VLM

Recently, VLMs and LLMs have been leveraged to reason about object affordances and functional parts [7, 22, 23, 24] and plan functional robotic manipulation [25, 26]. UAD [7], distill affordance priors from foundation models to generalize without manual labeling. Geometry-semantic alignment methods like One-Shot 3D Canonicalization [8] further advance functional region understanding. While these advances demonstrate the potential of language-grounded affordance reasoning, most methods have yet to integrate these insights tightly with policy learning for generalizable robotic manipulation, increasing the instance-level and category-level generalization abilities.

III. PROBLEM STATEMENT AND METHODS

Robotic manipulation in unstructured environments requires policies that generalize across diverse objects, poses, and tasks. Our goal is to learn generalizable manipulation policies that enable trajectory transfer across object instances without

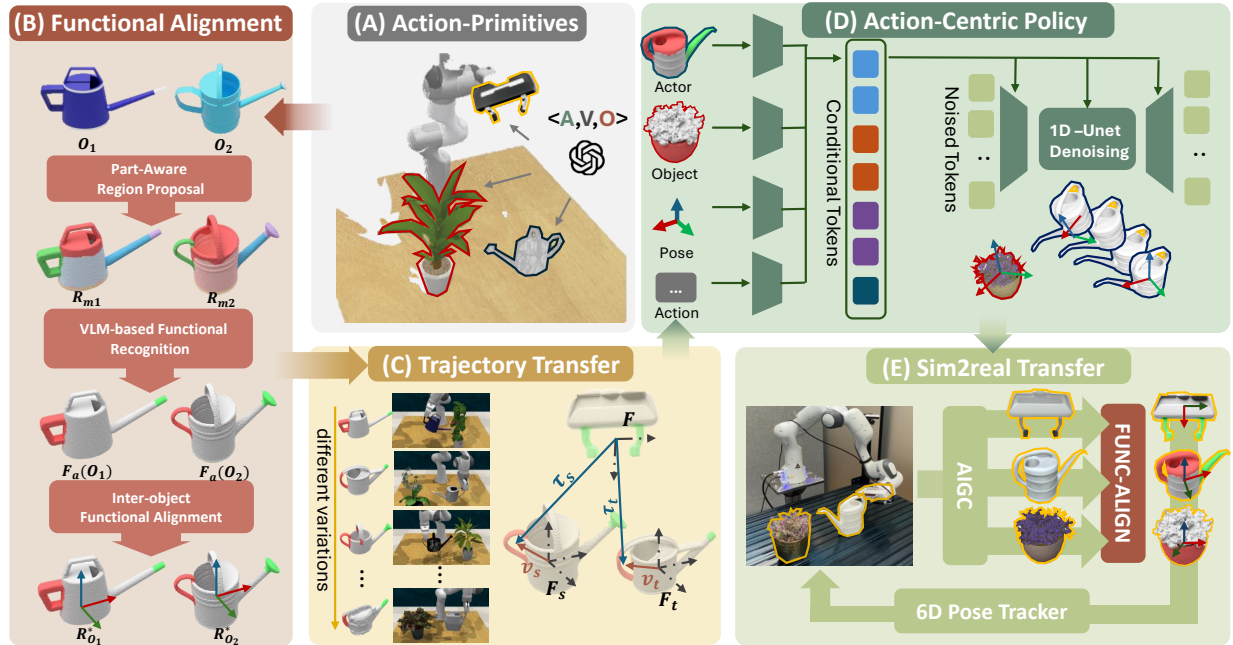


Fig. 2: FunCanon Architecture. Our method consists of five components: (A) **Action Primitive Modeling**, where vision-language models perform AVO segmentation and task decomposition; (B) **Functional Alignment**, producing function-based canonicalization of objects; (C) **Automatic Trajectory Transfer**, which augments training data across instances and categories using online 3D datasets; (D) **Action-Centric Diffusion Policy learning** with FuncDiffuser, trained on this augmented data to predict AVO-centric actions; and (E) **Sim-to-Real Inference**, enabling robust transfer of the learned policy to real-world scenarios.

requiring additional human supervision. To this end, we propose an unified framework that integrates five key components: (A) Action primitive modeling and long-horizon task decomposition, (B) Functional alignment via vision-language and large language models, (C) Automatic and generalizable trajectory transfer, (D) Object-centric and action-centric diffusion policy learning, and (E) Sim-to-real inference for robust deployment on physical robots.

A. Action Primitives and Task Decomposition

We model manipulation as sequences of reusable action primitives. Formally, an action primitive is defined as a triplet $a = \langle \mathcal{A}, \mathcal{V}, \mathcal{O} \rangle$, where \mathcal{A} represents the actor (e.g., kettle, pitcher), \mathcal{V} denotes the action verb (e.g., grasp, pour), and \mathcal{O} specifies the target object (e.g., cup, bowl). A long-horizon manipulation task can then be expressed as a sequence of action primitives $\mathcal{T} = \{a_1, a_2, \dots, a_N\}$.

We leverage large-scale multimodal language models \mathcal{M} combined with semantic and functional cues extracted from vision models \mathcal{V} to automatically perform task decomposition. This step not only focuses action learning on individual action semantics and object interactions but also facilitates subsequent action reuse and composition.

$$a_1, a_2, \dots, a_N = \mathcal{D}(\mathcal{T}_e, \mathcal{M}, \mathcal{V}) \quad (1)$$

where \mathcal{D} represents the decomposition function that takes task demonstrations \mathcal{T}_e and models \mathcal{M} and \mathcal{V} as input, outputting action primitive sequences.

With this decomposition, our original goal of learning generalizable manipulation policies is reformulated into the problem of learning robust AVO-based action primitives. By treating each action primitive as a transferable motion segment, the challenge of generalization is shifted from entire trajectories to the representation and reuse of AVO triplets across objects and tasks.

B. Functional Alignment

Despite vast differences in appearance and category, different objects may assume similar operational functions. To achieve cross-category generalization of action primitives, we propose a functional alignment mechanism that unifies functional pose expressions across different objects. The key components contain of part-aware region proposal, VLM-based functional recognition and inter-object functional alignment.

1) Part-aware Region Proposal

Given an object o with its 3D model M_o , we first perform RGB-D rendering from multiple viewpoints to obtain corresponding RGB images and depth maps. Based on [7], we use a self-supervised visual feature encoder [27] to extract features from each view, and these 2D features are lifted to 3D space using the depth information, forming feature representations for candidate functional regions. We then cluster these features using KMeans to obtain M functional region candidates, each represented as:

$$R_m = \{(x_i, f_i) \mid i \in C_m\}, \quad (2)$$

where x_i denotes the 3D coordinates of a point, f_i denotes the corresponding visual feature, and C_m denotes the set of point indices belonging to the m -th region. Visualization of the feature space and clustering results are provided in the supplementary material.

2) VLM-based Functional Recognition

For each candidate functional region R_m , we extract its visual features V_m and combine them with the object category information c_o . We then define a multimodal large language model as a *binary classifier*:

$$\Phi(R_m, a, r, c_o) \rightarrow \{\text{True}, \text{False}\}, \quad (3)$$

where a denotes an action (e.g., *grasp*, *pour*), and r denotes a role hypothesis (active or passive). The function Φ evaluates whether region R_m is functionally relevant to performing action a under role r , conditioned on the object category c_o .

The functional region set for object o under action a is then defined as:

$$F_a(o) = \{(R_m, a, r) \mid \Phi(R_m, a, r, c_o) = \text{True}\}. \quad (4)$$

This formulation explicitly links candidate regions to action–role pairs by verifying their functional relevance through binary decisions. For example, for a kettle, the handle region may be evaluated with the hypothesis (*grasp*, *active*), which yields *True*, while the same region under (*pour*, *passive*) yields *False*. As a result, the handle is explicitly associated with the *grasp*–*active* primitive but excluded from the *pour*–*passive* mapping.

3) Inter-object Functional Alignment

After obtaining the functional region set $F_a(o)$ for each object, we leverage these regions to achieve functional alignment between different objects, ensuring that action primitives can be transferred reliably across object instances. We assume that all object models have been normalized and centered, so that each object is consistently scaled and positioned at the origin.

For each object, we define a *functional direction vector* that captures the overall spatial location of its functional regions:

$$v_o = \frac{1}{|F_a(o)|} \sum_{(x_i, f_i) \in F_a(o)} x_i, \quad (5)$$

where x_i denotes the 3D coordinate of a point in a functional region. Intuitively, v_o represents the “functional center” of the object, serving as a reference for cross-object alignment.

Given a source object o_s and a target object o_t , with functional vectors v_s and v_t , respectively, we aim to align them by optimizing the object orientation. Specifically, the alignment problem is formulated as:

$$R^* = \arg \min_{R \in SO(3)} \|Rv_s - v_t\|_2^2. \quad (6)$$

Since most object models are aligned along the Z-axis, the only remaining rotational degree of freedom is the heading direction, i.e., rotation around Z. Therefore, finding the optimal rotation can be simplified to searching for a single angle around the Z-axis that best aligns the source and target vectors.

This formulation provides a lightweight yet effective strategy for functionally consistent cross-object alignment, enabling

robust transfer of action primitives.

C. Automatic and Generalizable Trajectory Transfer

With functional alignment, we can leverage RL Bench demonstrations for effective data augmentation. In RL Bench, a complete robotic trajectory is typically composed of multiple sub-trajectories, where each sub-trajectory is generated by planning from a start pose to an end pose defined by relative constraints between a source object and a target object.

Crucially, this decomposition naturally aligns with our AVO-based action primitives: each sub-trajectory corresponds to an action involving a specific actor, verb, and object. We exploit this alignment to augment training data as follows. Given a sub-trajectory τ_s associated with a source object o_s , we first transform the object coordinate system into its functional coordinate system defined by the functional vector v_s . Within this functional frame, the relative pose between the source and target objects is preserved, and the sub-trajectory is represented accordingly. To transfer the trajectory to a new object o_t , we simply apply the corresponding functional vector v_t to map the trajectory back into the target object’s coordinate frame.

Formally, the transferred trajectory τ_t is obtained as:

$$\tau_t = \tau_s + (v_t - v_s), \quad (7)$$

where the addition denotes transforming the trajectory from the source functional frame to the target functional frame. By repeating this process across objects within a category, we can generate multiple variations of each action primitive, substantially increasing the diversity of training data while ensuring functional consistency. This enables the model to learn robust, generalizable AVO-based policies across different object instances and categories.

D. Action-centric Policy Training

We train a diffusion-based policy, named FuncDiffuser, to generate future object trajectories corresponding to action primitives, rather than directly predicting robot joint motions. Our state representation encodes the relative pose between actor and object as pose feature, the point clouds with associated functional areas as instance feature, and a CLIP-encoded embedding of the action verb to guide the motion. Formally, the state is defined as

$$s = (h_\Delta, f_A, f_O, v), \quad (8)$$

where h_Δ denotes the *relative pose feature* between the actor and object, f_A, f_O are their visual functional features from PointNet++, and v is the verb embedding. The policy $\pi_\theta(a \mid s)$ is trained using a diffusion-based denoising loss (MSE) to align predicted action trajectories with demonstration trajectories.

To enhance generalization, we leverage functional alignment and relative-pose features to perform *trajectory augmentation*. RL Bench trajectories are first mapped into their functional frames, and the relative poses are transferred across objects. This process generates multiple variations of each action primitive, enabling the diffusion policy to naturally respect

object affordances and generalize across object instances and categories.

E. Sim-to-Real Inference

To enable robust deployment of our action-centric policies, we design a sim-to-real inference pipeline based on functional alignment. High-level human instructions are first processed by a large language model (GPT-4o), which decomposes the task into sub-tasks with actor-object pairs. Detailed decompositions examples are shown in appendix material.

For each target object, a 3D mesh is generated using an AIGC model [28] to compensate for the lack of high-quality CAD models. Functional alignment is then applied to perform segmentation and alignment on the generated mesh, identifying regions relevant to the intended action.

Using these aligned functional regions, an off-the-shelf pose estimator [29] produces an object-centric scene representation, including the poses and spatial relations of functional regions. This representation, together with the planned actor-object pairs from task decomposition, is fed into the action-centric diffusion policy, which predicts trajectories for each sub-task. This pipeline enables accurate trajectory prediction and robust sim-to-r

IV. EXPERIMENT

A. Experiments Setup

Simulation Environment: Experiments are conducted in RL Bench [30], a standard robotic manipulation benchmark providing diverse object categories, tasks, and demonstrations. This allows systematic evaluation of generalization under **pose-level**, **instance-level**, and **category-level** variations.

Real-world Environment: The real-world setup consists of a Franka Emika robot and multiple Intel RealSense cameras, as illustrated in Fig. 3. This setup enables evaluation of sim-to-real transfer of our manipulation policies.

B. Simulation Experiments: Data Generation

We conduct all **pose-level**, **instance-level** and **category-level** manipulation experiments in a simulation environment based on RL Bench [30]. During data collection, we monitor trajectory execution to annotate sub-task ground truth for each manipulation behavior, and record the corresponding **actor-manipulated object pairs**. We leverage **functional alignment** to automatically transfer trajectories between objects with similar functional roles. By aligning trajectories based on **functional alignment** of objects, our method generates new trajectory variants that preserve functional behavior rather than just low-level attributes like position or color.

Task and dataset Variation: To systematically evaluate generalization ability, we consider three levels of task variation:

Pose-level variation: Using the original RL Bench instances, tasks are executed with random initial poses of the objects and/or robot end-effector to test robustness to configuration changes.




Fig. 3: Real-World Experiment Setup with Franka Emika Robot and Intel RealSense Cameras.




Fig. 4: Task Variants in Simulation with Generated Trajectories based on our Functional Alignment method.

Instance-level variation: objects within the same category are replaced with various shapes, sizes, or textures. All instance-level objects are sampled from Objaverse [31].

Category-level variation: objects are substituted with different but related categories or subcategories, introducing variations in geometry and functional parts (e.g., watering can vs. teapot). Fig. 5 shows examples of functional segmentation used to guide trajectory transfer and Fig. 4 shows the generated trajectories for different variants.

C. Implementation Details

The pipeline of real-world experiments consists of three main components: LLM-based task decomposition, pose estimation, and policy inference. We employ GPT-4O for task decomposition, where the LLM takes human instructions and breaks down the overall task into sub-tasks with a list of actor-object pairs. We model the manipulation pipeline with multiple stages with different types of action primitives. Then, we leverage FoundationPose [29] for pose estimation. For our diffusion policy network, the state encoder is a lightweight three-layer MLP that encodes the pose into a 64-dimensional representation. We additionally use a pretrained CLIP model to process the language instructions into sentence embeddings. The diffusion module follows a UNet-style backbone, optimized with AdamW at a learning rate of 1e-4. A DDIM scheduler with

TABLE I: Success rate (%) in simulation under pose, instance-level, and category-level substitutions across different methods (25 trials, three repetitions). Numbers indicate mean \pm std. Avg row shows improvement over 3DA in green. Evaluation of 25 episodes using 3 random seeds for each task.

Task	Ours (Funcanon)			SPOT			3DA		
	Pose	Inst.	Cat.	Pose	Inst.	Cat.	Pose	Inst.	Cat.
T1	60.0 \pm 3.27	61.3 \pm 3.77	60.0 \pm 3.27	48.0 \pm 3.27	50.7 \pm 4.99	48.0 \pm 3.27	28.0 \pm 3.27	33.3 \pm 1.89	32.0 \pm 3.27
T2	77.3 \pm 4.99	80.0 \pm 3.27	72.0 \pm 3.27	68.0 \pm 3.27	64.0 \pm 3.26	70.7 \pm 1.89	24.0 \pm 3.27	25.3 \pm 1.89	22.7 \pm 1.89
T3	80.0 \pm 3.27	84.0 \pm 3.27	72.0 \pm 3.27	68.0 \pm 3.27	72.0 \pm 3.27	64.0 \pm 3.27	36.0 \pm 3.27	40.0 \pm 3.27	32.0 \pm 3.27
Avg.	72.4 (+43.1)	75.1 (+42.2)	68.0 (+39.1)	61.3 (+32.0)	62.2 (+29.3)	60.9 (+32.0)	29.3	32.9	28.9

100 timesteps during training and 10 at inference is employed. Training runs for 5000 epochs with a batch size of 64, ensuring stable convergence across tasks.

Based on the recognized poses and the planned actor-object pairs, the policy module infers the robot actions. Please refer to Supplementary Material for more implementation details. An example of our task decomposition method is shown as follows.

Context:

Consider you are a 6-DOF robotic arm. I will give you a daily-life task. Decompose this task into a sequence of **core, action-centered steps**.

Examples: { "task": "pour water", "steps": [{ "step": 1, "action": "grasp", "actor": "robot gripper", "object": "teapot" }, { "step": 2, "action": "pour", "actor": "teapot", "object": "cup" }] }

Output Format:

```
{ "task": "<task description>",
  "steps": [ {
    "step": 1,
    "action": "<primitive action>",
    "actor": "<agent performing the action,
              ↪ e.g., robot gripper or tool>",
    "object": "<object being manipulated or
              ↪ affected>"
  },
  ... ] }
```

Instruction: Stuck red cup into blue cup, then stuck green cup into blue cup.

Results:

```
{ "task":
  "stuck red cup into blue cup then stuck green
  ↪ cup into blue cup",
  "steps": [
    { "step": 1,
      "action": "grasp",
      "actor": "robot gripper",
      "object": "red cup" },
    { "step": 2,
      "action": "insert",
      "actor": "red cup",
      "object": "blue cup" }
    ... ]
}
```

D. Simulation Experiments: Policy Training Results

We evaluate our policy’s generalization through simulation experiments at both instance and category levels. For instance-level generalization, we train on a subset of objects within each category and validate on unseen instances from the same

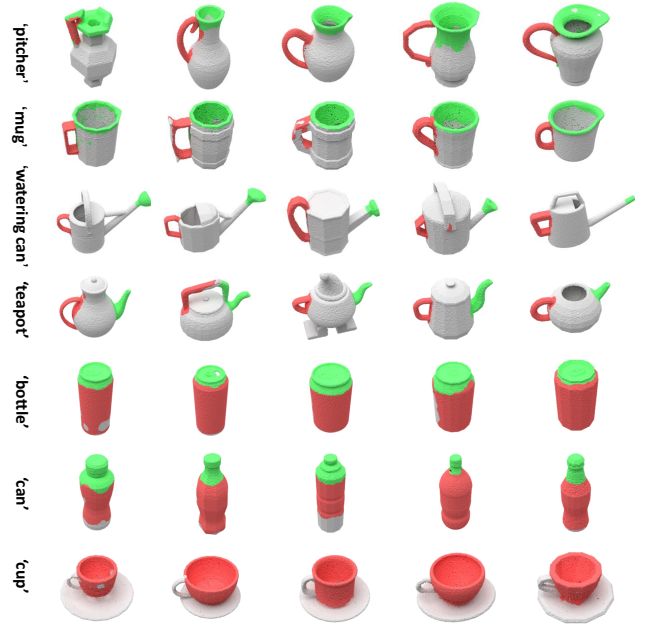


Fig. 5: Qualitative Results of Functional Alignment.

categories. For category-level generalization, training involves a few categories with representative instances, while validation includes unseen instances and novel categories to assess cross-category generalization.

Evaluation Metrics: Success Rate (SR) measures the overall task completion, and Sub-task Success Rate (**Sub-task SR**) evaluates the completion of individual steps within each task.

Tasks: We evaluate our policy on three core manipulation tasks: Put A in B, Pour A in B, and Water B with A. For trajectory transfer, we adopt the corresponding RLBench demonstrations from Put Groceries in Cupboard, Pour from Cup to Cup, and Water Plants. Specifically, Put A in B maps to the Pick & Place task, Pour A in B to Pick, Pour (Level 1), and Water B with A to Pick, Pour (Level 2). The key difference between the two pouring tasks is the container design: Pour A in B uses a generic cup without a spout, while Water B with A employs a watering can whose spout directs the liquid flow. We compare our method with several SOTA baseline approaches, including SPOT [4], 3DA [3], under all pose-level instance-level and category-level variations.

Results and Analysis: The results are summarized in

TABLE II: Policy training of Pick, Pour (L2) under category-level substitutions (25 trials, one variation). SR: Success Rate; Sub-SR: success rate of each sub-task. Sub2-SR = Sub2 success / Sub1 success. Numbers in parentheses indicate performance drop relative to Ours (FunCanon).

Method	Sub1-SR (%)	Sub2-SR (%)	SR (%)
Ours (FunCanon)	64.0	62.5	40.0
Geom-only	60.0	60.0	36.0 (-4.0)
No-Func-Frame	44.0	54.5	24.0 (-16.0)
Monolithic	36.0	44.4	16.0 (-24.0)

Tab. I. The results demonstrate that our policy achieves high success rates across tasks and object variations, consistently outperforming existing methods, and effectively generalizing to unseen instances and categories.

Our method outperforms SPOT and 3DA for several reasons. First, all methods share the same base instance, ensuring a fair comparison. Second, both our method and SPOT leverage pose-aware object policies, which gives them an advantage over 3DA. However, our evaluation is more challenging than SPOT, as we assess the complete manipulation pipeline starting from grasping, whereas SPOT begins from the post-grasp stage. Notably, the gap between instance-level and category-level success rates is larger for SPOT and 3DA, whereas our approach maintains high performance. This improvement is largely due to our functional alignment module, which enhances generalization to unseen object instances and categories. In contrast, SPOT relies on per-instance tracking and struggles to generalize to objects with novel coordinate conventions.

E. Ablation Study

To analyze the contribution of key components in our policy, we perform an ablation study focusing on three main variants: (A) **Monolithic Trajectory**, (B) **No Functional Frame**, and (C) **Geometry-only**. Monolithic Trajectory removes multi-stage modeling by merging all stages into a single trajectory. No Functional Frame uses only world or camera coordinates, omitting functional keypoints. Geometry-only relies solely on geometric information, ignoring affordance and functional cues.

The results show that each component plays a crucial role in policy performance. Monolithic Trajectory performs poorly on complex multi-stage tasks, demonstrating the necessity of stage modeling. No Functional Frame reduces the success rate of grasping and manipulation steps, highlighting the importance of functional keypoints for pose-aware execution. Geometry-only also lowers generalization to unseen instances and categories, confirming the value of affordance and functional information. Overall, these findings validate that stage modeling, functional frames, and functional cues are essential for achieving high Success Rate (SR) and Sub-task Success Rate (Sub-task SR).

F. Real-World Experiments (SIM2REAL Transfer)

To validate the generalization of our policy, we conduct SIM2REAL transfer experiments in the real world. The policy is trained entirely in simulation and deployed directly on the




Fig. 6: Qualitative Results of Real-World Experiments.

TABLE III: Real World Experiment Results.

Methods	Pick, Place	Pick, Pour (L1)	Pick, Pour (L2)
3DA [3]	48%	54%	62%
SPOT [4]	52%	76%	60%
FunCanon (Ours)	88%	90%	88%

real robot without any fine-tuning. We conducted 50 real-world trials with unknown objects, as shown in Fig. IV. The success rates are reported in Tab. III. Our method consistently outperforms the baselines. Compared with SPOT, FunCanon improves success by at least +14%. The advantage is even more pronounced over 3DA, with improvements of at least +26%. Notably, the performance gap is largest in the pouring tasks, where functional alignment provides a strong inductive bias for handling different container geometries. These results validate that FunCanon generalizes more robustly across object variations and achieves reliable sim-to-real transfer.

V. CONCLUSION AND FUTURE WORK

In this work, we introduced **FunCanon**, a framework that bridges object-centric and action-centric learning via **functional object canonicalization**. By converting long-horizon manipulation into modular *actor-verb-object primitives* and aligning objects into functional frames, our method supports **automatic trajectory transfer** and policy learning that generalizes across objects, categories, and tasks. Simulation and real-world experiments validate that FunCanon substantially outperforms SOTA baselines and enables robust sim-to-real transfer. In simulation, our success rate improves by +14.6% over SPOT and +38.4% over 3DA. In real-world trials, FunCanon consistently outperforms baselines, surpassing SPOT by at least +14% and 3DA by at least +26% across all tasks. The largest improvements appear in pouring tasks, where functional alignment enables robust handling of container geometries and spout orientations. These results highlight that FunCanon not only achieves higher success rates but also delivers stronger instance- and category-level generalization through functional canonicalization.

In future work, we plan to extend FunCanon to articulated object manipulation, and explore scalable data generation

pipelines that further enhance cross-domain generalization.

REFERENCES

- [1] 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 International Journal of Robotics Research*, p. 02783649241273668, 2023. 1, 2
- [2] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, "3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations," *arXiv preprint arXiv:2403.03954*, 2024. 1, 2
- [3] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki, "3d diffuser actor: Policy diffusion with 3d scene representations," *arXiv preprint arXiv:2402.10885*, 2024. 1, 2, 6, 7
- [4] C.-C. Hsu, B. Wen, J. Xu, Y. Narang, X. Wang, Y. Zhu, J. Biswas, and S. Birchfield, "Spot: Se (3) pose trajectory diffusion for object-centric manipulation," *arXiv preprint arXiv:2411.00965*, 2024. 1, 2, 6, 7
- [5] C. Wang, R. Martín-Martín, D. Xu, J. Lv, C. Lu, L. Fei-Fei, S. Savarese, and Y. Zhu, "6-pack: Category-level 6d pose tracker with anchor-based keypoints," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 059–10 066. 1
- [6] Z. Xu, Z. He, J. Wu, and S. Song, "Learning 3d dynamic scene representations for robot manipulation," *arXiv preprint arXiv:2011.01968*, 2020. 1
- [7] Y. Tang, W. Huang, Y. Wang, C. Li, R. Yuan, R. Zhang, J. Wu, and L. Fei-Fei, "Uad: Unsupervised affordance distillation for generalization in robotic manipulation," *arXiv preprint arXiv:2506.09284*, 2025. 2, 3
- [8] L. Jin, Y. Wang, W. Chen, Q. Dai, Q. Gao, X. Qin, and B. Chen, "One-shot 3d object canonicalization based on geometric and semantic consistency," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 16 850–16 859. 2
- [9] Y. Wang, L. Zhang, Y. Tu, H. Zhang, K. Bai, Z. Chen, and J. Zhang, "Tooleenet: Tool affordance 6d pose estimation," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 10 519–10 526. 2
- [10] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396. 2
- [11] L. Zhang, K. Bai, G. Huang, Z. Bing, Z. Chen, A. Knoll, and J. Zhang, "Contactdextnet: Multi-fingered robotic hand grasping in cluttered environments through hand-object contact semantic mapping," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025, pp. 8495–8502. 2
- [12] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," in *Conference on Robot Learning*. PMLR, 2023, pp. 785–799. 2
- [13] U. A. Mishra, S. Xue, Y. Chen, and D. Xu, "Generative skill chaining: Long-horizon skill planning with diffusion models," in *Conference on Robot Learning*. PMLR, 2023, pp. 2905–2925. 2
- [14] K. Bai, L. Zhang, Z. Chen, F. Wan, and J. Zhang, "Close the sim2real gap via physically-based structured light synthetic data simulation," in *2024 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2024, pp. 17 035–17 041. 2
- [15] L. Zhang, K. Bai, Q. Li, Z. Chen, and J. Zhang, "A collision-aware cable grasping method in cluttered environment," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 2126–2132. 2
- [16] Q. Wang, K. Bai, L. Zhang, Q. Li, A. Knoll, J. Zhang, Y. Ying, and M. Zhou, "Toward collision-aware robotic fragile fruit grasping: A sim-to-real framework for perception, reasoning, and execution," *IEEE/ASME Transactions on Mechatronics*, 2025. 2
- [17] K. Bai, H. Zeng, L. Zhang, Y. Liu, H. Xu, Z. Chen, and J. Zhang, "Cleardepth: enhanced stereo perception of transparent objects for robotic manipulation," *arXiv preprint arXiv:2409.08926*, 2024. 2
- [18] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on robot learning*. PMLR, 2022, pp. 894–906. 2
- [19] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan, "Vima: Robot manipulation with multimodal prompts," 2023. 2
- [20] A. Simeonov, Y. Du, A. Tagliasacchi, J. B. Tenenbaum, A. Rodriguez, P. Agrawal, and V. Sitzmann, "Neural descriptor fields: Se (3)-equivariant object representations for manipulation," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6394–6400. 2
- [21] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei, "Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation," *arXiv preprint arXiv:2409.01652*, 2024. 2
- [22] H. Chu, X. Deng, Q. Lv, X. Chen, Y. Li, J. Hao, and L. Nie, "3d-affordancellm: Harnessing large language models for open-vocabulary affordance detection in 3d worlds," *arXiv preprint arXiv:2502.20041*, 2025. 2
- [23] H. Zhu, Q. Kong, K. Xu, X. Xia, B. Deng, J. Ye, R. Xiong, and Y. Wang, "Grounding 3d object affordance with language instructions, visual observations and interactions," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 17 337–17 346. 2
- [24] S. Qian, W. Chen, M. Bai, X. Zhou, Z. Tu, and L. E. Li, "Affordancellm: Grounding affordance from vision language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 7587–7597. 2
- [25] D. Li, C. Zhao, S. Yang, L. Ma, Y. Li, and W. Zhang, "Learning instruction-guided manipulation affordance via large models for embodied robotic tasks," in *2024 International Conference on Advanced Robotics and Mechatronics (ICARM)*. IEEE, 2024, pp. 662–667. 2
- [26] L. Zhang, D. Zheng, K. Bai, Z. Bing, Z.-C. Marton, Z. Chen, A. C. Knoll, and J. Zhang, "Omnidexvlg: Learning dexterous grasp generation from vision language model-guided grasp semantics, taxonomy and functional affordance," *arXiv preprint arXiv:2512.03874*, 2025. 2
- [27] M. Oquab, T. Darcet, T. Moutakanni, H. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby et al., "Dinov2: Learning robust visual features without supervision," *arXiv preprint arXiv:2304.07193*, 2023. 3
- [28] D. Tochilkin, D. Pankratz, Z. Liu, Z. Huang, A. Letts, Y. Li, D. Liang, C. Laforte, V. Jampani, and Y.-P. Cao, "Triposr: Fast 3d object reconstruction from a single image," *arXiv preprint arXiv:2403.02151*, 2024. 5
- [29] B. Wen, W. Yang, J. Kautz, and S. Birchfield, "Foundationpose: Unified 6d pose estimation and tracking of novel objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 17 868–17 879. 5
- [30] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, "Rlbench: The robot learning benchmark & learning environment," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3019–3026, 2020. 5
- [31] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi, "Objaverse: A universe of annotated 3d objects," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 13 142–13 153. 5