Novelty Assessment Report

Paper: Abstracting Robot Manipulation Skills via Mixture-of-Experts Diffusion Policies
PDF URL: https://openreview.net/pdf?id=VSWjHIveqZ
Venue: ICLR 2026 Conference Submission
Year: 2026
Report Generated: 2026-01-01

Abstract

Diffusion-based policies have recently shown strong results in robot manipulation, but their extension to multi-task scenarios is hindered by the high cost of scaling model size and demonstrations. We introduce Skill Mixture-of-Experts Policy (SMP), a diffusion-based mixture-of-experts policy that learns a compact orthogonal skill basis and uses sticky routing to compose actions from a small, task-relevant subset of experts at each step. A variational training objective supports this design, and adaptive expert activation at inference yields fast sampling without oversized backbones. We validate SMP in simulation and on a real dual-arm platform with multi-task learning and transfer learning tasks, where SMP achieves higher success rates and markedly lower inference cost than large diffusion baselines. These results indicate a practical path toward scalable, transferable multi-task manipulation: learn reusable skills once, activate only what is needed, and adapt quickly when tasks change.

Core Task Landscape

This paper addresses: **Multi-Task Robot Manipulation with Diffusion-Based Mixture-of-Experts Policies**

A total of **21 papers** were analyzed and organized into a taxonomy with **18 categories**.

Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Diffusion Policy Architectures for Multi-Task Manipulation**
- **Flow-Matching Alternatives to Diffusion Policies**
- **Representation Learning for Multi-Task Manipulation**
- **Dexterous Manipulation with Diffusion Policies**
- **Residual Learning with Mixture-of-Experts**
- **Multi-Task Reinforcement Learning with MoE**
- **Locomotion with Mixture-of-Experts**
- **Diffusion Policies for Hybrid Dynamical Systems**
- **Skill Composition and Specialization with MoE**
- **Vision-Language-Action Models with MoE**
- ... and 3 more categories

Complete Taxonomy Tree

- Multi-Task Robot Manipulation with Diffusion-Based Mixture-of-Experts Policies Survey Taxonomy
- Diffusion Policy Architectures for Multi-Task Manipulation
  - Mixture-of-Experts Integration in Diffusion Policies
  - Skill-Based MoE Diffusion Policies ★ (2 papers)
    - [0] Abstracting Robot Manipulation Skills via Mixture-of-Experts Diffusion Policies (Anon et al., 2026) View paper
    - [14] MoE-DP: An MoE-Enhanced Diffusion Policy for Robust Long-Horizon Robotic Manipulation with Skill Decomposition and Failure Recovery (Baiye Cheng, 2025) View paper
  - Denoiser-Level MoE for Diffusion Transformers (1 papers)
    - [20] Efficient and Scalable Diffusion Transformer Policies with Mixture of Expert Denoisers for Multitask Learning (M Reuss, n.d.) View paper
  - Language-Conditioned MoE Diffusion Policies (1 papers)
    - [4] Language-Conditioned Representations and Mixture-of-Experts Policy for Robust Multi-Task Robotic Manipulation (Zhang Xiucheng, 2025) View paper
  - Sparse Diffusion Policies (1 papers)
  - [3] Sparse diffusion policy: A sparse, reusable, and flexible policy for robot learning (Wang, 2024) View paper
  - Distillation Methods for Diffusion Policies (1 papers)
  - [2] Variational distillation of diffusion policies into mixture of experts (Denis Blessing, 2024) View paper
- Flow-Matching Alternatives to Diffusion Policies (1 papers)
  - [7] VFP: Variational Flow-Matching Policy for Multi-Modal Robot Manipulation (Xuanran Zhai, 2025) View paper
- Representation Learning for Multi-Task Manipulation (1 papers)
  - [8] Bridging Perception and Action: Spatially-Grounded Mid-Level Representations for Robot Generalization (Yang, 2025) View paper
- Dexterous Manipulation with Diffusion Policies
  - Functional Pre-Grasp Manipulation (2 papers)
  - [5] Dexterous Functional Pre-Grasp Manipulation with Diffusion Policy (Wu, 2024) View paper

- Residual Learning with Mixture-of-Experts (1 papers)
  - [12] Efficient Residual Learning with Mixture-of-Experts for Universal Dexterous Grasping (YUAN HaoQi, 2024) View paper
- Multi-Task Reinforcement Learning with MoE
  - Attention-Based MoE for Multi-Task RL (1 papers)
  - [10] Multi-Task Reinforcement Learning With Attention-Based Mixture of Experts (Guangran Cheng, 2023) View paper
  - MoE for Hybrid Dynamical Systems (1 papers)
  - [15] SAC-MoE: Reinforcement Learning with Mixture-of-Experts for Control of Hybrid Dynamical Systems with Uncertainty (Leroy D'Souza, 2025) View paper
- Locomotion with Mixture-of-Experts (1 papers)
  - [9] MoE-Loco: Mixture of Experts for Multitask Locomotion (Huang Runhan, 2025) View paper
- Diffusion Policies for Hybrid Dynamical Systems (1 papers)
  - [16] Learning Adaptive Diffusion Policies for Hybrid Dynamical Systems (L D'Souza, 2025) View paper
- Skill Composition and Specialization with MoE (1 papers)
  - [13] Specializing Versatile Skill Libraries using Local Mixture of Experts (Onur Ã�elik, 2021) View paper
- Vision-Language-Action Models with MoE (1 papers)
  - [6] Expertise need not monopolize: Action-specialized mixture of experts for vision-language-action learning (Shen Wei-jie, 2025) View paper
- Factorized and Modular Multi-Task Learning (1 papers)
  - [19] Flexible Multitask Learning with Factorized Diffusion Policy (Liu, n.d.) View paper
- Factor-Aware MoE for Generalization (1 papers)
  - [21] FAME: FACTOR-AWARE MIXTURE-OF-EXPERTS WITH (ENCODER, n.d.) View paper
- Survey Literature on Diffusion and VLA Models (3 papers)
  - [1] A survey on diffusion policy for robotic manipulation: Taxonomy, analysis, and future directions (Mingchen Song, 2025) View paper
  - [11] Recipe for Vision-Language-Action Models in Robotic Manipulation: A Survey (Tomohiro Motoda, 2025) View paper
  - [18] Diffusion Model in Robotics: A Comprehensive Review (Linxin Bai, n.d.) View paper

Narrative

Core task: multi-task robot manipulation with diffusion-based mixture-of-experts policies. The field combines two powerful paradigms—diffusion models for generating smooth, multimodal action distributions and mixture-of-experts (MoE) architectures for decomposing complex multi-task problems into specialized sub-policies. The taxonomy reflects a rich landscape organized around several complementary themes. One major branch focuses on diffusion policy architectures themselves, exploring how to integrate MoE gating and skill decomposition directly into the generative process (e.g., Sparse Diffusion Policy[3], MoE-DP Skill Decomposition[14]). Adjacent branches examine flow-matching alternatives (Variational Flow-Matching Policy[7]) and representation learning strategies (Spatially-Grounded Representations[8]) that provide the perceptual backbone for these policies. Other directions address dexterous manipulation (Dexterous Pre-Grasp Diffusion[5], UniDexFPM[17]), residual learning frameworks (Residual MoE Grasping[12]), and broader multi-task reinforcement learning with MoE (Attention MoE MTRL[10]). Additional branches cover locomotion (MoE Locomotion[9]), hybrid dynamical systems (Adaptive Diffusion Hybrid[16]), skill composition (Local MoE Skills[13]), and vision-language-action models (VLA Models Survey[11]), alongside survey literature (Diffusion Policy Survey[1], Diffusion Robotics Review[18]) that contextualizes these developments.

Within this landscape, a particularly active line of work centers on skill-based MoE diffusion policies, where the goal is to learn a set of expert diffusion models that specialize in distinct manipulation primitives and combine them via learned gating mechanisms. MoE Diffusion Skills[0] sits squarely in this cluster, emphasizing the decomposition of multi-task manipulation into interpretable skill modules within a unified diffusion framework. This approach contrasts with methods like Sparse Diffusion Policy[3], which uses sparsity to prune unnecessary network capacity rather than explicitly modeling skill boundaries, and with MoE-DP Skill Decomposition[14], which also pursues skill-level factorization but may differ in how experts are trained or gated. A key open question across these works is how to balance the expressiveness of individual expert policies against the complexity of the gating network, and whether skill discovery should be supervised, emergent, or guided by auxiliary objectives. By integrating MoE structure directly into the diffusion denoising process, MoE Diffusion Skills[0] aims to achieve both high performance on diverse tasks and interpretable specialization, positioning it as a representative example of this skill-based MoE diffusion paradigm.

Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

1. MoE-DP: An MoE-Enhanced Diffusion Policy for Robust Long-Horizon Robotic Manipulation with Skill Decomposition and Failure Recovery

Authors: Baiye Cheng, Huang Suning, Tianhai Liang, Suning Huang, Zhang Fei-hong, et al. (15 authors total) | **Year/Venue**: 2025 | **URL**: View paper

Abstract

Diffusion policies have emerged as a powerful framework for robotic visuomotor control, yet they often lack the robustness to recover from subtask failures in long-horizon, multi-stage tasks and their learned representations of observations are often difficult to interpret. In this work, we propose the Mixture of Experts-Enhanced Diffusion Policy (MoE-DP), where the core idea is to insert a Mixture of Experts (MoE) layer between the visual encoder and the diffusion model. This layer decomposes t...

Relationship Analysis

Both papers belong to the Skill-Based MoE Diffusion Policies category, using mixture-of-experts architectures to decompose robot manipulation tasks into reusable skills. While the original paper (SMP) focuses on learning a state-adaptive orthogonal skill basis with sticky routing for multi-task generalization and transfer learning, the candidate paper (MoE-DP) emphasizes robustness to disturbances and failure recovery in long-horizon tasks through explicit skill decomposition and inference-time control. The key difference is that SMP prioritizes geometric skill disentanglement via orthonormal bases for cross-task reuse, whereas MoE-DP targets stage-aware recovery and flexible subtask reordering within individual long-horizon tasks.

Contributions Analysis

Overall novelty summary. The paper introduces Skill Mixture-of-Experts Policy (SMP), a diffusion-based MoE architecture that learns orthogonal skill bases and uses sticky routing to compose actions from task-relevant expert subsets. It resides in the 'Skill-Based MoE Diffusion Policies' leaf, which contains only two papers total (including this one). This leaf sits within the broader 'Mixture-of-Experts Integration in Diffusion Policies' branch, indicating a relatively sparse but active research direction focused on embedding MoE structures directly into diffusion policy frameworks for multi-task manipulation.

The taxonomy reveals several neighboring approaches: 'Denoiser-Level MoE' applies MoE to denoising transformers rather than skill decomposition, 'Language-Conditioned MoE' uses language instructions for routing, and 'Sparse Diffusion Policies' achieves efficiency through pruning rather than explicit expert specialization. Adjacent branches explore 'Distillation Methods' and 'Flow-Matching Alternatives', while more distant nodes address dexterous manipulation, locomotion, and vision-language-action models. The paper's focus on orthogonal skill bases and sticky routing distinguishes it from these related but structurally different approaches to multi-task learning.

Among fifteen candidates examined across three contributions, no clearly refutable prior work was identified. The core SMP architecture examined three candidates with zero refutations, the adaptive expert activation strategy examined ten candidates with zero refutations, and the variational training objective with sticky routing examined two candidates with zero refutations. This suggests that within the limited search scope—primarily top-K semantic matches and citation expansion—the specific combination of orthogonal skill learning, sticky routing, and adaptive activation appears relatively unexplored, though the broader MoE-diffusion paradigm is established.

Based on the limited literature search, the work appears to occupy a distinctive position within skill-based MoE diffusion policies, particularly in its integration of sticky routing and adaptive activation. However, the analysis covers only fifteen candidates from semantic search, not an exhaustive survey of all multi-task manipulation or MoE literature. The sparse population of the immediate taxonomy leaf (two papers) and absence of refutable candidates suggest novelty in the specific technical approach, though broader claims would require more comprehensive coverage.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

Contribution 1: Skill Mixture-of-Experts Policy (SMP)

Description: SMP is a diffusion-based mixture-of-experts framework that explicitly abstracts reusable manipulation skills via a state-dependent orthonormal action basis with sticky routing. This design improves performance across multiple tasks by learning disentangled, phase-consistent behaviors that can be reused and transferred.

This contribution was assessed against **3 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. KungfuBot2: Learning Versatile Motion Skills for Humanoid Whole-Body Control

URL: View paper

Brief Assessment

KungfuBot Humanoid[22] focuses on humanoid whole-body motion tracking with an orthogonal mixture-of-experts architecture for motion skill learning, not on manipulation tasks with diffusion-based policies and state-dependent orthogonal action basis as in the original paper.

---

2. MoE-DP: An MoE-Enhanced Diffusion Policy for Robust Long-Horizon Robotic Manipulation with Skill Decomposition and Failure Recovery

URL: View paper

Brief Assessment

MoE-DP Skill Decomposition[14] focuses on robustness and failure recovery in long-horizon tasks through MoE-enhanced diffusion policies, while the original paper emphasizes learning reusable skills via state-dependent orthonormal action basis with sticky routing for multi-task generalization and transfer learning.

---

3. Flexible Multitask Learning with Factorized Diffusion Policy

URL: View paper

Brief Assessment

Factorized Diffusion Multitask[19] focuses on compositional diffusion modeling through score aggregation for multitask learning, not on orthonormal skill basis with sticky routing. The candidate's approach uses continuous score-weighted composition of diffusion components, while the original paper explicitly constructs state-adaptive orthogonal action primitives with Dirichlet-Markov sticky gates for phase-consistent skill abstraction.

Contribution 2: Adaptive expert activation strategy

Description: An inference-time mechanism that selects only a small, state-dependent subset of experts (via top-k or coverage selection) to activate at each step. This reduces active parameters and latency substantially while preserving policy quality.

This contribution was assessed against **10 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

---

1. LExI: Layer-Adaptive Active Experts for Efficient MoE Model Inference

URL: View paper

Brief Assessment

LExI[28] focuses on layer-wise expert allocation in language/vision transformers for inference efficiency, while the original paper addresses state-dependent expert selection in diffusion-based robot manipulation policies. The technical contexts and problem domains are fundamentally different.

---

2. ROUTERRETRIEVER: Routing over a Mixture of Expert Embedding Models

URL: View paper

Brief Assessment

RouterRetriever[34] focuses on routing between domain-specific expert embedding models for information retrieval tasks, not on adaptive expert activation during inference for robot manipulation policies. The routing mechanism selects one expert per query based on embedding similarity, rather than dynamically activating a subset of experts at each timestep to reduce computational cost while preserving policy quality.

---

3. Exploiting inter-layer expert affinity for accelerating mixture-of-experts model inference

URL: View paper

Brief Assessment

Inter-Layer Expert Affinity[27] focuses on exploiting inter-layer expert affinity to optimize expert placement across GPUs for reducing communication overhead, not on adaptive inference-time expert activation mechanisms that dynamically select subsets of experts per step.

---

4. MC-MoE: Mixture Compressor for Mixture-of-Experts LLMs Gains More

URL: View paper

Brief Assessment

MC-MoE[32] focuses on compression techniques (quantization and pruning) for MoE-LLMs in NLP, not on learning adaptive expert activation strategies for robot manipulation policies. The original paper's contribution is specific to inference-time routing mechanisms in diffusion-based manipulation policies with state-dependent orthogonal skill bases.

5. ExpertFlow: Adaptive Expert Scheduling and Memory Coordination for Efficient MoE Inference

URL: View paper

Brief Assessment

ExpertFlow Adaptive[31] focuses on adaptive expert prefetching and cache-aware routing for MoE inference optimization across hardware platforms, not on inference-time expert selection mechanisms that reduce active parameters while preserving policy quality in robot manipulation tasks.

6. ExpertFlow: Optimized Expert Activation and Token Allocation for Efficient Mixture-of-Experts Inference

URL: View paper

Brief Assessment

ExpertFlow[26] focuses on inference-time expert offloading and scheduling between GPU and CPU for sparse MoE models in language tasks, using predictive routing and dynamic token scheduling. The original paper's adaptive expert activation operates within a robotics manipulation context with state-dependent orthogonal skill bases and sticky routing for action generation, which is a fundamentally different application domain and technical approach.

7. A survey on inference optimization techniques for mixture of experts models

URL: View paper

Brief Assessment

MoE Inference Survey[25] is a survey paper that reviews existing optimization techniques for MoE models, including adaptive expert activation strategies. It does not present original research that would refute the novelty of the original paper's specific inference-time mechanism with top-k or coverage selection for diffusion-based manipulation policies.

8. Pre-gated moe: An algorithm-system co-design for fast and scalable mixture-of-expert inference

URL: View paper

Brief Assessment

Pre-gated MoE[33] focuses on inference-time expert selection for reducing latency in MoE models, not on learning reusable manipulation skills or action-space decomposition for robot control tasks.

9. Moe-lpr: Multilingual extension of large language models through mixture-of-experts with language priors routing

URL: View paper

Brief Assessment

Multilingual MoE LPR[30] focuses on language-specific routing in multilingual NLP, not adaptive expert activation for robot manipulation. The candidate's routing mechanism selects experts based on language priors, while the original paper's contribution addresses state-dependent expert selection for reducing inference latency in robotic control tasks.

10. AdaMV-MoE: Adaptive Multi-Task Vision Mixture-of-Experts

URL: View paper

Brief Assessment

AdaMV-MoE[29] focuses on determining the number of experts per task in multi-task vision recognition (classification, detection, segmentation), not on state-dependent expert selection during policy execution for robot manipulation.

Contribution 3: Variational training objective with sticky routing

Description: A principled variational lower-bound formulation that combines reconstruction in a whitened basis, gate regularization via sticky Dirichlet Markov dynamics, and router alignment. This objective enables stable training of the orthonormal skill basis and phase-consistent gating.

This contribution was assessed against **2 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

1. A Hybrid Model With Bayesian Nonparametric Inference for RF Fingerprint Identification

URL: View paper

Brief Assessment

Bayesian RF Fingerprint[23] applies sticky control to Bayesian nonparametric models for RF fingerprint identification, not mixture-of-experts diffusion policies for robot manipulation. The technical domains and applications are fundamentally different.

2. The Seeded Codebook Cortex

URL: View paper

Brief Assessment

Seeded Codebook Cortex[24] focuses on vector quantization with sticky cluster selection for compression and routing, whereas the original paper develops a variational lower-bound combining reconstruction in a whitened basis, gate regularization via sticky Dirichlet Markov dynamics, and router alignment for diffusion-based manipulation policies. The technical approaches and application domains differ substantially.

Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

References

- [0] Abstracting Robot Manipulation Skills via Mixture-of-Experts Diffusion Policies View paper
- [1] A survey on diffusion policy for robotic manipulation: Taxonomy, analysis, and future directions View paper
- [2] Variational distillation of diffusion policies into mixture of experts View paper
- [3] Sparse diffusion policy: A sparse, reusable, and flexible policy for robot learning View paper

- [4] Language-Conditioned Representations and Mixture-of-Experts Policy for Robust Multi-Task Robotic Manipulation View paper
- [5] Dexterous Functional Pre-Grasp Manipulation with Diffusion Policy View paper
- [6] Expertise need not monopolize: Action-specialized mixture of experts for vision-language-action learning View paper
- [7] VFP: Variational Flow-Matching Policy for Multi-Modal Robot Manipulation View paper
- [8] Bridging Perception and Action: Spatially-Grounded Mid-Level Representations for Robot Generalization View paper
- [9] MoE-Loco: Mixture of Experts for Multitask Locomotion View paper
- [10] Multi-Task Reinforcement Learning With Attention-Based Mixture of Experts View paper
- [11] Recipe for Vision-Language-Action Models in Robotic Manipulation: A Survey View paper
- [12] Efficient Residual Learning with Mixture-of-Experts for Universal Dexterous Grasping View paper
- [13] Specializing Versatile Skill Libraries using Local Mixture of Experts View paper
- [14] MoE-DP: An MoE-Enhanced Diffusion Policy for Robust Long-Horizon Robotic Manipulation with Skill Decomposition and Failure Recovery View paper
- [15] SAC-MoE: Reinforcement Learning with Mixture-of-Experts for Control of Hybrid Dynamical Systems with Uncertainty View paper
- [16] Learning Adaptive Diffusion Policies for Hybrid Dynamical Systems View paper
- [17] UniDexFPM: Universal Dexterous Functional Pre-grasp Manipulation Via Diffusion Policy View paper
- [18] Diffusion Model in Robotics: A Comprehensive Review View paper
- [19] Flexible Multitask Learning with Factorized Diffusion Policy View paper
- [20] Efficient and Scalable Diffusion Transformer Policies with Mixture of Expert Denoisers for Multitask Learning View paper
- [21] FAME: FACTOR-AWARE MIXTURE-OF-EXPERTS WITH View paper
- [22] KungfuBot2: Learning Versatile Motion Skills for Humanoid Whole-Body Control View paper
- [23] A Hybrid Model With Bayesian Nonparametric Inference for RF Fingerprint Identification View paper
- [24] The Seeded Codebook Cortex View paper
- [25] A survey on inference optimization techniques for mixture of experts models View paper
- [26] ExpertFlow: Optimized Expert Activation and Token Allocation for Efficient Mixture-of-Experts Inference View paper
- [27] Exploiting inter-layer expert affinity for accelerating mixture-of-experts model inference View paper
- [28] LExI: Layer-Adaptive Active Experts for Efficient MoE Model Inference View paper
- [29] AdaMV-MoE: Adaptive Multi-Task Vision Mixture-of-Experts View paper
- [30] Moe-lpr: Multilingual extension of large language models through mixture-of-experts with language priors routing View paper
- [31] ExpertFlow: Adaptive Expert Scheduling and Memory Coordination for Efficient MoE Inference View paper
- [32] MC-MoE: Mixture Compressor for Mixture-of-Experts LLMs Gains More View paper
- [33] Pre-gated moe: An algorithm-system co-design for fast and scalable mixture-of-expert inference View paper
- [34] ROUTERRETRIEVER: Routing over a Mixture of Expert Embedding Models View paper