

Quantifying the Generalization Value of Pretraining Tasks in Robotic Manipulation

Final Report – Undergraduate Thesis

Namra Patel

npate334@uwo.ca

Department of Computer Science
Western University



Thesis supervisor: Prof. Boyu Wang
Department of Computer Science

Course instructor: Prof. Nazim Madhavji
Department of Computer Science

April 4th, 2025

Abstract

As robots increasingly learn from experience, a critical question emerges: which skills provide the most generalizable foundations for learning new tasks? While significant effort has focused on scaling model size and dataset quantity, surprisingly little work has systematically quantified what robots should learn first to best enable learning something new next. This thesis introduces a framework for measuring the representational transfer value of robotic manipulation tasks—the extent to which learning one task produces internal representations that accelerate or improve performance on others.

Using the Action Chunking Transformer architecture, we implemented a three-task manipulation hierarchy (grasping, pick-and-place, bimanual transfer) and conducted systematic transfer experiments across all task pairs. We quantified transfer value through two complementary metrics: adaptation efficiency (learning speed) and maximum performance improvement (success rate gains). Models were trained from scratch on each task, then fine-tuned across all source-target combinations to measure forward and backward transfer effects.

Our results reveal striking asymmetries in transfer relationships. We observe symmetric adaptation efficiency across the task hierarchy—both foundational and complex tasks accelerate learning on others. However, maximum performance improvement is asymmetric: only complex tasks improve final performance on simpler ones, while foundational tasks provide little benefit for complex task performance ceilings. Additionally, despite hierarchical task structure, we find substantial catastrophic forgetting during sequential learning, with severe performance degradation on source tasks after fine-tuning.

This work provides the first quantitative framework to evaluate representational transfer value between manipulation tasks, distinguishing between tasks that accelerate learning and those that raise performance ceilings. Our findings inform the design of task curricula for robot learning and highlight the need for retention-aware architectures that preserve transferable representations across sequential task learning.

Contents

1	Introduction	3
2	Background and Related Work	4
2.1	Robotic Manipulation and Imitation Learning	4
2.2	Transfer Learning	4
2.3	Data Selection and Transfer Utility	4
2.4	Analysis and Research Gap	5
3	Research Objectives	5
4	Methodology	5
4.1	Task Selection	5
4.2	Task Implementation	6
4.3	Data Collection	6
4.4	Model Architecture and Training	6
4.5	Evaluation Methodology	7
5	Results	7
5.1	Forward Transfer Analysis	8
5.1.1	Efficiency at 50% Success Threshold	8
5.1.2	Performance at Higher Thresholds	9
5.1.3	Maximum and Average Success Rates	9
5.1.4	Task-Specific Analysis	10
5.2	Backward Transfer Analysis	10
5.3	Implications	11
5.4	Novelty and Significance of Results	11
6	Discussion	12
6.1	Threats to the Validity of the Results	12
6.2	Implications of the Research Results	12
6.3	Generalizability of the Results	12
7	Conclusions	13
8	Future Work and Lessons Learnt	13
8.1	Future Work	13
8.2	Lessons Learnt	14
A	Resources	14
B	Additional Experimental Details	15
C	Task-Specific Subtask Performance	15
D	Learning Curves	16
	References	17

1 Introduction

Robots that learn from experience have the potential to acquire a diverse repertoire of manipulation skills. In recent years, imitation learning and robotics foundation models have enabled robots to perform many tasks from raw sensor inputs, moving towards generalist manipulation agents. For example, new imitation learning algorithms have shown strong task-specific performance (Zhao, Kumar, Levine, & Finn, 2023; Chi et al., 2024) and transformer-based policies trained on massive demonstration datasets can achieve impressive generalization across tasks and objects (Brohan et al., 2023; Ghosh et al., 2024; Kim et al., 2025).

However, not all training tasks are equally useful for broad generalization. As foundation models continue to expand in robotics, a critical question emerges: *Which skills provide the most generalizable foundations for learning other tasks?* While significant effort has focused on scaling model size and dataset quantity, surprisingly little work has systematically quantified what a robot should learn first to best enable learning something new next (Hejna, Bhateja, Jiang, Pertsch, & Sadigh, 2025).

In this work, we introduce a framework that measures the representational transfer value (RTV) of manipulation tasks, defined as the extent to which learning one task produces internal representations that accelerate or improve learning of another task. Our approach differs from traditional generalization studies that focus only on novel environments or objects; instead, we explicitly quantify cross-task transfer potential through two complementary metrics: (1) adaptation efficiency (how quickly pretrained representations enable learning a new task) and (2) final performance improvement (whether pretraining enables higher ultimate success).

By applying this framework to a hierarchy of manipulation skills (grasping, pick-and-place, and bimanual transfer), we reveal several key insights. First, we find striking asymmetry in transfer relationships: pre-training on foundational skills dramatically accelerates learning more complex tasks, while complex skills can elevate performance ceilings on simpler tasks. Second, we observe that hierarchical structure does not prevent catastrophic forgetting, with substantial performance degradation on source tasks after fine-tuning. These findings provide concrete evidence about which manipulation skills serve as effective foundations for others, and how task complexity relates to transferability.

Our experiments utilize the Action Chunking Transformer (ACT) architecture (Zhao et al., 2023), a sequence model that generates chunks of low-level actions from visual input. By maintaining consistent architecture, training procedures, and evaluation metrics across all tasks, we create a controlled experimental framework that isolates the impact of task-specific representations on transfer performance. We examine both forward transfer (how well representations adapt to new tasks) and backward transfer (how fine-tuning affects performance on original tasks), providing a comprehensive analysis of cross-task representational dynamics.

The contributions of this thesis are:

1. A quantitative framework and metrics for measuring the representational transfer value of robotic manipulation tasks,
2. An empirical mapping of transfer relationships across a hierarchy of manipulation skills, revealing that both foundational and complex tasks accelerate learning across tasks (via adaptation efficiency), but only complex tasks elevate the final performance of simpler tasks (via maximum performance improvement),
3. Evidence that hierarchical structure alone is insufficient to prevent catastrophic forgetting during sequential learning, highlighting the need for specialized approaches to knowledge retention in robotic skill acquisition.

Our findings deepen the understanding of how robots can leverage prior learned skills to learn new tasks more efficiently, moving a step closer to lifelong learning robots that can build progressively more complex capabilities from foundational skills.

The remainder of this report is organized as follows. Chapter 2 reviews background and related work in imitation learning, task transfer learning, and dataset curation for robotic skills. Chapter 3 details our methodology, including the definitions of transfer metrics, the ACT-based learning setup, and the experimental design. Chapter 4 presents the results of our transfer experiments and a discussion of the observed

transfer patterns. Chapter 5 concludes with insights, limitations, and future directions for generalist robot learning.

2 Background and Related Work

2.1 Robotic Manipulation and Imitation Learning

Robotic manipulation involves learning control policies for high-dimensional, contact-rich tasks such as object grasping, relocation, and tool use. These tasks are often sequential and require the coordination of visual feedback, proprioceptive sensing, and motor commands in real time. IL is a common approach to acquiring such behaviors, where a model learns to reproduce expert demonstrations by minimizing a supervised loss between predicted and demonstrated actions.

ACT is a transformer-based model for learning fine-grained bimanual manipulation from demonstrations. It leverages a ResNet-based visual encoder (He, Zhang, Ren, & Sun, 2016), a transformer encoder-decoder architecture (Vaswani et al., 2017), and a conditional variational autoencoder (CVAE) (Kingma & Welling, 2022) used only during training to encode action style. ACT predicts fixed-length action chunks, sequences of target joint positions, rather than single-step actions, reducing compounding errors during long-horizon execution. At inference, the style variable z is set to the mean of the prior and action chunks are generated autoregressively based on current joint positions and multi-view visual observations.

While large-scale generalist policies such as Octo (Ghosh et al., 2024) and OpenVLA (Kim et al., 2025) offer impressive multi-robot capabilities and pretrained visual-language-action representations, they are not directly suitable for our study. Both models are pretrained on heterogeneous datasets and contain pretrained submodules, such as vision-language encoders or diffusion heads, that confound our ability to isolate the transfer value of specific manipulation tasks. Moreover, their considerable resource demands and architectural complexity make them impractical for targeted representational analysis. In contrast, ACT offers a lightweight and self-contained architecture that enables controlled experimentation under consistent training and evaluation settings.

2.2 Transfer Learning

Transfer Learning (Pan & Yang, 2010) enables a model to reuse knowledge from a source task or domain to accelerate or improve learning on a related target task. This is particularly valuable in robotics, where collecting high-quality demonstrations can be costly.

In the context of robotic manipulation, transfer learning offers a way to reduce sample complexity by leveraging structural similarity across tasks. For instance, the ability to grasp an object may be foundational for learning more complex tasks such as pick-and-place or bimanual coordination.

Our work is conceptually inspired by the Taskonomy framework (Zamir et al., 2018), which introduces a computational taxonomy of task transfer relationships in computer vision. Taskonomy constructs a *task affinity matrix* by empirically measuring how well representations trained for one task transfer to others. While their analysis is restricted to static image-based tasks, the idea of mapping directional transferability serves as a precedent for our exploration in robotic domains.

2.3 Data Selection and Transfer Utility

As the scale and heterogeneity of robot demonstration datasets grow, identifying which data is most useful becomes increasingly important. Recent work such as Re-Mix (Hejna et al., 2025) proposes a framework for optimizing dataset mixtures to maximize downstream performance across diverse robotic domains. By formulating data curation as a distributionally robust optimization (DRO) problem, Re-Mix learns domain weights that balance performance across potential tasks and embodiments.

While Re-Mix focuses on selecting *where to learn from* in large-scale datasets, our work instead evaluates *what to learn*—which *tasks*, when used as pretraining objectives, offer the most transferable representational structure. Both efforts share the goal of quantifying the structure of generalization, but differ in granularity: Re-Mix operates over datasets, while our approach pairs of tasks, and can be extended to pairs of groups of tasks.

2.4 Analysis and Research Gap

Despite growing interest in multi-task and transfer learning for robot manipulation, there is currently no method in the literature that quantitatively evaluates the generalization value a task provides to others.

Our work addresses this gap by introducing a systematic framework for evaluating RTV across manipulation tasks. We define metrics for both adaptation efficiency (how quickly a task can be learned with transfer) and performance improvement (whether a higher final performance can be achieved). By structuring tasks into a skill hierarchy and analyzing both forward and backward transfer, we provide a novel empirical method to quantify how learned task representations support or interfere with others during fine-tuning.

This fills a crucial gap in robotic learning research: the absence of tools to understand which tasks are most useful to pretrain on, and how those choices influence performance and retention in downstream tasks.

3 Research Objectives

Our research is guided by the following objectives:

- O1:** Develop a quantitative framework for measuring the RTV of robotic manipulation tasks, that is, the extent to which training on one task facilitates learning others.
- O2:** Map the directional flow of transferable representations across a structured hierarchy of manipulation skills.
- O3:** Analyze how a task’s position in the skill hierarchy correlates with different forms of transfer value, such as adaptation efficiency and final performance.
- O4:** Evaluate the retention of prior task capabilities during sequential learning, to understand trade-offs between adaptability and knowledge preservation.

4 Methodology

4.1 Task Selection

We define three manipulation tasks forming a progressive skill hierarchy. These tasks were selected to reflect increasing manipulation complexity. Grasp Cube (Figure 1) introduces object interaction, Pick and Place (Figure 2) requires object relocation to a target, and Transfer Cube (Figure 3) involves bimanual coordination. This progression allows us to analyze how learned representations from foundational tasks contribute to performance on more advanced capabilities, and vice versa. Each task is visualized below with annotated keyframes and a brief description.




Figure 1: **Grasp Cube.** Reach for and lift the red cube. The right arm approaches the cube (#1), grasps it (#2), and lifts it off the table (#3).




Figure 2: **Pick and Place.** Move the red cube onto the blue goal pad. The right gripper approaches the cube (#1), lifts it (#2), and places it on the pad (#3).




Figure 3: **Transfer Cube.** Pass the red cube between arms. The right arm approaches (#1), grasps the cube (#2), and hands it to the left arm (#3).

4.2 Task Implementation

We implemented all tasks in the MuJoCo physics simulation environment. Each task implementation consists of a task class definition inheriting from a base `BimanualViperXTask` class, with an episode initialization function that sets robot and object configurations, and a reward function that analyzes physics contacts to determine task progress and success states.

To ensure consistent evaluation, all tasks feature controlled randomization of initial object positions across a fixed range, encouraging learned representations that can support generalization across spatial variations. The environment design maintains consistent robot kinematics, camera viewpoints, and control interfaces across all tasks, ensuring that observed transfer effects stem from representational reuse rather than confounding implementation factors.

4.3 Data Collection

For each task, we collected 50 successful demonstrations using task-specific scripted policies. Our collection process involved executing end-effector space policies designed with appropriate waypoints for each task, converting these trajectories to joint-space commands, and replaying joint trajectories while recording multi-view observations and corresponding actions.

Each demonstration contains RGB images from multiple camera perspectives (480640 resolution), 14-dimensional joint position and velocity vectors, and corresponding target joint positions. Small noise injection during data collection ensured sufficient diversity while maintaining task success. The final dataset consists of approximately 20,000 state-action pairs per task (50 episodes 400 timesteps), providing a consistent foundation for studying transferability between learned task representations.

4.4 Model Architecture and Training

We employed the Action Chunking Transformer (ACT) architecture for all experiments, which predicts sequences of actions rather than single-step commands. This architecture is particularly well-suited for manipulation tasks due to its ability to mitigate compounding errors through action sequence prediction and to learn temporally extended representations useful for multi-step behaviors.

Our transfer learning methodology follows this algorithmic structure:

1. **Baseline Training:** For each task in our hierarchy, train a model from scratch until convergence (10,000 epochs) to establish performance benchmarks.

2. **Transfer Matrix Construction:** Create all possible source-target task pairs (33 matrix), where each cell represents transferring from a source to a target task.
3. **Forward Transfer:** For each source-target pair, initialize a model with weights from the source task’s trained model, then fine-tune on the target task’s dataset.
4. **Snapshot Preservation:** Save model checkpoints at regular intervals (every 500 epochs) during fine-tuning to capture the evolution of transfer benefits over time.
5. **Backward Transfer:** After fine-tuning on target tasks, evaluate each model on its original source task to assess retention of earlier-learned skills.

To ensure fair comparison, all models were trained using identical hyperparameters, detailed in Table 1 in the Appendix. Models were trained for up to 10,000 epochs on an NVIDIA A40 GPU, with each complete training run taking approximately 3 hours. This methodical approach allows us to isolate the effects of task hierarchy position on the transfer utility of each task.

4.5 Evaluation Methodology

Our evaluation methodology quantifies two key aspects of RTV:

- **Adaptation Efficiency:** Measured by the number of epochs required to reach specific success rate thresholds (50%, 60%, 75%) when fine-tuning on a target task, compared to learning that task from scratch.
- **Maximum Performance Improvement:** Calculated as the increase in peak success rate achieved on the target task due to transfer learning, compared to training from scratch.

The evaluation procedure for each model checkpoint follows a consistent protocol:

1. Load model weights from the specified checkpoint
2. Run 50 rollouts with randomized initial object positions
3. For each rollout:
 - (a) Reset the environment with a new random object configuration
 - (b) Run the policy for the full episode length (400 timesteps)
 - (c) Record success/failure and the highest reward achieved
 - (d) Track subtask-specific completion rates
4. Calculate the overall success rate and average episode return
5. Generate success rate statistics for each subtask component

Each model evaluation consists of 50 rollouts with randomized initial conditions to ensure reliable performance statistics. This comprehensive framework allows us to systematically quantify the *utility* of learned task representations by observing their impact on learning speed, final performance, and retention during task transfer in a manipulation hierarchy.

5 Results

In this section, we present the results of our transfer learning experiments across three hierarchically related manipulation tasks: Grasp Cube, Pick & Place, and Transfer Cube. These tasks form a natural skill hierarchy, where grasping is a prerequisite for pick and place, which in turn is a component of the bimanual transfer task. We use this hierarchical structure as a framework to systematically quantify the RTV of each task, specifically measuring how effectively representations learned from one task accelerate learning (adaptation efficiency) and improve final performance (maximum performance improvement) on other tasks. We analyze both forward transfer (how pretraining on a source task supports learning of a target task) and backward transfer (how fine-tuning on a new task affects performance on the original source task).

5.1 Forward Transfer Analysis

We begin by analyzing forward transfer: models pretrained on one task are fine-tuned on a different, related task. This quantifies the transfer utility of the source task, revealing how its learned features support learning in the target domain.




Figure 4: Training efficiency metrics for forward transfer, showing the number of epochs required to reach different success rate thresholds when fine-tuning from different source tasks. Empty cells indicate the threshold was never reached.

5.1.1 Efficiency at 50% Success Threshold

The results in Figure 4a highlight patterns of adaptation efficiency across the task hierarchy. Models pretrained on either Pick & Place or Transfer Cube achieve 50% success on the Grasp task within 500 epochs, faster than training from scratch (1000 epochs). This demonstrates that tasks higher in the hierarchy confer representational features that transfer efficiently to simpler, lower-level tasks.

Notably, pretraining on Transfer Cube improves adaptation efficiency for Pick & Place, allowing the model to reach 50% success in only 500 epochs, compared to 1500 epochs when trained from scratch. This once again reflects strong transfer from higher to lower positions in the hierarchy, where complex task

representations encapsulate component behaviors.

Perhaps more surprising is the high adaptation efficiency provided by the foundational Grasp task. Pretraining on Grasp enables Pick & Place to reach 50% success in just 500 epochs, a dramatic improvement over training from scratch. This suggests that even basic manipulation skills encode inductive structure that benefits more complex tasks higher in the hierarchy.

For Transfer Cube, pretraining on Grasp results in achieving 50% success within 2000 epochs, substantially faster than the 4000 epochs required from scratch. However, models pretrained on Pick & Place never reach the 50% success threshold for Transfer Cube. This may indicate representational over-specialization: intermediate tasks may encode narrow behaviors that limit generalization to more demanding downstream tasks.

5.1.2 Performance at Higher Thresholds

At the 60% success threshold (Figure 4b), the same general adaptation patterns persist, but limitations become more apparent. Neither Grasp nor Pick & Place pretraining enables models to surpass 60% success on Transfer Cube, highlighting that improvements in adaptation efficiency do not necessarily translate into better final performance.

At 75% (Figure 4c), the transfer landscape narrows further. Only models pretrained on Grasp or Transfer Cube achieve 75% success on Pick & Place. Notably, while Pick & Place fails to reach 75% when trained from scratch, Transfer Cube pretraining eventually enables it, demonstrating meaningful gains in maximum performance. Still, this takes nearly 9000 epochs, compared to just 500 when pretraining on Grasp, an order-of-magnitude difference in adaptation efficiency.

This result reinforces that while multiple source tasks can confer performance benefits, those benefits vary widely in both efficiency and eventual payoff. Foundational tasks like Grasp deliver faster adaptation; higher-level tasks like Transfer Cube may support stronger final outcomes, but only with extensive fine-tuning.

5.1.3 Maximum and Average Success Rates

We now examine RTV through the lens of maximum performance improvement. Figure 5a shows that pretraining on Pick & Place or Transfer Cube improves Grasp’s peak success rate to 0.98 and 0.96 respectively, compared to 0.84 when trained from scratch. These gains confirm that higher-level tasks can yield reusable representations that significantly elevate performance ceilings on simpler tasks.




Figure 5: Performance metrics for forward transfer, showing maximum and average success rates achieved when fine-tuning from different source tasks.

Likewise, pretraining on Grasp raises Pick & Place’s peak performance from 0.74 to 0.86. This supports the idea that foundational skills encode reusable primitives that contribute to long-term performance on more complex behaviors. In contrast, none of the source tasks produce substantial maximum performance gains for Transfer Cube, suggesting that task-specific skills such as bimanual coordination require dedicated learning and cannot be easily composed from lower-level representations.

Figure 5b further illustrates the stability and robustness of transfer. Pick & Place pretraining yields the highest average success for Grasp over the first 5000 epochs (0.86), indicating consistent and efficient early-stage learning. In combination, these metrics offer a multidimensional perspective on task transfer: not just how well models ultimately perform, but how consistently and efficiently they improve.

5.1.4 Task-Specific Analysis

Detailed breakdowns of subtask performance for each manipulation task at 500 and 10,000 training epochs are included in Appendix C (Tables 2-4). These results isolate performance on individual components of each task, such as grasping or transferring, and may provide additional insight into where representational transfer succeeds or fails.

5.2 Backward Transfer Analysis




Figure 6: Performance metrics for backward transfer, showing maximum and average success rates on the original source task after fine-tuning on a target task.

Backward transfer assesses representational retention: how fine-tuning on a new task impacts performance on the original task. This complements our forward analysis by quantifying interference effects and identifying stability constraints.

As shown in Figure 6, fine-tuning typically leads to substantial degradation on the source task. Maximum success rates remain in a partial retention regime (0.24 to 0.64), while average success drops near zero, consistent with catastrophic forgetting. Despite conceptual overlap between tasks, transfer often overwrites prior knowledge rather than building upon it.

This finding is particularly notable given the task hierarchy: one might expect that fine-tuning on higher-level tasks would preserve low-level competencies. Instead, we observe representational interference across all directions, suggesting that task-specific tuning dominates representational stability in current architectures.

5.3 Implications

Based on our experimental results, we identify several key insights with implications for understanding the generalization value of tasks within our manipulation hierarchy:

Adaptation efficiency is bidirectional across the task hierarchy. We observe that pretraining on either foundational or higher-level tasks improves adaptation efficiency when fine-tuning on the other. This indicates that both simple and complex tasks encode transferable structure, and that representational reuse is not strictly dependent on hierarchical position. These results challenge the assumption that only foundational skills are effective pretraining candidates for downstream adaptation.

Maximum performance improvement is asymmetric across the hierarchy. Pretraining on higher-level tasks consistently improves the final performance on simpler, lower-level tasks, but the reverse is not true: foundational tasks provide little to no benefit in improving the maximum performance of complex tasks. This asymmetry suggests that while higher-level skills encode reusable representations applicable to simpler behaviors, task-specific competencies required for complex manipulation cannot be easily composed from lower-level skills alone.

Hierarchical structure does not prevent catastrophic forgetting. Despite the hierarchical relationship between tasks, models exhibit severe performance degradation on source tasks after fine-tuning on new tasks. This suggests that adaptation to new tasks tends to overwrite rather than build upon previously learned skills, even when those skills remain conceptually relevant within the hierarchy.

5.4 Novelty and Significance of Results

In robot learning, there are no studies that provide a framework for determining what value pretraining on particular tasks provides for performance on other tasks. While prior work like Re-Mix (Hejna et al., 2025) proposes techniques for optimizing the composition of large-scale datasets in imitation learning, it does not address the core question of *which tasks are most useful to pretrain on*. Re-Mix operates at the level of dataset domains thus does it offer a method for identifying foundational versus enhancer tasks in a skill hierarchy. Our study is therefore novel both in its experimental design and in the results it uncovers. We highlight the following findings:

- **Asymmetry in maximum performance, but symmetry in adaptation efficiency.** We find that pretraining on tasks at *any level* of the hierarchy improves adaptation efficiency on others, indicating that both foundational and complex skills encode transferable structure useful for rapid learning. However, maximum performance gains are highly asymmetric: only pretraining on *higher-level* tasks (e.g., Transfer Cube) leads to improved final performance on simpler tasks, while foundational tasks offer little value in improving the performance ceiling of complex ones. This reveals a functional distinction between tasks that enable faster learning and those that elevate task performance, both of which are critical, yet distinct, forms of generalization value not previously formalized in robotic manipulation.
- **Hierarchical task structure does not prevent catastrophic forgetting.** Despite the conceptual overlap between tasks in our manipulation hierarchy, we observe severe performance degradation on source tasks after fine-tuning on new targets. This indicates that hierarchical structure alone is insufficient to prevent representational interference, challenging the intuition that progressive task complexity scaffolds help preserve learned skills in neural networks.
- **Systematic evaluation of representational transfer can yield valuable insight.** By studying the structure and magnitude of transfer between tasks, we are able to discover emergent patterns such as the asymmetry of transfer and the instability of hierarchical retention. These findings demonstrate that transferability is not simply a function of task similarity or complexity, but reflects deeper properties of learned representations. Our framework thus serves not only as a diagnostic tool for evaluating pretraining strategies, but also as a foundation for future work in curriculum design, skill discovery, and lifelong learning in robotics.

6 Discussion

6.1 Threats to the Validity of the Results

While our experiments were designed for consistency and interpretability, several limitations constrain the confidence and scope of our findings.

First, each model was trained using a single random seed, and evaluation was based on 50 rollouts per checkpoint. While this produces reasonably stable trends, it does not capture variance across repeated training runs. Future work should introduce statistical averaging over multiple seeds to more rigorously characterize differences in representational transfer across tasks.

Second, all experiments were performed in simulation using MuJoCo. While this allows for precise control and repeatability, it abstracts away many complexities present in physical systems, including sensor noise, hardware variability, and subtle embodiment constraints. It is unclear to what extent the representational transfer patterns we observed would persist in real-world robotic systems.

Third, our task set was deliberately narrow: a three-task hierarchy with limited variation in object properties, scene layout, and embodiment. While this was necessary to control for confounds and isolate representational factors, it limits the generalizability of our observations. Larger and more diverse task suites would allow stronger claims about representational utility at scale.

Finally, our findings are specific to the Action Chunking Transformer architecture, which imposes particular representational biases and temporal priors. It remains an open question whether the transfer patterns we observed are robust across model families, especially architectures with radically different memory mechanisms or inductive biases.

6.2 Implications of the Research Results

Despite these limitations, our results suggest several actionable insights for the design of training pipelines and pretraining curricula in robot learning.

First, the concept of RTV provides a useful lens through which to evaluate training data. Rather than asking whether a model performs well on a given task, we can ask how well that task serves as a source of reusable structure for learning other tasks. This distinction reframes pretraining from being a means to an end, to being a diagnostic probe into the content and reusability of task-specific representations.

Second, our findings reveal that both foundational and high-level tasks offer meaningful adaptation benefits, but in complementary ways. Both types of tasks accelerate learning on others, indicating symmetric adaptation efficiency. However, only high-level tasks enhance the final performance on simpler targets. This suggests that pretraining strategies should be informed by the desired downstream outcome: when the goal is fast adaptation, either category may be suitable; when maximizing asymptotic performance is critical, higher-level tasks are more advantageous. Designing curricula that strategically sequence both may lead to more general and capable manipulation agents.

Third, our analysis of backward transfer highlights an important architectural challenge: even when tasks are hierarchically related, current models fail to retain useful representations across sequential training. Catastrophic forgetting persists despite conceptual overlap between tasks. This points toward a need for memory-stable architectures or training procedures explicitly designed to preserve useful internal representations across sequential task updates.

6.3 Generalizability of the Results

Caution is warranted when generalizing these findings beyond our experimental setup. While the idea of RTV is broadly applicable, its specific empirical expression may vary significantly across task domains, model types, and training paradigms.

For example, locomotion or navigation tasks may rely more on spatial priors or global planning, and less on fine-grained temporal chunking. Representational transfer in such domains may require entirely different metrics or methods. Similarly, multi-task or continual learning settings, where tasks are learned jointly or interleaved, introduce new dynamics that our sequential transfer protocol does not address.

Moreover, while our study focuses on imitation learning from demonstrations, reinforcement learning and self-supervised learning may yield fundamentally different forms of representation that support transfer

through reward shaping, exploration biases, or latent goal inference. Investigating representational reuse under these paradigms remains an important open direction.

Nonetheless, our framework offers a general approach: define tasks as probes, measure directional transfer, and interpret representational reuse not as an artifact of performance, but as a signal of informational structure. This core methodology can be extended, adapted, and stress-tested across domains to further our understanding of how learning experiences encode generalizable knowledge.

7 Conclusions

This thesis addressed the critical question of which manipulation skills provide the most generalizable foundations for learning other tasks. We developed a systematic framework for quantifying the representational transfer value of tasks (O1), mapped directional transferability across a manipulation hierarchy (O2), analyzed how hierarchical position affects different types of transfer (O3), and evaluated retention during sequential learning (O4).

Our results revealed a key asymmetry in maximum performance transfer: models pretrained on complex tasks like Transfer Cube significantly improved final performance on simpler tasks, elevating Grasp’s peak success rate from 0.84 to 0.96. However, the reverse did not hold, foundational tasks like Grasp provided little improvement in the maximum achievable performance of complex tasks. In contrast, adaptation efficiency exhibited directional symmetry: pretraining on either simpler or more complex tasks enabled faster learning on the target task. This suggests a functional distinction between tasks that accelerate learning and those that enhance performance ceilings, highlighting the multifaceted nature of transfer value in robotic manipulation.

Furthermore, we found that hierarchical structure alone is insufficient to prevent catastrophic forgetting. Despite conceptual overlap between tasks, fine-tuning led to substantial performance degradation on source tasks, with average success rates dropping near zero across most transfer directions. This highlights a fundamental tension between adaptation and retention in current neural architectures.

We conclude that effective pretraining strategies must consider both forms of transfer value. Because adaptation efficiency benefits arise regardless of task complexity, either foundational or high-level tasks may serve as effective starting points for fast learning. However, when the goal is to maximize final performance, high-level tasks offer a distinct advantage. These insights support the design of curricula that strategically sequence tasks to balance speed and capability. Additionally, the severe forgetting observed during sequential fine-tuning emphasizes the need for architectural or algorithmic advances to preserve useful representations over time, especially in hierarchical task settings.

8 Future Work and Lessons Learnt

8.1 Future Work

Several promising research directions could extend and enhance the framework established in this thesis:

1. **Scaling to larger task hierarchies:** Our three-task hierarchy provided a controlled testbed, but expanding to broader and deeper skill taxonomies would strengthen generalizability claims. Future work should explore more diverse object types, environmental conditions, and manipulation primitives to create a more comprehensive map of cross-task transfer.
2. **Real-world validation:** While simulation offers controlled experimentation, validating our findings on physical robot systems would address questions about embodiment effects, sensor noise, and the persistence of observed transfer patterns under real-world constraints.
3. **Architectural innovations for retention:** The severe catastrophic forgetting observed across hierarchically related tasks demands specialized architectural solutions. Future work should investigate memory consolidation mechanisms, modular network designs, or replay-based techniques specifically targeting representational stability in manipulation settings.

4. **Automated curriculum design:** Our transfer metrics could inform algorithms that automatically sequence learning experiences to maximize transfer benefits. By quantifying which tasks serve as effective stepping stones to others, future systems could dynamically construct optimal learning curricula.
5. **Cross-embodiment transfer:** Extending our framework to study transfer across different robot morphologies would reveal whether representational utility generalizes across embodiments, a crucial question for scalable robot learning.

8.2 Lessons Learnt

This research has yielded several novel insights that contribute to the broader understanding of transfer learning in robotic manipulation:

1. **Complementary roles of foundational and high-level tasks:** We identified a functional distinction between tasks that contribute to fast adaptation and those that elevate final task performance. While both foundational and complex tasks improve adaptation efficiency, only high-level tasks consistently enhance maximum performance on simpler tasks. This distinction offers a new lens for evaluating task utility beyond traditional generalization metrics.
2. **Asymmetry in performance transfer, not in learning speed:** Contrary to prior assumptions, adaptation efficiency is symmetric across the task hierarchy: pretraining on either simple or complex tasks accelerates learning on others. However, performance transfer is asymmetric, only high-level tasks confer improvements in final success rates when transferred to simpler tasks. This nuanced view challenges conventional interpretations of task hierarchies in robotics.
3. **Representational stability is not guaranteed by task similarity:** Despite conceptual and structural overlap between tasks, we observed significant degradation in performance on previously learned tasks after fine-tuning. This indicates that catastrophic forgetting can occur even within closely related task families, highlighting the need for explicit retention mechanisms in lifelong learning systems.
4. **Multiple transfer metrics are essential for insight:** Our dual-metric evaluation framework, adaptation efficiency and maximum performance improvement, enabled us to capture different dimensions of transfer value. This approach reveals patterns that would be invisible under a single evaluation criterion and underscores the importance of multidimensional analysis in transfer learning research.

These lessons extend beyond our specific task hierarchy, offering practical guidance for the design of transfer learning algorithms, pretraining curricula, and memory-stable architectures in robotic manipulation.

Acknowledgments

I am deeply grateful to my supervisor, Dr. Boyu Wang, for his trust in allowing me the freedom to explore and define my own research direction and checking in on me throughout the process. I also thank Dr. Nazim Madhavji and Lisa Ann Moszczynski for enabling me to formally pursue this thesis despite administrative hurdles, their support made this thesis possible. I appreciate Tony Zhao and Kevin Rohling for openly sharing their ACT implementations, which served as a crucial starting point for this work.

To Aminah Merchant, thank you for your steady encouragement, thoughtful conversations, and for being a sounding board when the ideas were fragile and in need of clarity. Your presence made the intellectual solitude of research feel less solitary.

To my parents, thank you for your unwavering belief, for your understanding, and for the quiet strength you have always offered. This work, like every other thing in my life, rests on the foundation of your sacrifices.

A Resources

Models and Datasets: The models and datasets we produced as a part of this project are available for download on the HuggingFace Hub:

- <https://huggingface.co/collections/namrapatel/gen-value-67e201296ab798eedc839116>

Codebase: <https://github.com/namrapatel/gen-value>

Docker Image: A Docker image deployed from the most recent commit to the GitHub codebase is available at: <https://hub.docker.com/r/namra2000/act-training>

B Additional Experimental Details

Table 1: Hyperparameters used for all ACT model training

Parameter	Value
Learning rate	1e-5
Batch size	8
Chunk size	100
Hidden dimension	512
Feedforward dimension	3200
KL weight	10
Encoder layers	4
Decoder layers	7
Number of epochs	10,000
Seed	0

C Task-Specific Subtask Performance

Table 2: Success rate (%) for Grasp Cube Task. For each cell, left of vertical bar shows results at 500 epochs, right shows 10000 epochs.

Model	Grasp Cube	
	Touched	Grasped
Grasp (from scratch)	74 88	44 82
Pick & Place → Grasp	100 100	92 92
Transfer → Grasp	0 90	0 86

Table 3: Success rate (%) for Pick And Place Task. For each cell, left of vertical bar shows results at 500 epochs, right shows 10000 epochs.

Model	Pick And Place			
	Touched	Lifted	Placed	Released
Grasp → Pick & Place	96 100	82 82	78 80	78 80
Pick & Place (from scratch)	52 100	36 74	24 66	24 66
Transfer → Pick & Place	0 98	0 70	0 70	0 70

Table 4: Success rate (%) for Transfer Cube Task. For each cell, left of vertical bar shows results at 500 epochs, right shows 10000 epochs.

Model	Transfer Cube			
	TouchedRight	LiftedRight	TouchedLeft	Transferred
Grasp $\rightarrow$ Transfer	80 90	42 38	24 36	24 36
Pick & Place $\rightarrow$ Transfer	90 90	44 48	38 40	38 40
Transfer (from scratch)	86 88	46 48	10 44	10 44

These subtasks highlight where representational transfer is most beneficial, or where it breaks down. For Grasp Cube, Pick & Place pretraining yields immediate benefits across all components, while Transfer pretraining shows delayed gains. For Transfer Cube itself, grasping and lifting remain key failure points, reflecting a need for more targeted representation learning.

D Learning Curves

To visualize the adaptation efficiency patterns discussed in the main text, Figures 7–9 show the validation loss curves for all transfer learning experiments. These plots directly illustrate how different pretraining sources affect convergence rates on each target task.




Figure 7: Validation loss when fine-tuning on Grasp Cube task. The red line shows training from scratch, while blue and green lines represent pretraining on Pick & Place and Transfer Cube, respectively.




Figure 8: Validation loss when fine-tuning on Pick & Place task. The red line shows training from scratch, while yellow and green lines represent pretraining on Grasp and Transfer Cube, respectively.




Figure 9: Validation loss when fine-tuning on Transfer Cube task. The red line shows training from scratch, while yellow and blue lines represent pretraining on Grasp and Pick & Place, respectively.

These curves provide visual confirmation of the adaptation efficiency metrics reported in Figure 4. Models pretrained on tasks higher in the hierarchy (e.g., Transfer Cube) initially converge faster when fine-tuned on simpler tasks (e.g., Grasp). Similarly, foundational tasks like Grasp provide strong adaptation benefits when used as pretraining for more complex tasks like Pick & Place.

References

Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., ... Zitkovich, B. (2023,

- July). RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation. Retrieved 2025-04-05, from <http://www.roboticsproceedings.org/rss19/p025.pdf> doi: 10.15607/RSS.2023.XIX.025
- Chi, C., Xu, Z., Feng, S., Cousineau, E., Du, Y., Burchfiel, B., ... Song, S. (2024, October). Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 02783649241273668. Retrieved 2025-04-06, from <https://doi.org/10.1177/02783649241273668> (Publisher: SAGE Publications Ltd STM) doi: 10.1177/02783649241273668
- Ghosh, D., Walke, H., Pertsch, K., Black, K., Mees, O., Dasari, S., ... Levine, S. (2024, July). Octo: An Open-Source Generalist Robot Policy. In *Robotics: Science and Systems XX*. Robotics: Science and Systems Foundation. Retrieved 2025-04-05, from <http://www.roboticsproceedings.org/rss20/p090.pdf> doi: 10.15607/RSS.2024.XX.090
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In (pp. 770–778). Retrieved 2025-04-05, from https://openaccess.thecvf.com/content_cvpr_2016/html/He_Deep_Residual_Learning_CVPR_2016_paper.html
- Hejna, J., Bhateja, C. A., Jiang, Y., Pertsch, K., & Sadigh, D. (2025, January). ReMix: Optimizing Data Mixtures for Large Scale Imitation Learning. In *Proceedings of The 8th Conference on Robot Learning* (pp. 145–164). PMLR. Retrieved 2025-04-05, from <https://proceedings.mlr.press/v270/hejna25a.html> (ISSN: 2640-3498)
- Kim, M. J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., ... Finn, C. (2025, January). OpenVLA: An Open-Source Vision-Language-Action Model. In *Proceedings of The 8th Conference on Robot Learning* (pp. 2679–2713). PMLR. Retrieved 2025-04-05, from <https://proceedings.mlr.press/v270/kim25c.html> (ISSN: 2640-3498)
- Kingma, D. P., & Welling, M. (2022, December). *Auto-Encoding Variational Bayes*. arXiv. Retrieved 2025-04-05, from <http://arxiv.org/abs/1312.6114> (arXiv:1312.6114 [stat]) doi: 10.48550/arXiv.1312.6114
- Pan, S. J., & Yang, Q. (2010, October). A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359. Retrieved 2025-04-04, from <https://ieeexplore.ieee.org/document/5288526> (Conference Name: IEEE Transactions on Knowledge and Data Engineering) doi: 10.1109/TKDE.2009.191
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is All you Need. In *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc. Retrieved 2025-04-05, from https://proceedings.neurips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html
- Zamir, A. R., Sax, A., Shen, W., Guibas, L., Malik, J., & Savarese, S. (2018, June). Taskonomy: Disentangling Task Transfer Learning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 3712–3722). Retrieved 2025-04-04, from <https://ieeexplore.ieee.org/document/8578489> (ISSN: 2575-7075) doi: 10.1109/CVPR.2018.00391
- Zhao, T., Kumar, V., Levine, S., & Finn, C. (2023, July). Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation. Retrieved 2025-04-04, from <http://www.roboticsproceedings.org/rss19/p016.pdf> doi: 10.15607/RSS.2023.XIX.016