Novelty Assessment Report

Paper: DataMIL: Selecting Data for Robot Imitation Learning with Datamodels

PDF URL: https://openreview.net/pdf?id=AcTsKglDdh

Venue: ICLR 2026 Conference Submission

Year: 2026

Report Generated: 2025-12-30

Abstract

Recently, the robotics community has amassed ever larger and more diverse datasets to train generalist policies. However, while these policies achieve strong mean performance across a variety of tasks, they often underperform on individual, specialized tasks and require further tuning on newly acquired task-specific data. Combining task-specific data with carefully curated subsets of large prior datasets via co-training can produce better specialized policies, but selecting data naively may actually harm downstream performance. To address this, we introduce DataMIL, a data selection framework built on the datamodels paradigm that reasons about data selection in an end-to-end manner, using the policy itself to identify which data points will most improve performance. Unlike standard practices that filter data using human notions of quality (e.g., based on semantic or visual similarity), DataMIL directly optimizes data selection for task success, allowing us to select data that improves the policy while dropping data that degrade it. To avoid performing expensive rollouts in the environment during selection, we introduce a surrogate loss function on task-specific data, allowing us to use DataMIL in the real world without degrading performance. We validate our approach on 60+ simulation and real-world manipulation tasks, notably showing successful data selection from the largest open collections of robot datasets (OXE); demonstrating consistent gains in success rates over prior works. Our results underscore the importance of end-to-end, performance-aware data selection for unlocking the potential of large prior datasets in robotics.

Disclaimer

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

Core Task Landscape

This paper addresses: **Data Selection for Robot Imitation Learning**

A total of **50 papers** were analyzed and organized into a taxonomy with **16 categories**.

Taxonomy Overview

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

- **Data Quality Assessment and Curation**
- **Data Retrieval and Augmentation from Prior Datasets**
- **Active Data Collection and Human Interaction**
- **Data Scaling and Efficiency Analysis**
- **Representation and Feature Selection for Learning**
- **Domain Adaptation and Transfer Learning**
- **Specialized Learning Paradigms and Applications**
- **Foundational Methods and Frameworks**

Complete Taxonomy Tree

- Data Selection for Robot Imitation Learning Survey Taxonomy
- Data Quality Assessment and Curation
  - Quality Metrics and Estimation (4 papers)
  - [1] Data quality in imitation learning (Belkhale, 2023) View paper
  - [8] Robot data curation with mutual information estimators (Hejna, 2025) View paper
  - [16] CUPID: Curating Data your Robot Loves with Influence Functions (Agia, 2025) View paper
  - [42] Quantifying Demonstration Quality for Robot Learning and Generalization (Maram Sakr, 2022) View paper
  - Handling Imperfect and Imbalanced Demonstrations (4 papers)
  - [7] Trend: Tri-teaching for robust preference-based reinforcement learning with demonstrations (Shuaiyi Huang, 2025) View paper
  - [22] Imitation learning from imperfection: Theoretical justifications and algorithms (Li, 2023) View paper
  - [28] Towards balanced behavior cloning from imbalanced datasets (Parekh, 2025) View paper
  - [44] Identifying Expert Behavior in Offline Training Datasets Improves Behavioral Cloning of Robotic Manipulation Policies (Qiang Wang, 2023) View paper
- Data Retrieval and Augmentation from Prior Datasets
  - Retrieval-Based Few-Shot Learning (4 papers)
  - [2] Data retrieval with importance weights for few-shot imitation learning (Xie, 2025) View paper
  - [10] COLLAGE: Adaptive Fusion-based Retrieval for Augmented Policy Learning (Kumar, 2025) View paper
  - [25] Task-Unaware Lifelong Robot Learning with Retrieval-based Weighted Local Adaptation (Pengzhi Yang, 2024) View paper
  - [41] Learning and Retrieval from Prior Data for Skill-based Imitation Learning (Nasiriany, 2022) View paper
  - Data Mixing and Weighting Strategies ★ (2 papers)
  - [0] DataMIL: Selecting Data for Robot Imitation Learning with Datamodels (Anon et al., 2026) View paper
  - [4] Re-mix: Optimizing data mixtures for large scale imitation learning (Hejna, 2024) View paper
  - Synthetic Data Generation and Augmentation (2 papers)
  - [6] Interventional data generation for robust and data-efficient robot imitation learning (R Hoque, 2023) View paper

- [30] IntervenGen: Interventional Data Generation for Robust and Data-Efficient Robot Imitation Learning (Ryan Hoque, 2024) View paper
- Active Data Collection and Human Interaction
  - Active Querying and Intervention (4 papers)
  - [11] Selective sampling and imitation learning via online regression (Sekhari, 2023) View paper
  - [15] Safari: Safe and active robot imitation learning with imagination (Di Palo, 2020) View paper
  - [26] Adversarial Data Collection: Human-Collaborative Perturbations for Efficient and Robust Robotic Imitation Learning (Huang, 2025) View paper
  - [29] Active Robot Curriculum Learning from Online Human Demonstrations (Mu-Han Hou, 2025) View paper
  - Demonstration Modality and Interface Design (4 papers)
  - [12] Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations (Michael Laskey, 2017) View paper
  - [31] Leveraging Haptic Feedback to Improve Data Quality and Quantity for Deep Imitation Learning Models. (Catie Cuan, 2025) View paper
  - [39] A robot learning from demonstration framework to perform force-based manipulation tasks (Leonel Rozo, 2013) View paper
  - [40] A Process-Oriented Framework for Robot Imitation Learning in Human-Centered Interactive Tasks (Mu-Han Hou, 2023) View paper
  - Preference-Based and Batch Active Learning (2 papers)
  - [18] Batch active preference-based learning of reward functions (Erdem BÄ±yÄ±k, 2018) View paper
  - [47] Active learning of probabilistic movement primitives (Conkey, 2019) View paper
- Data Scaling and Efficiency Analysis (2 papers)
  - [5] Data scaling laws in imitation learning for robotic manipulation (Hu Ying-dong, 2024) View paper
  - [38] JUICER: Data-Efficient Imitation Learning for Robotic Assembly (Lars Ankile, 2024) View paper
- Representation and Feature Selection for Learning
  - Task Space and Feature Selection (2 papers)
  - [14] Automatic selection of task spaces for imitation learning (Manuel MÃ¼hlig, 2009) View paper
  - [24] Human-Driven Feature Selection for a Robotic Agent Learning Classification Tasks from Demonstration (Kalesha Bullard, 2018) View paper
  - Skill Segmentation and Hierarchical Learning (3 papers)
  - [36] A novel trajectory learning method for robotic arms based on Gaussian Mixture Model and k-value selection algorithm. (Jingnan Yan, 2025) View paper
  - [37] Robot learning from demonstration by constructing skill trees (George Konidaris, 2012) View paper
  - [43] LeSkill: Structured Skill Learning for Long-Horizon Robotic Manipulation Tasks (Xiucai Huang, 2025) View paper
- Domain Adaptation and Transfer Learning (3 papers)
  - [19] Metamvuc: Active learning for sample-efficient sim-to-real domain adaptation in robotic grasping (Maximilian Gilles, 2025) View paper
  - [46] Active image sampling on canonical views for novel object detection (Qianli Xu, 2020) View paper
  - [49] Sample Efficient Robot Learning in Supervised Effect Prediction Tasks (Mehmet Eren, 2024) View paper
- Specialized Learning Paradigms and Applications
  - Video and Vision-Language Demonstrations (3 papers)
  - [9] Learning Adaptive Dexterous Grasping from Single Demonstrations (Liu Yu-lin, 2025) View paper
  - [32] MA-ROESL: Motion-aware Rapid Reward Optimization for Efficient Robot Skill Learning from Single Videos (Wang Xiang-Hui, 2025) View paper
  - [33] Meta-imitation learning by watching video demonstrations (J Li, 2021) View paper
  - World Models and Imagination-Based Learning (2 papers)
  - [21] Subconscious Robotic Imitation Learning (Xie Jun, 2024) View paper
  - [34] Dream to Manipulate: Compositional World Models Empowering Robot Imitation Learning with Imagination (Barcellona, 2024) View paper
  - Specialized Task Domains (3 papers)
  - [13] A practical roadmap to learning from demonstration for robotic manipulators in manufacturing (Alireza Barekatain, 2024) View paper
  - [35] Wish you were here: Hindsight Goal Selection for long-horizon dexterous manipulation (Davchev, 2021) View paper
  - [50] Robotic Constrained Imitation Learning for the Peg Transfer Task in Fundamentals of Laparoscopic Surgery (Kento KAWAHARAZUKA, 2024) View paper
- Foundational Methods and Frameworks (7 papers)
  - [3] Robot learning from human teachers (Sonia Chernova, 2022) View paper
  - [17] Toward Probabilistic Safety Bounds for Robot Learning from Demonstration. (Daniel S. Brown, 2017) View paper
  - [20] Robot learning from failed demonstrations (D. Grollman, 2012) View paper
  - [23] Natural methods for robot task learning: Instructive demonstrations, generalization and practice (Monica N. Nicolescu, 2003) View paper
  - [27] Behavior cloning and replay of humanoid robot via a depth camera (Quan-tao Wang, 2023) View paper
  - [45] Deep Q-learning from Demonstrations (Todd Hester, 2017) View paper
  - [48] Learning robot behaviour and skills based on human demonstration and advice: the machine learning paradigm (R. Dillmann, 2000) View paper

Narrative

Core task: data selection for robot imitation learning. The field addresses how to choose, curate, and combine demonstration data so that robots can learn manipulation skills efficiently and robustly. The taxonomy organizes research into several major branches. Data Quality Assessment and Curation focuses on filtering or ranking demonstrations by metrics such as success likelihood or mutual information (e.g., Data Quality Imitation[1], Mutual Information Curation[8]). Data Retrieval and Augmentation from Prior Datasets explores how to mine and mix existing corpora, including strategies for weighting or blending heterogeneous sources (e.g., Re-mix[4], Data Retrieval Weights[2]). Active Data Collection and Human Interaction examines methods that query human teachers or adaptively gather new demonstrations (e.g., Learning Human Teachers[3], Batch Active Preference[18]). Data Scaling and Efficiency Analysis investigates how performance changes with dataset size and composition (e.g., Data Scaling Laws[5]). Representation and Feature Selection for Learning considers which state or action features matter most for generalization. Domain Adaptation and Transfer Learning tackles distribution shifts across tasks or embodiments. Specialized Learning Paradigms and Applications cover niche settings such as dexterous grasping or surgical robotics, while Foundational Methods and Frameworks provide core algorithmic building blocks.

A particularly active line of work centers on mixing and weighting strategies within the Data Retrieval and Augmentation branch, where researchers debate how to balance diverse demonstration sources—some high-quality, some suboptimal—to maximize policy performance. DataMIL[0] sits squarely in this cluster, proposing a principled approach to data mixing that accounts for varying demonstration quality and task relevance. It contrasts with Re-mix[4], which emphasizes replay-buffer blending for continual learning, and with Data Retrieval Weights[2], which focuses on retrieval-based weighting from large prior datasets. Meanwhile, works in Data Quality Assessment (e.g., Data Quality Imitation[1], Mutual Information Curation[8]) offer complementary perspectives by first filtering demonstrations before any mixing occurs. Across these branches, a central open question remains: whether to curate aggressively upfront or to rely on adaptive weighting during training, and how scaling laws (Data Scaling Laws[5]) inform these trade-offs in practice.

Related Works in Same Category

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

1. Re-mix: Optimizing data mixtures for large scale imitation learning

Authors: Hejna, Joey, Bhateja, Chethan, Joey Hejna, et al. (14 authors total) | **Year/Venue**: 2024 | **URL**: View paper

Abstract

Increasingly large imitation learning datasets are being collected with the goal of training foundation models for robotics. However, despite the fact that data selection has been of utmost importance in vision and natural language processing, little work in robotics has questioned what data such models should actually be trained on. In this work we investigate how to weigh different subsets or ``domains'' of robotics datasets for robot foundation model pre-training. Concrete, we use distributio...

Relationship Analysis

Both papers belong to the Data Mixing and Weighting Strategies category, focusing on optimally combining multiple data sources for robot imitation learning. They overlap in addressing how to weight and select data from heterogeneous robot datasets to improve policy performance. However, DataMIL uses datamodels to predict individual datapoint influence on task-specific performance through end-to-end optimization, while Re-Mix employs distributionally robust optimization (DRO) to learn domain-level mixture weights that maximize worst-case performance across all domains for generalist pre-training, without targeting specific downstream tasks.

Contributions Analysis

Overall novelty summary. The paper introduces DataMIL, a framework for selecting and mixing demonstration data to train specialized robot policies. It sits within the 'Data Mixing and Weighting Strategies' leaf of the taxonomy, which contains only one other sibling paper (Data Retrieval Weights). This leaf is notably sparse compared to neighboring branches such as 'Quality Metrics and Estimation' (four papers) or 'Retrieval-Based Few-Shot Learning' (four papers), suggesting that principled data mixing for robot imitation learning remains an underexplored research direction despite the growing availability of large-scale datasets.

The taxonomy reveals that DataMIL occupies a position between two related but distinct research threads. Upstream, the 'Data Quality Assessment and Curation' branch (eight papers across two leaves) focuses on filtering demonstrations before training, using metrics like mutual information or success likelihood. Downstream, the 'Retrieval-Based Few-Shot Learning' leaf (four papers) emphasizes selecting relevant subsets from prior datasets using distance metrics. DataMIL bridges these directions by reasoning about data selection in an end-to-end manner during policy training, rather than relying solely on upfront filtering or retrieval heuristics.

Among thirty candidates examined, none clearly refute any of the three core contributions. The DataMIL framework itself (ten candidates examined, zero refutable) appears novel in its application of the datamodels paradigm to robot imitation learning. The surrogate loss function (ten candidates, zero refutable) addresses a tractability challenge specific to robotic settings, where rollout costs make standard datamodels approaches prohibitive. The adaptations of datamodels for robotic contexts (ten candidates, zero refutable) also show no substantial prior overlap within the limited search scope, though the analysis acknowledges this reflects top-K semantic matches rather than exhaustive coverage.

Based on the limited literature search, the work appears to introduce a relatively fresh perspective on data mixing for robot learning. The sparse taxonomy leaf and absence of refutable candidates suggest novelty, though the search scope (thirty candidates) leaves open the possibility of relevant prior work in adjacent machine learning communities. The framework's distinctiveness lies in its end-to-end optimization approach, which contrasts with the filtering-then-training or retrieval-then-training paradigms prevalent in neighboring taxonomy branches.

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

Contribution 1: DataMIL framework for robot imitation learning data selection

Description: The authors propose DataMIL, a framework that extends the datamodels paradigm to robotics by directly optimizing data selection for task success rather than using human-defined heuristics like semantic or visual similarity. The method uses the policy to evaluate which data points improve performance in an end-to-end manner.

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. Interventional data generation for robust and data-efficient robot imitation learning

URL: View paper

Brief Assessment

Interventional Data Generation[6] focuses on generating corrective interventional data to handle distribution shift from pose estimation errors, not on optimizing data selection from large prior datasets using policy-based end-to-end evaluation.

---

2. Hierarchical Human Demonstration Toward Imitation Learning of Generalist Robot Planner

URL: View paper

Brief Assessment

Hierarchical Human Demonstration[66] focuses on hierarchical demonstration learning for robot planning with human guidance, not on policy-based end-to-end data selection optimization. The candidate does not address the DataMIL framework's core contribution of using datamodels to directly optimize data selection for task success.

---

3. Learning and Retrieval from Prior Data for Skill-based Imitation Learning

URL: View paper

Brief Assessment

Learning Retrieval Prior[41] focuses on skill-based imitation learning with retrieval mechanisms for policy learning, not on direct policy-based end-to-end optimization for data selection as in DataMIL.

---

4. Map-based deep imitation learning for obstacle avoidance

URL: View paper

Brief Assessment

Map-based Obstacle Avoidance[64] focuses on obstacle avoidance using local occupancy maps for mobile robots, not on data selection frameworks for imitation learning. The candidate addresses a different problem domain (navigation) with different methods (value iteration networks for map-based planning).

---

5. FAGR: Feature-Action Generative Replay for Robot Lifelong Imitation Learning

URL: View paper

Brief Assessment

FAGR[65] addresses lifelong imitation learning with generative replay in feature-action space to mitigate catastrophic forgetting across sequential tasks. This is fundamentally different from DataMIL's focus on selecting optimal subsets from large prior datasets using policy-based end-to-end optimization for single-task specialization.

---

6. Goal-conditioned imitation learning using score-based diffusion policies

URL: View paper

Brief Assessment

Goal-conditioned Diffusion[61] focuses on learning goal-conditioned policies using score-based diffusion models for action generation, not on data selection or curation for robot imitation learning. The candidate addresses a fundamentally different problem (policy architecture) than the original's contribution (data selection optimization).

---

7. Towards imitation learning to branch for mip: A hybrid reinforcement learning based sample augmentation approach

URL: View paper

Brief Assessment

Imitation Branch MIP[63] focuses on variable selection for mixed integer programming using reinforcement learning to augment training samples for branching policies. This is fundamentally different from DataMIL's robot imitation learning data selection framework, as it addresses a completely different domain (mathematical optimization vs. robotics) and problem (branching decisions in MIP solvers vs. selecting demonstration data for robot policies).

---

8. Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations

URL: View paper

Brief Assessment

Better-than-Demonstrator[68] focuses on achieving better-than-demonstrator performance through noise injection into behavioral cloning policies to generate automatic rankings, rather than directly optimizing data selection for task success using the policy itself as DataMIL does.

---

9. Coherent soft imitation learning

URL: View paper

Brief Assessment

Coherent Soft Imitation[62] focuses on combining behavioral cloning with inverse reinforcement learning for policy improvement, not on data selection for imitation learning. The candidate addresses policy optimization methods rather than selecting training data from large datasets.

---

10. Behavior imitation for manipulator control and grasping with deep reinforcement learning

URL: View paper

Brief Assessment

Behavior Imitation Manipulator[67] focuses on motion imitation from video using 3D pose estimation and reinforcement learning for manipulator control, not on data selection optimization for imitation learning datasets.

Contribution 2: Surrogate loss function for tractable data selection without rollouts

Description: The authors introduce a proxy metric based on validation loss on held-out target demonstrations that replaces expensive real-world rollouts during datamodel estimation. This makes the approach tractable and fully differentiable while maintaining sufficient correlation with true policy performance.

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. Learn to grasp with less supervision: A data-efficient maximum likelihood grasp sampling loss

URL: View paper

Brief Assessment

Maximum Likelihood Grasp[58] addresses grasp planning with sparse labels in robotics, not general data selection for imitation learning. The candidate's proxy metric is task-specific (grasp quality estimation), whereas the original contribution concerns general policy performance estimation across diverse robotics tasks.

---

2. ASP: Automatic Selection of Proxy dataset for efficient AutoML

URL: View paper

Brief Assessment

ASP Proxy AutoML[55] focuses on proxy dataset selection for AutoML tasks (NAS/HPO) using training loss as a surrogate metric, not on robotics policy learning or imitation learning rollouts. The candidate addresses a fundamentally different domain and problem setting.

---

3. Predicting with proxies: Transfer learning in high dimension

URL: View paper

Brief Assessment

Predicting with Proxies[51] addresses transfer learning in high-dimensional regression settings using proxy data, not robotic policy learning or data selection for imitation learning. The candidate focuses on combining proxy and gold data in supervised learning contexts, whereas the original contribution specifically develops a surrogate loss based on validation loss on held-out target demonstrations to replace expensive real-world rollouts in robotics.

---

4. Loss function considering dead zone for neural networks

URL: View paper

Brief Assessment

Dead Zone Loss[53] addresses a completely different problem domain - handling dead zones in robotic actuators during inverse dynamics computation. It does not address data selection for imitation learning or provide alternatives to environment rollouts.

5. When Can Proxies Improve the Sample Complexity of Preference Learning?

URL: View paper

Brief Assessment

Proxies Sample Complexity[54] focuses on theoretical conditions for when proxy feedback can improve sample complexity in preference learning, not on practical surrogate metrics for robotics data selection without environment rollouts.

6. Reinforcement neural fuzzy surrogate-assisted multiobjective evolutionary fuzzy systems with robot learning control application

URL: View paper

Brief Assessment

Reinforcement Neural Fuzzy[57] focuses on surrogate-assisted evolutionary optimization for robot control, not data selection for imitation learning. The surrogate estimates objective function values for controller evolution, not validation loss for selecting training data.

7. Approximate selection with guarantees using proxies

URL: View paper

Brief Assessment

Approximate Selection Proxies[60] addresses a fundamentally different problem domain: selecting data records from databases using proxy models to approximate expensive oracle predicates (e.g., human labelers or expensive DNNs for classification tasks). The original paper focuses on robot imitation learning with policy rollouts and validation loss on demonstrations, which is a distinct application area and technical approach.

8. Sample selecting method based on feature density for pest identification in smart agriculture

URL: View paper

Brief Assessment

Feature Density Selection[56] focuses on pest identification in agriculture using ternary loss functions for feature aggregation. This is fundamentally different from the original paper's surrogate loss for robotics data selection that replaces environment rollouts with validation loss on held-out demonstrations.

9. Accelerating neural architecture search via proxy data

URL: View paper

Brief Assessment

Proxy Data Acceleration[59] focuses on data selection for neural architecture search (NAS) using proxy datasets to reduce computational cost, not on robot imitation learning or policy rollouts. The surrogate metrics serve different purposes in fundamentally different domains.

10. Choosing a proxy metric from past experiments

URL: View paper

Brief Assessment

Choosing Proxy Metric[52] addresses proxy metric construction for randomized experiments in recommendation systems, not surrogate losses for robotic data selection. The domains and objectives are fundamentally different.

Contribution 3: Adaptations of datamodels for robotic settings

Description: The authors develop several modifications to make datamodels work in robotics, including clustering training examples at different temporal scales to reduce variance, using a proxy metric to avoid rollouts, and incorporating target data during estimation to reduce distribution shift in heterogeneous datasets.

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. Domain Adaptive Imitation Learning with Visual Observation

URL: View paper

Brief Assessment

Domain Adaptive Visual[74] focuses on domain adaptation in imitation learning with visual observations using dual feature extraction and cycle-consistency, not on adapting datamodels for robotic settings with clustering or proxy metrics.

2. Robot See, Robot Do: On the Development of Robust and Adaptive Imitation Learning for Robots

URL: View paper

Brief Assessment

Robot See Do[75] focuses on imitation learning from human feedback (demonstrations, corrections, preferences) and stability analysis, not on datamodels or data selection techniques for robotics.

3. Off-dynamics reinforcement learning via domain adaptation and reward augmented imitation

URL: View paper

Brief Assessment

Off-dynamics Domain Adaptation[72] focuses on domain adaptation in reinforcement learning with dynamics shifts, not on adapting datamodels for robotic imitation learning. The candidate addresses different technical challenges (off-dynamics RL via reward modification and imitation learning) rather than datamodel estimation techniques for robotics.

4. Enhancing visual domain robustness in behaviour cloning via saliency-guided augmentation

URL: View paper

Brief Assessment

Saliency-guided Augmentation[69] focuses on visual domain robustness through image augmentation techniques in behaviour cloning, not on data selection or datamodel adaptation for robotics. The technical approaches are fundamentally different.

5. Dida: Denoised imitation learning based on domain adaptation

URL: View paper

Brief Assessment

Dida[70] focuses on denoising noisy expert demonstrations through domain adaptation techniques, not on adapting datamodels for data selection in robotics. The candidate addresses a fundamentally different problem (learning from noisy demonstrations) using different technical approaches (domain adversarial training, noise discriminators).

6. Data quality in imitation learning

URL: View paper

Brief Assessment

Data Quality Imitation[1] focuses on fundamental properties of demonstration data quality (action divergence and transition diversity) rather than adapting datamodels framework for robotics. The paper does not address datamodel estimation methods, clustering strategies, or proxy metrics for avoiding rollouts as described in the original contribution.

7. Grasping with chopsticks: Combating covariate shift in model-free imitation learning for fine manipulation

URL: View paper

Brief Assessment

Grasping Chopsticks[71] addresses covariate shift in imitation learning for fine manipulation through noise injection and ensemble methods, not datamodel adaptations. The paper focuses on model-free imitation learning without interactive experts, using different techniques (transform to object-centric frames, synthetic corrective labels via noise injection, and ensemble of parametric/non-parametric methods) that are unrelated to the datamodels framework for data selection.

8. Imitation learning for sim-to-real adaptation of robotic cutting policies based on residual Gaussian process disturbance force model

URL: View paper

Brief Assessment

Sim-to-real Robotic Cutting[76] focuses on sim-to-real transfer for cutting tasks using Gaussian process regression to model disturbance forces, not on adapting datamodels for robotic imitation learning with clustering, proxy metrics, or distribution shift reduction.

9. Efficient guided policy search via imitation of robust tube MPC

URL: View paper

Brief Assessment

Guided Policy Search[73] focuses on demonstration-efficient MPC compression using robust tube MPC for trajectory tracking, not on adapting datamodels for robotic imitation learning with clustering, proxy metrics, or distribution shift reduction.

10. Safari: Safe and active robot imitation learning with imagination

URL: View paper

Brief Assessment

Safari[15] focuses on safe imitation learning using uncertainty estimation and world models for failure prediction and active learning, not on adapting datamodels for data selection in robotics. The technical approaches are fundamentally different.

Appendix: Text Similarity Detection

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

References

- [0] DataMIL: Selecting Data for Robot Imitation Learning with Datamodels View paper
- [1] Data quality in imitation learning View paper
- [2] Data retrieval with importance weights for few-shot imitation learning View paper
- [3] Robot learning from human teachers View paper
- [4] Re-mix: Optimizing data mixtures for large scale imitation learning View paper
- [5] Data scaling laws in imitation learning for robotic manipulation View paper
- [6] Interventional data generation for robust and data-efficient robot imitation learning View paper
- [7] Trend: Tri-teaching for robust preference-based reinforcement learning with demonstrations View paper
- [8] Robot data curation with mutual information estimators View paper
- [9] Learning Adaptive Dexterous Grasping from Single Demonstrations View paper
- [10] COLLAGE: Adaptive Fusion-based Retrieval for Augmented Policy Learning View paper
- [11] Selective sampling and imitation learning via online regression View paper
- [12] Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations View paper
- [13] A practical roadmap to learning from demonstration for robotic manipulators in manufacturing View paper
- [14] Automatic selection of task spaces for imitation learning View paper
- [15] Safari: Safe and active robot imitation learning with imagination View paper
- [16] CUPID: Curating Data your Robot Loves with Influence Functions View paper
- [17] Toward Probabilistic Safety Bounds for Robot Learning from Demonstration. View paper
- [18] Batch active preference-based learning of reward functions View paper
- [19] Metamvuc: Active learning for sample-efficient sim-to-real domain adaptation in robotic grasping View paper
- [20] Robot learning from failed demonstrations View paper
- [21] Subconscious Robotic Imitation Learning View paper
- [22] Imitation learning from imperfection: Theoretical justifications and algorithms View paper
- [23] Natural methods for robot task learning: Instructive demonstrations, generalization and practice View paper
- [24] Human-Driven Feature Selection for a Robotic Agent Learning Classification Tasks from Demonstration View paper
- [25] Task-Unaware Lifelong Robot Learning with Retrieval-based Weighted Local Adaptation View paper
- [26] Adversarial Data Collection: Human-Collaborative Perturbations for Efficient and Robust Robotic Imitation Learning View paper

- [27] Behavior cloning and replay of humanoid robot via a depth camera View paper
- [28] Towards balanced behavior cloning from imbalanced datasets View paper
- [29] Active Robot Curriculum Learning from Online Human Demonstrations View paper
- [30] IntervenGen: Interventional Data Generation for Robust and Data-Efficient Robot Imitation Learning View paper
- [31] Leveraging Haptic Feedback to Improve Data Quality and Quantity for Deep Imitation Learning Models. View paper
- [32] MA-ROESL: Motion-aware Rapid Reward Optimization for Efficient Robot Skill Learning from Single Videos View paper
- [33] Meta-imitation learning by watching video demonstrations View paper
- [34] Dream to Manipulate: Compositional World Models Empowering Robot Imitation Learning with Imagination View paper
- [35] Wish you were here: Hindsight Goal Selection for long-horizon dexterous manipulation View paper
- [36] A novel trajectory learning method for robotic arms based on Gaussian Mixture Model and k-value selection algorithm. View paper
- [37] Robot learning from demonstration by constructing skill trees View paper
- [38] JUICER: Data-Efficient Imitation Learning for Robotic Assembly View paper
- [39] A robot learning from demonstration framework to perform force-based manipulation tasks View paper
- [40] A Process-Oriented Framework for Robot Imitation Learning in Human-Centered Interactive Tasks View paper
- [41] Learning and Retrieval from Prior Data for Skill-based Imitation Learning View paper
- [42] Quantifying Demonstration Quality for Robot Learning and Generalization View paper
- [43] LeSkill: Structured Skill Learning for Long-Horizon Robotic Manipulation Tasks View paper
- [44] Identifying Expert Behavior in Offline Training Datasets Improves Behavioral Cloning of Robotic Manipulation Policies View paper
- [45] Deep Q-learning from Demonstrations View paper
- [46] Active image sampling on canonical views for novel object detection View paper
- [47] Active learning of probabilistic movement primitives View paper
- [48] Learning robot behaviour and skills based on human demonstration and advice: the machine learning paradigm View paper
- [49] Sample Efficient Robot Learning in Supervised Effect Prediction Tasks View paper
- [50] Robotic Constrained Imitation Learning for the Peg Transfer Task in Fundamentals of Laparoscopic Surgery View paper
- [51] Predicting with proxies: Transfer learning in high dimension View paper
- [52] Choosing a proxy metric from past experiments View paper
- [53] Loss function considering dead zone for neural networks View paper
- [54] When Can Proxies Improve the Sample Complexity of Preference Learning? View paper
- [55] ASP: Automatic Selection of Proxy dataset for efficient AutoML View paper
- [56] Sample selecting method based on feature density for pest identification in smart agriculture View paper
- [57] Reinforcement neural fuzzy surrogate-assisted multiobjective evolutionary fuzzy systems with robot learning control application View paper
- [58] Learn to grasp with less supervision: A data-efficient maximum likelihood grasp sampling loss View paper
- [59] Accelerating neural architecture search via proxy data View paper
- [60] Approximate selection with guarantees using proxies View paper
- [61] Goal-conditioned imitation learning using score-based diffusion policies View paper
- [62] Coherent soft imitation learning View paper
- [63] Towards imitation learning to branch for mip: A hybrid reinforcement learning based sample augmentation approach View paper
- [64] Map-based deep imitation learning for obstacle avoidance View paper
- [65] FAGR: Feature-Action Generative Replay for Robot Lifelong Imitation Learning View paper
- [66] Hierarchical Human Demonstration Toward Imitation Learning of Generalist Robot Planner View paper
- [67] Behavior imitation for manipulator control and grasping with deep reinforcement learning View paper
- [68] Better-than-Demonstrator Imitation Learning via Automatically-Ranked Demonstrations View paper
- [69] Enhancing visual domain robustness in behaviour cloning via saliency-guided augmentation View paper
- [70] Dida: Denoised imitation learning based on domain adaptation View paper
- [71] Grasping with chopsticks: Combating covariate shift in model-free imitation learning for fine manipulation View paper
- [72] Off-dynamics reinforcement learning via domain adaptation and reward augmented imitation View paper
- [73] Efficient guided policy search via imitation of robust tube MPC View paper
- [74] Domain Adaptive Imitation Learning with Visual Observation View paper
- [75] Robot See, Robot Do: On the Development of Robust and Adaptive Imitation Learning for Robots View paper
- [76] Imitation learning for sim-to-real adaptation of robotic cutting policies based on residual Gaussian process disturbance force model View paper