

Scaling Trusted Autonomy: Towards Generally Capable Verifiable Robots

by
Benedict Quartey

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Computer Science at Brown University

Providence, Rhode Island
October 2026

© Copyright 2026 by Benedict Quartey

This dissertation by Benedict Quartey is accepted in its present form by
the Department of Computer Science as satisfying the dissertation requirement
for the degree of Doctor of Philosophy.

Date _____

George Konidaris, Advisor

Recommended to the Graduate Council

Date _____

Stefanie Tellex, Reader

Date _____

Srinath Sridhar, Reader

Approved by the Graduate Council

Date _____

David P. Lindstrom
Dean of the Graduate School

For Aleen, my beautiful wife, and my parents, Janis and Ignatius Quartey

Abstract of “Scaling Trusted Autonomy: Towards Generally Capable Verifiable Robots” by Benedict Quartey, Ph.D., Brown University, October 2026.

Robots powered by foundation models are increasingly able to interpret open-ended language, recognize unfamiliar objects, and operate across diverse environments. Yet broader capability does not by itself establish that a robot has understood a request, that its actions will produce the intended effects, or that its internal representation still corresponds to reality. Classical symbolic planning provides explicit structure for composition and verification, but its application to robotics has traditionally depended on hand-engineered skills and representations specified in advance. Learned policies can broaden the robot’s repertoire, but provide weak guarantees about what will happen and why. This thesis argues that scaling trusted autonomy requires a neurosymbolic approach that unifies the generality of learned models with the symbolic structure needed for robust robot behavior. It explores three complementary directions toward this goal. First, LIMP uses foundation models and spatial reasoning to ground natural-language instructions in explicit task constraints and instruction-specific semantic maps, enabling task-and-motion planning while exposing the robot’s interpretation for inspection. Second, REALM extends this framework to an evolving world. It maintains an editable, photorealistic 3D Gaussian scene, predicts the consequences of planned manipulation, and locally reconciles the scene so that subsequent instructions obtain their grounding and planning geometry from an updated state. Third, PACTS learns visuomotor skills that jointly generate action trajectories and predictions of their symbolic effects, supporting skill composition, progress monitoring, and recovery through replanning. Together, these contributions provide mechanisms for constructing task-relevant abstractions when needed, connecting learned behavior to its predicted symbolic effects, and maintaining an evolving world representation for subsequent tasks. This body of work preserves explicit structure for reasoning and verification while broadening how it is obtained, providing a foundation for generally capable robots whose behavior can be assessed under the assumptions of their models.

Acknowledgments

Long before this dissertation took shape, there were people who believed in me and helped me find my way into research. I am deeply grateful to them, and to everyone who has made these years so meaningful.

My advisor, George Konidaris, took a chance on a young man from Ghana during a pandemic that shook the entire world. We had never met in person, and my path into research had been far from conventional. Yet, George looked beyond what could be seen on paper and believed in what I could become. He gave me the opportunity to come to Brown and supported me at every stage of this PhD. I am eternally grateful for his support, advice, and mentorship.

Stefanie Tellex has been an extraordinary mentor at Brown and beyond. Stefanie gave me a shot to do research at the Robotics and AI Institute, believing in my ideas before they could materialize; her encouragement and mentorship have been invaluable. I also thank Srinath Sridhar, whose lessons and intuitive explanations gave me the foundations to pursue the vision work in this thesis, and Michael Littman, who introduced me to reinforcement learning in my first year and served on my first candidacy committee.

Many years ago during my undergraduate studies at Ashesi, Dr. Ayorkor Korsah introduced me to research and helped me see a PhD as a possibility. Thank you for encouraging that curiosity before I knew where it might lead.

I am also very grateful to Eric Rosen, who worked with me on my first major PhD paper; I learned a great deal from our collaboration. Thank you for sharing your time and knowledge so generously. To Sebastian Castro, Wil Thomason, and my colleagues at the Robotics and AI Institute: thank you for making my time there so rewarding. Working together was a pleasure and I will always value the friendships that grew alongside research.

To my friends, labmates, and colleagues in the Intelligent Robot Lab and the Humans to Robots Lab, thank you for the conversations, shared enthusiasm, and laughter that made Brown a place where I belonged.

Wasiwasi and Amina, we began our PhDs together as three Africans finding our way in a new place. You have grown from friends into family, and I am grateful that we have shared so much of this experience. Tuluhan and Yong, our friendship began even before our Ph.D. journey. Being roommates and spending those early years together gave me brothers, not just friends. Thank you for the companionship and all the memories I will carry with me.

To my mom and dad, thank you for the love and faith that have sustained me throughout these years. You nourished my curiosity from childhood and went above and beyond, sacrificing so much to give me the opportunities that led me here. Thank you so much. Living so far from home has made me appreciate even more how much of who I am comes from you. I hope you can feel your part in this achievement. To my friends near and far, thank you for keeping me connected to the people and places I love.

Finally, to my wife, Aleen: thank you for your patience through the long nights and for the love that has filled these years with happiness. So much of the joy and enthusiasm that others see in me begins with you. You have shared the demands of this PhD with a generosity I will never take for granted. I am grateful that I get to share this moment, and my life, with you. I love you. This dissertation is for you and for our family. Thank you Jesus, we made it!

Contents

List of Tables	xii
-----------------------	------------

List of Figures	xiii
------------------------	-------------

1 Introduction	1
1.1 From Tasks to Robot Behavior	2
1.1.1 Explicit Structure and Compositional Planning	2
1.1.2 Learning Behavior from Experience	4
1.2 The Capability–Assurance Gap	5
1.3 Three Interfaces for Trusted Autonomy	6
1.3.1 Intent: Did the Robot Understand What the User Meant?	6
1.3.2 Action: Did the Robot’s Actions Produce the Intended Effects?	6
1.3.3 World: Does the Robot’s Internal World Still Match Reality?	7
1.4 Thesis and Contributions	7
1.4.1 LIMP: Grounded Language Understanding	8
1.4.2 REALM: Evolving Representations of the World	8
1.4.3 PACTS: Composable State-Reporting Skills	8
2 Background	10
2.1 Foundation Models and Broader Generality	10
2.1.1 Language and Visual Representations	11
2.1.2 From Pretrained Knowledge to Robot Behavior	11
2.2 Symbolic State, Planning, and Task Specifications	12
2.2.1 Predicates and Action Models	12
2.2.2 Requirements over Time	13
2.2.3 Automata, Synthesis, and the Scope of Verification	14
2.3 Task and Motion Planning with Reusable Skills	15
2.3.1 Geometric Feasibility	15
2.3.2 Skills as Temporally Extended Actions	16
2.4 Semantic Mapping and Spatial Grounding	16
2.4.1 From Images to a Metric Scene	17

2.4.2	Resolving Objects, Relations, and Identity	17
2.5	World Models and State Change	18
2.6	Gaussian Scene Representations	19
2.6.1	Gaussian Primitives and Image Formation	19
2.6.2	Object Motion, Local Updates, and Inspection	20
2.7	Learning Visuomotor Skills from Demonstrations	21
2.7.1	Behavior Cloning and Closed-Loop Execution	21
2.7.2	Denoising Diffusion	22
2.7.3	Conditional Flow Matching	22
2.8	Symbolic Interfaces for Composition and Monitoring	23
2.8.1	State Estimates and Predicted Outcomes	23
2.8.2	Composition, Progress, and Recovery	24
3	Grounding Language in the Real World	25
3.1	From an Instruction to a Grounded Plan	25
3.2	Related Work	27
3.2.1	Temporal Specifications for Instruction Following	27
3.3	Problem Definition	28
3.4	Language Instruction Grounding for Motion Planning	28
3.4.1	Language Instruction Module	28
3.4.2	Spatial Grounding Module	30
3.4.3	Task and Motion Planning Module	32
3.5	Evaluation	32
3.6	Discussion	34
3.7	From Grounded Intent to Continued Autonomy	34
4	Evolving World Representations for Planning	36
4.1	Planning Across World Changes	36
4.2	Related Work	38
4.3	One World, Many Instruction-Specific Views	39
4.4	REALM	41
4.4.1	From the Persistent World to the Current Instruction	42
4.4.2	Grounding Says What the Task Means; Ownership Says What Moves	42
4.4.3	Predict Physical Change (Known from the Plan)	43
4.4.4	Reconcile Visual Consequences (Inferred from Evidence)	44
4.4.5	Commit (Trusted for Reuse)	45
4.5	Following the World Across Two Instructions	46
4.5.1	Timestep 0: Build the First Task from W_0	46
4.5.2	From the First Plan to the Next World	46
4.5.3	Timestep 1: Rebuild the Second Task from W_1	46

4.5.4	What the Sequence Must Preserve	47
4.6	Results	47
4.6.1	Grounding Becomes Action Support Through Ownership	47
4.6.2	Known, Inferred, and Trusted States Remain Local and Reversible	47
4.6.3	The Next Instruction Rebuilds Its World from W_1	48
4.6.4	History Preserves How the Current World Was Reached	49
4.7	Discussion	49
4.7.1	From Verifiable Plans to an Evolving World	49
4.7.2	One Representation, Four Connected Uses	50
4.7.3	From a Renderable Scene to a Spatial World Model	50
4.7.4	From an Evolving World to Composable Skills	50
5	Learning Composable Skills	52
5.1	Learning Actions Together with Their Symbolic Outcomes	52
5.2	Related Work	53
5.3	Jointly Modeling Predicates and Actions	54
5.3.1	Problem definition: joint action–outcome skill model	54
5.3.2	DDPM formulation	55
5.3.3	Conditional Flow Matching formulation	56
5.3.4	Architecture	56
5.4	Planning-based Skill Composition with Predicate Beliefs	56
5.4.1	Skill segmentation and labeling toolkit.	57
5.5	Evaluation	58
5.5.1	Tasks and datasets	58
5.5.2	Baselines and ablations	59
5.5.3	Metrics	60
5.6	Results and Discussion	61
5.7	Limitations	64
5.8	From Composable Skills to Scaling Trusted Autonomy	64
6	Scaling Trusted Autonomy	65
6.1	From Greater Capability to Trusted Autonomy	65
6.2	Constructing the Abstractions a Task Requires	66
6.3	Trust Across Intent, Action, and World	68
6.4	Toward Continued Autonomy	69
6.4.1	Recognizing When an Abstraction Is Insufficient	70
6.4.2	Connecting Predictions to Observations	70
6.4.3	Expanding Skills and Their Symbolic Meaning	71
6.4.4	Managing Uncertainty Across the Interfaces	71
6.5	Concluding Perspective	72

A	Supplementary Methods for LIMP	73
A.1	Overview	73
A.2	Language Instruction Module	73
A.2.1	Interactive Symbol Verification	75
A.3	Spatial Grounding Module	77
A.3.1	3D Spatial Comparators	77
A.4	Task and Motion Planning Module	78
A.5	Robot Skills	81
A.6	Evaluation and Baseline Details	81
A.6.1	NLMap-Saycan Implementation Prompt	83
A.6.2	Code-as-Policies Implementation Prompt	84
A.6.3	Instruction set	86
B	Supplementary Methods for REALM	91
B.1	State Lifecycle and Promotion	91
B.2	Ownership Diagnostics	92
B.3	Task Specifications and Plan Evidence	92
B.4	Observation and Reconciliation Protocol	93
B.5	Complete State Sequence	93
	Bibliography	95

List of Tables

3.1	Performance comparison of one-shot instruction understanding and spatial resolution.	34
3.2	Performance comparison of one-shot temporal alignment and plan success rate.	34
4.1	Design requirements for an evolving robot world.	42
4.2	Ownership selector analysis for the garden ball.	47
4.3	Spatial and visual continuity across the two world transitions	48
5.1	Single-policy rollout evaluation on PushBarrier	61
5.2	Single-policy rollout evaluation on RoboMimic	61
B.1	Logical contents of a persistent world revision.	91
B.2	The five-state mainline produced by the sequential garden study.	94

List of Figures

1.1	Explicit representations for the book-transfer task	3
1.2	Learning the book-transfer behavior from demonstrations	4
2.1	An automaton for an ordered-visitation requirement	14
3.1	Grounded instruction following with LIMP	26
3.2	The LIMP language-to-plan framework	29
3.3	Two-stage translation from language to task specifications	30
3.4	Spatial grounding of instruction referents	30
3.5	Task and motion planning with task progression semantic maps	31
4.1	A robot can invalidate the map it just used	37
4.2	One persistent world, rebuilt for each instruction	43
4.3	How the world evolves	44
4.4	Ownership defines what the action moves	48
4.5	Sequential world evolution across two instructions	49
5.1	Predicate-Action Skills	53
5.2	Architectures and baselines	55
5.3	Evaluation environments	58
5.4	Predicate-action coherence of predicted belief rollouts	62
5.5	Skill recomposition and adversarial recovery in PushBarrier	62
6.1	The three trust interfaces and their complementary contributions	68
A.1	Example plans for navigation and mobile pick-and-place	82
A.2	Plan success with interactive reprompting	82

Chapter 1

Introduction

A central ambition of robotics is to build machines that can perform a wide range of useful tasks in the environments people inhabit. Such robots would allow people to specify goals and constraints without prescribing each movement or engineering a separate system for every activity. Their usefulness would depend on the ability to interpret unfamiliar requests, apply learned abilities in new circumstances, and sustain purposeful behavior as their surroundings change. The challenge is to make this breadth of capability compatible with behavior that people can understand, assess, and rely upon.

Foundation models have opened new possibilities for pursuing this goal. Their training on broad collections of data allows knowledge acquired across tasks and settings to inform the behavior of individual robot systems (Bommasani et al., 2021). They support interfaces through which people can describe tasks and objects beyond a fixed vocabulary, and connect pretrained language and visual knowledge to robot control across diverse tasks and settings (Minderer et al., 2022; Zitkovich et al., 2023; Kim et al., 2025). This progress changes what can reasonably be asked of a robot. It also brings a longstanding requirement into sharper focus: as robots undertake a broader range of activities, there must be a corresponding basis for assessing whether their behavior meets the requirements of those activities.

Greater capability does not automatically provide that assurance. A robot may respond fluently to an instruction while leaving its interpretation implicit, or produce effective movements without exposing the conditions those movements establish. The consequences matter when people delegate tasks in homes and workplaces, where intentions must be communicated, constraints respected, and decisions made from incomplete information. The scientific problem is therefore to extend the scope of autonomous behavior while retaining the means to examine what the robot understood, reason about what it will do, and assess whether continued reliance on its decisions is warranted. Generality and assurance must advance together.

Classical planning provides an important foundation for this objective. Explicit task representations and action models support the composition of behaviors, while formal specifications make requirements available for checking and behavior synthesis (Fikes and Nilsson, 1971; Kress-Gazit et al., 2009). Task and motion planning connects this reasoning to geometric feasibility (Garrett

et al., 2021). These strengths, however, depend on suitable descriptions of the task, the robot’s skills, and the world. Constructing and maintaining those representations across varied environments has traditionally required substantial engineering, limiting the settings in which structured reasoning can be applied. Learned models offer ways to broaden perception and behavior; planning offers structure for reasoning about their use. The challenge is to bring these strengths together without allowing the assumptions that support reliable reasoning to become hidden inside increasingly capable components.

The central research question of this thesis is therefore: *how can robot capability be extended across tasks and environments while preserving the explicit structure needed to interpret, compose, and check its behavior?* This thesis develops an answer through grounded language understanding, composable state-reporting skills, and evolving representations of the world. These contributions connect learned capabilities to representations that a planner can inspect and use.

1.1 From Tasks to Robot Behavior

Two complementary approaches help explain both the opportunity and the difficulty of this problem. Planning constructs behavior by reasoning over a model of the task and environment, whereas policy learning acquires behavior from experience. The former makes the organization of behavior explicit; the latter provides a way to acquire the perception and control needed to realize it. Understanding their respective strengths requires examining how each turns a requested outcome into physical action, and how knowledge acquired for one task becomes useful in another.

1.1.1 Explicit Structure and Compositional Planning

A physical scene contains far more information than is needed for any particular decision. Symbolic planning makes reasoning manageable by representing the objects, properties, and relationships relevant to a task, together with models of the actions that can change them (Fikes and Nilsson, 1971). Predicates express conditions that can be evaluated as true or false, such as whether an object is on a surface or held by the robot. An action model specifies the conditions under which an action is applicable and the effects it is expected to produce. The initial state describes the conditions that currently hold, while the goal describes those the planner must establish. This abstraction allows the planner to reason about changes in the world without representing every detail of their physical realization.

For example, the instruction “bring the red book on the table to the sofa” can be expressed using objects for the book, table, and sofa, and predicates describing the book’s location and whether the robot is holding it. Picking up the book might require an available gripper and a reachable book; its modeled effects include holding the book and removing it from the table. The planner can then search for a sequence in which the effects of each action establish the conditions required by the next: approach the table, pick up the book, approach the sofa, and place the book. Figure 1.1 illustrates this connection between a requested outcome and a compositional description of behavior.




Figure 1.1: Explicit representations for the book-transfer task. A symbolic description connects the observed scene and desired goal to a sequence of task-level actions, whose physical realization requires motion planning and control. The book-and-vase extension illustrates how the same action models support a new composition. Learned policies can implement individual skills within this structure.

The value of this representation extends beyond finding one successful sequence. If the request changes to transporting both the book and the vase, the planner can reuse the same action models with different object instances and arrange their applications to satisfy the new goal. There is no need to store a separate procedure for every combination of familiar objects and goals: the model describes reusable relationships between conditions and actions (Konidaris et al., 2018). This form of composition gives structured planning its own source of generality: a domain can support many tasks that were never individually specified as complete behaviors.

Physical realization requires a further connection between discrete decisions and continuous control. The robot must select grasps, configurations, and collision-free motions that support the proposed sequence; a failed geometric search may require the task plan itself to change. Task and motion planning reasons across these levels (Garrett et al., 2021). Its actions can be implemented through reusable skills, such as grasping or navigation, that unfold over multiple control steps and have conditions governing when they can begin and terminate (Sutton et al., 1999). A skill may use an analytic controller, a motion planner, or a learned policy. Planning and learning can therefore contribute to the same behavior, with planning organizing the skills and learning supplying some of their control policies.

Explicit task structure also permits requirements to shape the course of an activity. For the book-transfer task, reaching the sofa while avoiding a specified region is a different requirement from merely ending at the sofa. A final-state goal captures the destination but does not, by itself, express every constraint on the intervening behavior. Formal methods ideas like temporal logic provide a language for requirements such as ordering and avoidance, and support the synthesis of behavior that satisfies a specification under a given model (Pnueli, 1977; Kress-Gazit et al., 2009). Planners can thus reason about both what should be achieved and which sequences of events are acceptable.

The reach of this reasoning depends on the representations supplied to it. The planning system needs a formal task grounded in the relevant objects and skills equipped with appropriate action models. A prebuilt semantic map can supply associations between spatial structure and meaning, but future requests may refer to properties or objects that were not recorded when the map was created. Such a map is not inherently closed-world: that assumption arises when represented facts are treated as complete and unrepresented facts as false. Even a detailed map may require revision as the environment changes. Extending planning to less structured settings therefore requires ways to construct and maintain its inputs without engineering them anew for every task.

1.1.2 Learning Behavior from Experience

Learning from demonstration offers a way to acquire control from examples of successful behavior. For the same book-transfer task, an expert supplies trajectories containing observations and robot actions, from which a visuomotor policy learns to produce actions in response to visual input and, where available, robot state and a task description (Mandlekar et al., 2022). Figure 1.2 illustrates this route from demonstrated experience to a control policy. During deployment, a closed-loop policy repeatedly receives current observations, allowing its movements to respond to the scene as execution proceeds. The acquired behavior is consequently a mapping from situations to actions, rather than a recording of one sequence of movements.




Figure 1.2: Learning the book-transfer behavior from demonstrations. Observation–action examples train a policy that produces robot actions. The schematic emphasizes behavior acquisition; a deployed closed-loop policy receives current observations. Comparing this output with Figure 1.1 identifies the additional task and symbolic-state information needed to use the learned behavior within a planner.

A useful control representation must accommodate the fact that several different movements can accomplish the same task. There may be multiple valid approaches to an object, and averaging their trajectories need not yield a valid approach. Generative policies address this multiplicity by modeling distributions over actions or action sequences. Diffusion Policy, for example, generates observation-conditioned action trajectories and executes them through receding-horizon control (Chi et al., 2025). By generating a trajectory, executing part of it, and then responding to another observation, the policy combines coordinated movement with feedback. Such methods provide a basis for acquiring manipulation skills whose detailed motions would be difficult to specify by hand.

The same policy can also cover several tasks when its output is conditioned on the requested behavior. A task identifier can select an activity, a goal image can depict an outcome, and language can describe what should be accomplished. In the shared example, conditioning distinguishes bringing only the book from bringing both the book and the vase. How well the policy responds to an unfamiliar request depends on the coverage of its training experience and its ability to generalize beyond that experience. This makes the source and organization of training data central to the breadth of learned behavior.

Longer activities place additional demands on that behavior. Errors can change the observations encountered later in an execution, taking a policy into situations poorly represented in its demonstrations (Ross et al., 2011). Feedback may correct a deviation within a skill, while recovery at the task level may require selecting a different skill or revising their order. A planner making that decision needs to know which relevant conditions hold and which actions can restore progress. An action-producing policy may use such information internally without exposing it in a form that another component can interpret. The ability to perform a behavior and the ability to supply the information needed to compose it are therefore distinct aspects of a learned skill.

1.2 The Capability–Assurance Gap

Planning allows familiar actions to be composed for new goals; learning broadens the behavior a robot can acquire from experience. Both extend what the robot can do, but neither alone establishes whether a particular behavior satisfies the requirements of the task. Capability concerns what a robot can accomplish and how broadly that competence extends across tasks and conditions. Assurance concerns the evidence and reasoning that support confidence in its behavior under stated conditions. For example, recognizing the objects mentioned in an instruction does not establish that the robot has resolved the instructor’s intended referents, and a high task success rate does not establish compliance with every constraint in a particular execution. An account of dependable autonomy must connect evidence about individual capabilities to the requirements of the activity in which they are used.

Symbolic representations make these requirements available for reasoning, but the conclusions of that reasoning remain conditional on the model. A plan can satisfy a specification even when the specification misrepresents the instruction; similarly, a valid transition in an action model may fail to describe what occurred during execution. Closing the capability–assurance gap requires connecting an instructor’s intended task to the robot’s interpretation, and the effects anticipated by the planner

to what happens during execution. The interfaces between perception, planning, and control must expose enough information to check these connections.

This thesis uses *trusted autonomy* to describe the objective of making reliance on task-level robot behavior warranted by such evidence and reasoning. The term *verifiable* refers to explicit specifications and representations against which aspects of behavior can be checked. Formal reasoning establishes properties within a model, while empirical evaluation examines how well grounding, control, and state predictions correspond to the environment. These forms of evidence support different claims; together, they identify both the basis for a decision and the assumptions on which that basis depends.

1.3 Three Interfaces for Trusted Autonomy

These connections raise three questions about the information on which a robot relies. Before planning, the robot needs a representation of what the user intends. During execution, it needs to determine whether the effects of its actions support continued progress. Across actions and successive tasks, it needs to maintain the world state on which further interpretation and planning depend. These responsibilities define the *Intent*, *Action*, and *World* interfaces developed in this thesis. Each interface addresses a different part of the same activity: interpreting the task, assessing progress, and maintaining the context for further decisions.

1.3.1 Intent: Did the Robot Understand What the User Meant?

Language communicates a task through descriptions whose meaning depends on the environment and on their relationships to one another. In the book example, “the red book on the table” identifies an instance through an attribute and a spatial relation. An instruction may additionally constrain the order of operations or specify what the robot must avoid. Grounding must connect these expressions to objects, locations, and behavioral requirements, so that the planner reasons about the requested task.

An explicit task specification exposes the interpreted requirements, while grounded referents make the selected objects and locations available for inspection. This permits the interpretation to be examined separately from the plan constructed to satisfy it. The research problem at the intent interface is to obtain these representations from expressive language and perceptual observations without assuming that every relevant concept was encoded in a complete semantic map beforehand. The corresponding contribution is *grounded language understanding*.

1.3.2 Action: Did the Robot’s Actions Produce the Intended Effects?

Composition relies on the effects of one skill establishing the conditions required by another. For the robot to carry the book, its preceding grasp must establish that the book is held; a command to close the gripper is not sufficient evidence of that condition. The planner needs a state description at the level of its action models, even when the skill itself operates through images and continuous motor

commands. This creates an interface between the quantities used for control and the conditions used for task reasoning.

For learned behavior, the research problem is to acquire skills that expose those conditions as execution unfolds. The thesis develops *composable state-reporting skills*, which generate control actions together with predictions of the associated symbolic state. Such reports allow a planner to monitor expected effects and revise a skill sequence using updated state estimates. Their usefulness depends on how well they agree with the consequences of the generated actions: evaluating predicate accuracy alone does not establish action–predicate coherence. The symbolic report must therefore be studied as part of the skill’s behavior.

1.3.3 World: Does the Robot’s Internal World Still Match Reality?

Task-level state reports do not by themselves maintain the spatial context required for sustained autonomy. A later instruction may refer to an object by its relationship to other objects, while geometric planning needs its current shape and location. Although these decisions use different abstractions, they must draw on a compatible account of the environment. As the robot intervenes in that environment, the representation must preserve relevant changes, object identity, and unaffected scene structure across tasks.

The research problem at the world interface is to maintain that continuity through a representation that can predict task-relevant changes and supply updated inputs to grounding and planning. This motivates *evolving representations of the world*, in which action-conditioned transitions connect the state used for one task to the state available for the next. The world interface provides the persistent spatial context that complements the action interface’s reports of skill-level conditions. For these interfaces to support continued operation in the physical world, predicted effects must be reconciled with observations, so that later decisions reflect what actually happened.

1.4 Thesis and Contributions

The central argument of this dissertation is:

Scaling trusted autonomy requires unifying grounded language understanding, composable state-reporting skills, and evolving representations of the world.

The common design principle is to introduce explicit symbolic and spatial structure where learned outputs become inputs to robot reasoning. At these interfaces, a system can expose its interpreted requirements, expected action effects, and persistent world state for inspection, composition, and checking. The contributions develop this principle through three complementary forms of generality: broader instruction grounding, continuity across world changes, and new compositions of learned skills.

The technical chapters follow the order **LIMP** → **REALM** → **PACTS**. LIMP establishes the language-to-plan framework and exposes two assumptions that limit its continued use: a world representation that remains valid after mapping, and an available library of skills with suitable planning

interfaces. REALM addresses the first by evolving the world state used for subsequent tasks; PACTS addresses the second by learning skills that expose planning-relevant state.

Chapter 2 develops the shared foundations needed to examine the contributions: foundation models, task specifications, task and motion planning, semantic grounding, spatial world representations, and learned visuomotor skills. These concepts provide the basis for the technical chapters, where each method is also situated within its related work. Chapters 3–5 present LIMP, REALM, and PACTS in the dependency order described above. The concluding Chapter 6 brings these contributions together around the construction and maintenance of task-relevant abstractions. It returns to the capability–assurance question through the three trust interfaces and develops a research agenda for continued autonomy, including physical integration, uncertainty management, and recovery.

1.4.1 LIMP: Grounded Language Understanding

Chapter 3 introduces Language Instruction Grounding for Motion Planning (LIMP), which connects expressive instructions to explicit task constraints and spatially grounded plans. Language and vision-language models supply descriptions of the task and its referents, while spatial reasoning resolves object instances and connects the resulting specification to motion planning. The central advance is to construct the semantic information required by an instruction at runtime, using observations retained during metric mapping, while preserving an explicit representation of its temporal and spatial requirements. This makes learned semantic breadth available to structured planning without requiring a complete prebuilt semantic map.

1.4.2 REALM: Evolving Representations of the World

Chapter 4 develops REALM, an action-conditioned spatial world model that carries a persistent representation across successive tasks. It derives the semantic and geometric abstractions needed for grounding, planning, and rendered inspection from a common 3D Gaussian scene (Kerbl et al., 2023). Action-conditioned transitions predict geometric change and update the affected parts of the scene while preserving object identity and unaffected scene content. These predicted changes provide the world state used to ground and plan subsequent tasks. The scene becomes a world model through these predictive transitions and its reuse for later decisions. The contribution is a persistent spatial basis for planning whose evolution remains explicit.

1.4.3 PACTS: Composable State-Reporting Skills

Chapter 5 introduces Predicate-Action Skills (PACTS), which learn control and its symbolic interface through joint generation of action trajectories and temporally aligned predicate-belief trajectories. Conditioned on current observations, each generated sample pairs proposed actions with the symbolic conditions the skill expects as those actions unfold. These reports supply an off-the-shelf planner with state estimates for sequencing skills, checking effects, and replanning. They are refreshed through

observation-conditioned rollouts, but remain predictions rather than independent measurements of action success.

Together, these contributions address three requirements for extending robot capability while preserving a basis for assurance. Broader language understanding must remain connected to explicit task requirements; learned skills must expose the state needed to assess and compose their behavior; and the world representation used for planning must account for change. By making these connections explicit, this thesis develops a foundation for robots whose capabilities can expand while their interpretation and behavior remain open to examination.

Chapter 2

Background

Robots become more useful when they can carry what they have learned into tasks that were never individually programmed or demonstrated. The introduction argued for pursuing this breadth through a combination of learning and planning. Learned models can bring new instructions, objects, and skills within a planner’s reach, while planning can organize those capabilities to accomplish new goals. The challenge is to preserve the meaning of a task as it passes from a person’s request into the robot’s decisions and, ultimately, into physical action.

Making this connection requires choices about what the robot represents and what it can infer from those representations. A planner can reason that a grasp must precede a placement only if the consequences of grasping are expressed in terms it can use. Perception must establish which object those terms refer to, and execution must supply information about whether the conditions for continuing the task hold. The representations connecting these processes determine both how their capabilities can be combined and what can be checked about the resulting behavior.

This chapter develops the ideas needed to understand and construct those connections. It begins with the breadth of knowledge available through foundation models and examines how that knowledge can enter explicit reasoning about robot tasks. Following this connection from interpretation to action leads to the physical grounding of task specifications, the maintenance of a changing world, and the acquisition of skills that a planner can reuse. Together, these foundations explain how the opportunities for broader capability give rise to the interface problems addressed by the technical chapters.

2.1 Foundation Models and Broader Generality

The knowledge required for a robot task often extends beyond the robot’s own experience. A person may describe an object through its appearance, its use, or its relationship to another object, using words that were never included in a task-specific training set. Foundation models provide a way to bring knowledge learned from broad collections of data into such settings. The term describes models trained at scale on broad data and adapted to a range of downstream uses (Bommasani et al.,

2021). For robotics, the important change is that much of the representation learning can precede the definition of a particular task or deployment environment.

2.1.1 Language and Visual Representations

A language model can be asked to continue an ordinary sentence or to express an instruction in a formal language. An autoregressive model approaches both through the same prediction problem: generating the next token, a unit of text, from the prompt and the tokens already produced. With learned parameters θ , the probability of a sequence $w_{1:n}$ conditioned on a prompt c factorizes as

$$p_{\theta}(w_{1:n} \mid c) = \prod_{i=1}^n p_{\theta}(w_i \mid c, w_{1:i-1}). \quad (2.1)$$

The prompt may include an instruction, descriptions of available tools, and examples of the desired output. This makes it possible to request a structured expression, such as a program or a logical formula, using the same sequence-generation mechanism that produces ordinary text. In-context examples specify how the model should respond without themselves changing its trained parameters (Brown et al., 2020). They can guide the form and content of a translation, but the probability assigned to a token sequence is not a proof that the translation preserves the user’s meaning.

Visual representations supply a complementary connection to the scene. In an open-vocabulary detector, a text query can specify the object category or description to localize, rather than selecting only from a fixed list of detector labels. OWL-ViT, for example, adapts image-text pretraining to text-conditioned object detection (Minderer et al., 2022). Detection proposes image regions; segmentation identifies the pixels belonging to a candidate region. These are distinct operations: a box around a book may include part of the table, whereas a segmentation mask can provide more selective support for projecting that book into three dimensions (Kirillov et al., 2023).

Open vocabulary expands the descriptions that can be submitted to perception; it does not imply exhaustive recognition of every concept or reliable identification of every instance. Several books may match “red book,” and a correct category detection may still select the wrong book for a relational instruction. Language understanding, object detection, and referent resolution therefore contribute different information. The task becomes grounded only when the expressions in the instruction are associated with the relevant entities and conditions in the robot’s environment (Tellex et al., 2020).

2.1.2 From Pretrained Knowledge to Robot Behavior

Pretrained representations can enter robot behavior at different levels. Vision-language-action models connect language and visual knowledge to control learned from robot experience. RT-2 studies this transfer by incorporating robot actions into a vision-language model’s output representation (Zitkovich et al., 2023). OpenVLA develops policies from diverse robot demonstrations and supports adaptation to additional robot setups (Kim et al., 2025), while π_0 combines pretrained vision-language representations with continuous action generation (Black et al., 2025). In each case, robot data connects semantic knowledge to the actions available to a particular embodiment.

The same knowledge can also support explicit intermediate reasoning. SayCan combines language-model scores with learned skill value functions to select skills using estimates of their relevance to the instruction and their feasibility in the current state (Ichter et al., 2023). Code as Policies generates programs that call perception and control interfaces (Liang et al., 2023). Language-to-temporal-logic methods instead produce task specifications that can be passed to a planner (Liu et al., 2023). These approaches differ in what the learned model supplies: motor commands, skill choices, executable code, or formal requirements. That choice determines which decisions remain available for separate examination.

A language-conditioned action policy can use task meaning without expressing it as a checkable specification. Conversely, a generated specification can be checked for logical properties while still referring to the wrong object. The value of an explicit interface is that it separates these questions and makes the relevant outputs accessible. To understand what a planner can establish from such an interface, it is first necessary to describe the states, actions, and requirements over which planning operates.

2.2 Symbolic State, Planning, and Task Specifications

To use knowledge about a task in planning, the robot needs a way to reason about what its actions would change. For the book-transfer task, the relevant connection is that grasping the book makes it available to be carried, and carrying it to the sofa makes the requested placement possible. Symbolic planning captures such relationships while leaving much of the detail of their physical realization to perception and control. The resulting abstraction allows the planner to explore how familiar actions could accomplish a new goal. Its construction begins with the conditions that distinguish one task-relevant state from another.

2.2.1 Predicates and Action Models

Let ξ_t denote the physical state of the robot and its environment at time t . It may include robot configuration, object poses, contact conditions, and quantities that are not directly observed. A grounded predicate is a Boolean function of this state. For example, $\text{on}(b, T)$ describes whether book b is on table T , and $\text{holding}(b)$ describes whether the robot holds it. The predicate name and its arguments define the condition of interest; an evaluator connects that condition to geometry, sensor data, or another state estimate.

Given a finite set $\mathcal{P}$ of grounded predicates, a symbolic state can be represented by the subset that is true:

$$s_t = \{p \in \mathcal{P} : p(\xi_t) = 1\}. \quad (2.2)$$

This expression defines an abstraction of physical state, not a claim that the robot observes ξ_t exactly. In practice, the planner receives an estimate $\hat{s}_t$ constructed from available evidence. Two physically different scenes can have the same symbolic description: the book may lie at several positions on

the table while $\text{on}(b, T)$ remains true. Such abstraction is useful when the omitted differences do not invalidate the reasoning being performed (Konidaris et al., 2018).

In a basic STRIPS-style representation, an action operator A has a set of preconditions $\text{Pre}(A)$, an add set $\text{Add}(A)$, and a delete set $\text{Del}(A)$ (Fikes and Nilsson, 1971). For the positive-precondition case, the operator is applicable when $\text{Pre}(A) \subseteq s$ and predicts the successor

$$\Gamma(s, A) = (s \setminus \text{Del}(A)) \cup \text{Add}(A). \quad (2.3)$$

Conditions outside the add and delete sets retain their values. An illustrative pick operator might require $\text{on}(b, T)$, $\text{near}(b)$, and handempty ; it adds $\text{holding}(b)$ and deletes $\text{on}(b, T)$ and handempty . This model describes the changes needed for task reasoning. It leaves the choice of grasp and the detailed motion to the mechanisms that realize the operator.

A planning problem supplies an initial state s_0 and a goal condition g in addition to the action model. A sequential plan $(A_1, \dots, A_n)$ is valid within this model if each operator is applicable in the state reached by its predecessors and the final state satisfies g . Planning Domain Definition Language (PDDL) separates reusable domain descriptions from problem instances containing objects, initial conditions, and goals (Ghallab et al., 1998). This separation supports reuse: changing the requested arrangement need not change the operators that describe picking and placing.

The usual set-of-true-facts representation adopts a closed-world interpretation: a grounded fact absent from s is false. That convention is convenient for search, but it must be distinguished from incomplete perception. Failure to detect a book is not evidence that no book exists. Similarly, a fixed semantic map is not inherently closed-world; the assumption arises when its contents are treated as complete for the relevant reasoning. The symbolic domain must be understood together with the procedure that supplies its facts.

2.2.2 Requirements over Time

A final-state goal specifies an outcome but does not describe every requirement on the way to it. The robot might be required to visit one location before another, avoid a region while carrying an object, or repeat an activity indefinitely. Temporal Logics, such as Linear Temporal Logic (LTL), express such requirements over sequences of propositional valuations (Pnueli, 1977; Emerson, 1990).

Let AP be a set of atomic propositions and let $\sigma = \sigma_0 \sigma_1 \dots$ be an infinite trace, with $\sigma_t \subseteq \text{AP}$ containing the propositions true at step t . An atomic proposition may be a grounded state predicate or an event associated with a skill, provided its meaning is specified. LTL combines Boolean connectives with temporal operators. The expression $\mathbf{X}\varphi$ requires φ at the next step; $\mathbf{F}\varphi$ requires it eventually; and $\mathbf{G}\varphi$ requires it at every step. The strong-until expression $\varphi \mathbf{U} \psi$ requires ψ at some step and φ at every preceding step. For example,

$$\begin{aligned} \sigma, t \models \varphi \mathbf{U} \psi &\iff \exists j \geq t : \sigma, j \models \psi \text{ and} \\ &\forall i \in \{t, \dots, j-1\}, \sigma, i \models \varphi. \end{aligned} \quad (2.4)$$

Here $\models$ denotes satisfaction. The operators describe logical ordering over trace positions; they do not by themselves specify durations in seconds.

For example, let k mean that the robot is at the kitchen and f that it is at the fridge. The formula

$$\varphi = \mathbf{F}(k \wedge \mathbf{F}f) \quad (2.5)$$

requires a kitchen visit with a fridge visit at that time or later. It permits earlier fridge visits and permits simultaneous satisfaction if both propositions can hold at once. Replacing the inner $\mathbf{F}f$ with $\mathbf{X}\mathbf{F}f$ requires a later fridge visit, but still does not forbid an earlier one. These differences illustrate why fluent language and a plausible subgoal list are insufficient to establish that a translation expresses the intended constraint.

2.2.3 Automata, Synthesis, and the Scope of Verification

Temporal specifications can be converted into automata that track progress as propositions become true (Vardi, 1996). For a co-safe specification, satisfaction is witnessed by a finite *good prefix*: once that prefix has occurred, every continuation satisfies the specification (Kupferman and Vardi, 2001). Such requirements can be represented by a finite automaton accepting their good prefixes. This differs from claiming that every LTL requirement is decidable from a completed finite task. For example, observing a safe run so far cannot establish $\mathbf{G}\neg b$ for an unbounded future. An avoidance requirement limited to reaching a goal can instead be written $(\neg b) \mathbf{U} (g \wedge \neg b)$, which requires eventual achievement of g and excludes b through that achievement.




Figure 2.1: An automaton for $\mathbf{F}(k \wedge \mathbf{F}f)$. State 2 awaits a kitchen visit; state 1 records that the kitchen requirement has been met; state 0 accepts after a qualifying fridge visit. The direct transition from 2 to 0 permits simultaneous satisfaction. The self-loop at 2 also permits a fridge visit before the kitchen, so the formula does not prohibit such a visit. Edge labels describe proposition valuations, not physical actions.

A deterministic automaton can be written as $\mathcal{A}_\varphi = (Q, q_0, 2^{\text{AP}}, \delta, F)$, where Q is its state set, q_0 its initial state, δ its transition function, and F its accepting states. After reading a valuation σ_t , its progress state changes according to

$$q_{t+1} = \delta(q_t, \sigma_t). \quad (2.6)$$

Figure 2.1 shows how this state records which parts of a requirement have already been satisfied. The automaton does not itself identify the fridge, compute a path, or cause the robot to move. Those operations supply and act on the propositions that drive its transitions.

Verification and synthesis use the specification in different ways. Model checking asks whether a supplied model satisfies a property; synthesis constructs behavior to satisfy the property under a model and stated assumptions (Vardi, 1996; Kress-Gazit et al., 2009). Combining a task model with an automaton allows a planner to reason about both the environment state and progress through the specification. In robot applications, the resulting logical plan must also admit a physically feasible realization. Thus a correct-by-construction plan is correct relative to its specification, grounding, and action model. Whether that specification expresses the user’s intent, and whether execution follows the modeled transitions, are additional questions.

2.3 Task and Motion Planning with Reusable Skills

A symbolic plan can establish that picking precedes placing without determining how either action is carried out. Physical execution introduces choices that affect the validity of the sequence: where the robot stands, which grasp it uses, how it approaches the object, and whether the object can pass through the available space. Task and motion planning (TAMP) treats the discrete action sequence and these continuous choices as connected parts of one problem (Garrett et al., 2021).

2.3.1 Geometric Feasibility

A robot configuration $q \in \mathcal{C}$ specifies the variables needed to determine its geometry, such as joint angles and mobile-base pose. For a fixed environment and manipulation mode, the collision-free subset $\mathcal{C}_{\text{free}}$ contains configurations that satisfy the chosen collision constraints. A geometric path $\gamma : [0, 1] \rightarrow \mathcal{C}_{\text{free}}$ connects a start configuration to a goal configuration or goal set. The path parameter specifies progression along the curve; a time-parameterized trajectory additionally specifies when configurations are reached. Kinematic feasibility, collision avoidance, and dynamic feasibility impose different restrictions on such a trajectory.

Manipulation changes the geometry relevant to those restrictions. An object attached to the gripper becomes part of the moving assembly, and a release changes its pose in the scene. A task-level choice of grasp can therefore affect the feasibility of later transport and placement. Methods such as PDDLStream connect symbolic search to procedures that generate and test continuous candidates, including grasps, poses, and motions (Garrett et al., 2020). The planner may need to revisit an earlier discrete or continuous choice when a later motion is infeasible.

Task requirements can also create geometric constraints beyond ordinary collision avoidance. An instruction to avoid a particular object or region changes which paths are acceptable even when those paths are physically traversable. Once a referent is localized, a proposition such as $\text{near}(o)$ can label a region of space. For a selected automaton transition, positions can then be distinguished according to whether they advance the task, violate its constraints, or leave its progress unchanged. This is the

role of the Task Progression Semantic Maps developed in Chapter 3. They connect logical progress to geometric search; they are not models of physical dynamics.

2.3.2 Skills as Temporally Extended Actions

The planner need not select every motor command individually. A skill can encapsulate a closed-loop behavior that runs over several control steps, such as approaching, grasping, or opening a door. The options framework formalizes this temporal abstraction through

$$O = (I_O, \pi_O, \beta_O), \quad (2.7)$$

where I_O is the initiation set, π_O the policy used during execution, and β_O the termination rule (Sutton et al., 1999). In the stochastic formulation, $\beta_O(\xi)$ is the probability of terminating in state ξ . The policy may itself use an observation or history rather than the full physical state. A deterministic stopping condition is a special case of this formulation.

An option’s duration can vary with its starting state and execution outcome. At the level of skill selection, a transition model may therefore describe both the terminal state and elapsed duration, $p(\xi', d \mid \xi, O)$. This is the source of the semi-Markov formulation associated with options: decisions occur after temporally extended behaviors, rather than necessarily after a single motor step. The options framework can support reinforcement learning, but the abstraction does not require that every policy be learned through reward optimization. An analytic controller, a motion planner, or a policy learned from demonstrations can each implement a skill.

An initiation set and a termination rule are not interchangeable with symbolic preconditions and successful effects. A symbolic precondition summarizes when an operator is applicable; its physical realization requires the corresponding skill to be able to run from the actual state. Likewise, a skill may stop because it succeeded, timed out, or encountered a failure. Termination alone cannot establish its intended effects. Composition requires a connection between the states one skill can produce and the states from which the next can operate (Burrige et al., 1999).

This connection explains why state abstraction and skill design belong together. The useful symbolic distinctions depend on the behaviors available to the robot and the questions the planner must answer about them (Konidaris et al., 2018). A representation that records only whether a book is on a table may suffice for one task model but omit the accessibility information needed by a particular grasping skill. The planner’s symbols must therefore be grounded at a level that supports the intended composition, rather than chosen solely because they are convenient labels.

2.4 Semantic Mapping and Spatial Grounding

The ability to reuse a skill depends on recognizing where it can be applied. A model of picking up a book can serve tasks in different rooms, but each use requires the robot to identify a particular book, locate it, and relate its geometry to the motion being considered. The same abstraction that makes a planning model reusable must therefore be connected to the details of the present scene. Semantic

mapping and grounding establish this connection, allowing task reasoning to draw on what the robot perceives.

2.4.1 From Images to a Metric Scene

Observations from different viewpoints must first be brought into a common spatial frame so that detections and planned motions refer to the same environment. Simultaneous localization and mapping (SLAM) addresses this problem by estimating the robot’s state while constructing a representation of the environment from its observations (Cadena et al., 2016). A metric map provides spatial relationships in that common coordinate system. A semantic map adds associations with categories, instances, attributes, or relations. These are complementary forms of information: a collision map may support navigation without identifying a book, while a text label alone does not provide the geometry needed to reach it.

Calibrated RGB-D observations connect visual detections to metric structure. Let K be a pinhole camera’s intrinsic matrix, (u, v) a pixel, and $D(u, v)$ its depth measured along the camera’s optical axis. For an undistorted image, backprojection gives the camera-frame point

$$\mathbf{p}_C = D(u, v)K^{-1}[u, v, 1]^\top. \quad (2.8)$$

If $(R_{WC}, \mathbf{t}_{WC})$ maps camera coordinates into the world frame, then

$$\mathbf{p}_W = R_{WC}\mathbf{p}_C + \mathbf{t}_{WC}. \quad (2.9)$$

The calibration and depth convention determine the meaning of these equations (Hartley and Zisserman, 2004). Applying them to pixels inside a segmentation mask produces candidate three-dimensional support for an object. Combining observations from several viewpoints can improve coverage, but also requires deciding which detections refer to the same instance.

An object record can contain a persistent identifier, pose, spatial extent, appearance features, and supporting observations. Scene graphs additionally represent relations among such entities. Hydra and ConceptGraphs illustrate ways to organize metric and semantic information into structures that support higher-level reasoning (Hughes et al., 2022; Gu et al., 2024). For this thesis, the useful distinction is between storing the environment’s geometry and deciding which parts of that geometry matter to an instruction. Task-relevant semantics can be constructed when needed from retained observations or a suitable scene representation, rather than exhaustively enumerated at mapping time.

2.4.2 Resolving Objects, Relations, and Identity

Referent resolution must distinguish an object’s description from its identity. “The red book” is a description that may match several objects; “the red book on the table” adds a relation that can restrict the candidates. An instruction can nest these relations, as in a book on a table beside a sofa. Resolving it requires evaluating candidate combinations, not simply detecting each noun

independently. The result should bind a task expression to an object instance that can be used consistently by the planner.

Spatial predicates also depend on their reference frame and geometric interpretation. “Left of” can refer to the observer’s viewpoint or an object’s intrinsic orientation, whereas “above” may be defined relative to gravity. Different languages and contexts make different use of spatial frames (Majid et al., 2004). A numerical comparison between object centers is meaningful only after the relevant frame has been fixed. Other relations, such as support or containment, may require shape, contact, or volume information rather than a comparison of centers alone.

The same distinction matters over time. If the robot moves a book, its description may change while its identity persists. Repeated detection provides new candidates but does not automatically establish their correspondence to earlier objects. Similarly, a semantic identification does not determine which geometric elements should move with the object. An instance record and its geometric support answer different questions: which object the task concerns, and which part of the scene represents that object. Maintaining both becomes necessary when the representation must account for the robot’s interventions.

2.5 World Models and State Change

Grounding makes a representation useful for action, but action can change the very conditions on which that grounding depends. Once the robot has moved the book, later instructions and plans must refer to the scene that results. Continued autonomy therefore requires a way to carry knowledge of the environment forward, preserving what remains useful while accounting for what has changed.

A world model supports this reasoning by describing how aspects of the scene evolve. For manipulation, this thesis adopts a functional view: the model should be grounded in the environment, predict task-relevant changes, and account for the actions that cause them. The appropriate mechanism depends on what the robot needs to predict. An analytic transformation, a learned transition distribution, or a combination of the two can supply the required information without modeling every physical process.

Let W_t be an internal representation at a task-relevant decision time, and let A_t describe an intervention. A deterministic transition model produces

$$\widehat{W}_{t+1} = F(W_t, A_t), \quad (2.10)$$

while a probabilistic model may represent $p(W_{t+1} \mid W_t, A_t)$. Both differ from a renderer, which predicts an image from a scene and camera pose without necessarily representing any change to the scene. Predicting a new view of a fixed object is not the same operation as predicting the consequences of moving it.

The appropriate prediction depends on the decision being supported. A task planner may need only whether a door becomes open; a motion planner needs the resulting geometry; a controller may require velocities and contact forces. Dynamic Scene Representation, for example, studies object permanence and continuity in action-conditioned manipulation prediction (Xu et al., 2021). Such

work makes explicit that a useful state must retain information about objects even when they are temporarily hidden. The level of detail can vary, but it must support the downstream question.

For a rigid object, a predicted change of pose can preserve its shape while relocating it in the scene. If the prediction comes from the endpoint of a manipulation plan, its meaning is conditional: this is where the object should be if the planned action succeeds. A failed grasp or slip can produce a different physical outcome. Observation-based correction therefore plays a different role from prediction. In a general state-estimation loop, new evidence o_{t+1} can revise the predicted representation through an update $W_{t+1} = U(\widehat{W}_{t+1}, o_{t+1})$. The prediction provides an expectation; the observation constrains whether that expectation describes the environment.

Not every image used to update a scene is an independent observation of the world. A rendered image depicts the representation, and a generated image supplies an inferred appearance consistent with its conditioning. Such images can support prediction or reconstruction, but agreement with them measures a different property from agreement with a physical sensor observation. Preserving the source of the evidence is necessary to interpret what an updated model establishes.

Continued planning also requires agreement among the representations derived from the world state. If an object moves, its semantic location, collision geometry, and rendered appearance should describe the same revision of the scene. A correct update to one representation cannot compensate for another component continuing to use an old one. Chapter 4 develops this problem using an explicit, renderable spatial representation. The properties of that representation explain both how local updates are possible and why visual reconstruction alone does not complete a world model.

2.6 Gaussian Scene Representations

For a world model to remain useful after manipulation, its scene representation must support the changes the robot needs to predict. When an object moves, the robot needs a way to change its represented position while retaining the surrounding scene. When a person inspects a proposed action, the representation should make the relevant geometry and spatial relationships visible. Local editing and rendered inspection place different demands on the information stored about the world.

A point cloud records samples of three-dimensional geometry; a mesh connects surface elements; and an occupancy representation describes whether regions are occupied. A radiance representation also models appearance as a function of viewpoint, allowing images to be synthesized from camera poses other than those used during capture. Combining this ability to render new views with an explicit spatial structure offers a basis for inspecting a scene and updating the parts affected by manipulation.

2.6.1 Gaussian Primitives and Image Formation

Three-dimensional Gaussian Splatting represents scene appearance using explicit primitives with position, shape, opacity, and color parameters (Kerbl et al., 2023). A primitive centered at $\mu_i \in \mathbb{R}^3$ has

spatial profile

$$G_i(\mathbf{p}) = \exp\left[-\frac{1}{2}(\mathbf{p} - \boldsymbol{\mu}_i)^\top \Sigma_i^{-1}(\mathbf{p} - \boldsymbol{\mu}_i)\right], \quad (2.11)$$

where the positive-definite covariance Σ_i controls its extent and orientation. These roles can be separated through

$$\Sigma_i = R_i \text{diag}(s_i^2) R_i^\top, \quad (2.12)$$

where the vector s_i contains three positive scales, squared componentwise, and R_i is a rotation matrix.

To render an image, the primitives are projected into the camera and their contributions are combined in visibility order. For a pixel $\mathbf{u}$, a simplified front-to-back compositing expression is

$$C(\mathbf{u}) = \sum_i T_i(\mathbf{u}) \alpha_i(\mathbf{u}) c_i(\mathbf{d}), \quad T_i(\mathbf{u}) = \prod_{j < i} (1 - \alpha_j(\mathbf{u})), \quad (2.13)$$

where $\alpha_i(\mathbf{u})$ combines opacity with the projected footprint, T_i is the remaining transmittance, and $c_i(\mathbf{d})$ is color for viewing direction $\mathbf{d}$. Spherical-harmonic coefficients provide one way to represent this directional appearance. The parameters can be optimized by comparing rendered and captured images, with the original method also adjusting the density of primitives during reconstruction (Kerbl et al., 2023).

The resulting representation is explicit in space, but its elements need not align with object boundaries. A primitive may contribute to several pixels, and a collection that looks like one object from a particular view may include support belonging to nearby surfaces. Good image reconstruction and correct object ownership are consequently different requirements. This distinction matters for manipulation: moving every primitive in a broad image crop can inadvertently move background content along with the target object.

2.6.2 Object Motion, Local Updates, and Inspection

Once the primitives belonging to an object have been identified, their explicit geometry makes it possible to move them together. For a rigid change of coordinates $(R, \mathbf{t})$, a Gaussian’s geometry transforms as

$$\boldsymbol{\mu}'_i = R\boldsymbol{\mu}_i + \mathbf{t}, \quad \Sigma'_i = R\Sigma_i R^\top. \quad (2.14)$$

Applying the same transformation to an object’s correctly identified primitives moves its represented geometry while preserving its internal arrangement. Other scene elements can remain fixed. When appearance is directional, its representation must also be handled consistently with the chosen rotation and lighting assumptions. The mathematical convenience of a rigid transformation does not establish that the physical object underwent that transformation; the distinction between a planned and an observed pose still applies.

Object motion also exposes parts of the scene that may not have been visible before. Relocating the book reveals the surface beneath it, and a new viewpoint may reveal an unobserved side of the book itself. Geometry alone cannot recover all of this appearance. Local reconstruction methods such as CL-Splats update changed regions while retaining the rest of an existing Gaussian scene (Ackermann et al., 2025). This provides a basis for maintaining scene continuity without fitting the entire environment

after each change. Which regions may change, which quantities remain constrained by the action model, and which observations support the update are additional design choices.

Rendered inspection can expose consequences of these choices. A user can view a proposed path relative to objects and obstacles, or examine whether an updated scene places an object where expected. The image makes a representation accessible, but its realism does not certify collision freedom, object identity, or physical correctness. A Gaussian scene therefore becomes useful as a world model through its transition and update mechanisms and through the geometric and semantic abstractions it supplies to later decisions.

2.7 Learning Visuomotor Skills from Demonstrations

A current world representation allows the planner to reconsider what can be done as the scene changes. The range of behaviors it can construct, however, still depends on the skills available to it. Extending that range requires a way to acquire control policies whose detailed movements would be difficult to engineer for every object and situation. Learning from demonstration approaches this problem through examples of successful behavior.

For manipulation, those examples can contain images, robot state, and the commands issued during execution. The policy learns how observations relate to actions without requiring the designer to specify every movement. The quality and coverage of the demonstration data, the observation representation, and the policy architecture all affect the resulting behavior (Mandlekar et al., 2022).

2.7.1 Behavior Cloning and Closed-Loop Execution

Let $\mathcal{D}$ contain demonstrated observation–action pairs (o_t, a_t) . Behavior cloning trains a conditional policy to reproduce those actions; a probabilistic formulation minimizes

$$\mathcal{L}_{\text{BC}}(\theta) = -\mathbb{E}_{(o_t, a_t) \sim \mathcal{D}} [\log \pi_{\theta}(a_t \mid o_t)]. \quad (2.15)$$

The conditioning input may include a recent observation history and a task description. At deployment, a closed-loop policy receives new observations as the robot acts. It therefore responds to the scene rather than replaying a stored sequence of commands. Its responsiveness nevertheless depends on whether its inputs distinguish the states that require different actions.

Supervised action prediction is evaluated on a distribution that execution can itself change. An early error may bring the robot to a state that the demonstrator rarely visited, where another error is more likely. This sequential distribution shift motivates approaches that collect corrective behavior on states visited by the learned policy (Ross et al., 2011). Feedback helps the robot respond to deviations, but does not ensure that every unfamiliar situation is covered by its demonstrations. Task-level recovery may also require replacing the current skill, rather than adjusting its next movement.

Several different action sequences may be appropriate for the same observation. For example, an object can be approached from either side, yet averaging the two approaches may produce an undesirable path between them. Generative policies address this multimodality by learning a distribution

over action trajectories. Let

$$\mathbf{x} = (a_t, \dots, a_{t+H-1}) \in \mathbb{R}^{H \times d_a} \quad (2.16)$$

be a length- H action chunk, with d_a control dimensions per step. A conditional generator samples $\mathbf{x} \sim p_\theta(\mathbf{x} \mid o_t)$. The robot can execute a short prefix, acquire another observation, and generate a new chunk. Diffusion Policy uses this receding-horizon organization to combine coordinated action generation with visual feedback (Chi et al., 2025).

2.7.2 Denoising Diffusion

Sampling a whole action trajectory requires a model that can preserve the relationships among its movements while accommodating several possible ways to act. Denoising diffusion learns such a distribution by learning to remove noise from corrupted training examples (Ho et al., 2020). For a clean action trajectory $\mathbf{x}^0$, the forward corruption process can be sampled directly at noise level k as

$$\mathbf{x}^k = \sqrt{\bar{\alpha}_k} \mathbf{x}^0 + \sqrt{1 - \bar{\alpha}_k} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (2.17)$$

where $\bar{\alpha}_k = \prod_{j=1}^k \alpha_j$ and the schedule $\alpha_j \in (0, 1)$ determines how much signal is retained. A conditional noise-prediction network is trained with the objective

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E} \left[\left\| \epsilon_\theta(o_t, \mathbf{x}^k, k) - \epsilon \right\|_2^2 \right]. \quad (2.18)$$

The expectation samples demonstrated trajectories, noise levels, and Gaussian noise. At inference, a noisy sample is iteratively refined into an action trajectory. Conditioning on the current observation guides the generated behavior without requiring the model to generate future images (Chi et al., 2025).

The diffusion index k and the physical action index t have different meanings. One trajectory contains actions for several future control steps; denoising repeatedly refines that entire trajectory before its selected prefix is executed. Neither a denoising step nor the reduction of noise in a sample is an event in the physical task. This distinction is essential when interpreting diagrams of generated actions and predicate trajectories in Chapter 5.

2.7.3 Conditional Flow Matching

The trajectory-generation problem can also be approached by learning how to move samples continuously from noise toward demonstrated behavior. Flow matching learns a vector field that transports samples between these distributions (Lipman et al., 2023; Tong et al., 2024). To match the convention used in Chapter 5, let $\mathbf{x}^0$ be a data trajectory, $\mathbf{x}^1 \sim \mathcal{N}(0, I)$ a noise sample, and define a linear interpolation

$$\mathbf{x}_\tau = (1 - \tau)\mathbf{x}^0 + \tau\mathbf{x}^1, \quad \tau \in [0, 1]. \quad (2.19)$$

For this sampled pair, the target velocity is $\mathbf{x}^1 - \mathbf{x}^0$. A conditional vector field v_θ can be trained by

$$\mathcal{L}_{\text{CFM}} = \mathbb{E} \left[\left\| v_\theta(o_t, \mathbf{x}_\tau, \tau) - (\mathbf{x}^1 - \mathbf{x}^0) \right\|_2^2 \right]. \quad (2.20)$$

The expectation samples trajectory–noise pairs and interpolation values τ uniformly from $[0, 1]$. Generation integrates the learned field from noise at $\tau = 1$ toward data at $\tau = 0$. The interpolation parameter is again a sampling coordinate, not physical time. Independent data–noise pairing defines one conditional flow-matching construction; optimal-transport variants change how the pairs are selected. The straight-line form alone does not establish an optimal transport coupling between the complete distributions.

Diffusion and flow matching are alternative ways to train and sample a conditional trajectory model. They do not specify what information the trajectory must contain. A temporal convolutional network can model local structure along the horizon, while a Transformer can model interactions among time-step representations through attention. Either architecture can condition on an observation encoding and a noise or flow parameter. The modeling target, training objective, and network architecture are therefore separate choices. For compositional autonomy, the central question is whether the generated behavior also supplies the state information needed to reason about its use.

2.8 Symbolic Interfaces for Composition and Monitoring

The value of an acquired skill extends beyond the movements it has learned to produce. To use a grasping skill as part of a longer activity, the planner needs to reason about whether the object will be held and whether subsequent transport or placement is appropriate. These are the conditions represented by the action models in Section 2.2. Connecting the learned policy to those conditions gives the planner a basis for composing the skill with others and assessing progress through the task. The remaining question is how the skill’s state reports relate to the behavior it produces.

2.8.1 State Estimates and Predicted Outcomes

A predicate evaluator can estimate current symbolic conditions from an observation, for example through scores $b_j(o_t) \in [0, 1]$ for predicates $j = 1, \dots, J$. A threshold converts those scores into a discrete valuation for a classical planner. This conversion is an interface decision: it does not make an uncertain estimate certain. Nor is a score between zero and one automatically a calibrated probability. Calibration concerns whether stated confidence agrees with empirical correctness frequencies, a property that must be assessed separately from classification accuracy (Guo et al., 2017).

A predicted predicate trajectory answers a different question from a current-state evaluator. For an action chunk $\mathbf{x}$, let

$$\mathbf{z} = (\mathbf{z}_{t+1}, \dots, \mathbf{z}_{t+H}) \in [0, 1]^{H \times J} \quad (2.21)$$

describe the conditions expected at the corresponding future steps. An observation-conditioned classifier and a trajectory predictor may use similar inputs, but their targets refer to different times. The vector $\mathbf{z}_{t+1}$ sampled before an action is taken is a prediction of its next state; it is not an independent observation confirming success after execution. A later rollout can incorporate a new observation without changing the provenance of the earlier prediction.

There is also a distinction between predicting action and outcome marginals and predicting their relationship. Independent samples from $p(\mathbf{x} \mid o_t)$ and $p(\mathbf{z} \mid o_t)$ may select incompatible alternatives when different action modes lead to different outcomes. A joint model $p(\mathbf{x}, \mathbf{z} \mid o_t)$ can represent this dependence. Sharing an observation encoder does not by itself require two independently sampled outputs to agree. Chapter 5 studies joint generation as a way to acquire this symbolic interface together with the control policy.

The distinction also determines what must be evaluated. Predicate accuracy compares a prediction with labeled conditions; action–predicate coherence compares predicted conditions with those produced by executing the associated actions. Accurate marginals do not establish that correspondence, and joint generation is a modeling choice rather than a proof of physical agreement. Evaluating paired rollouts test the relationship on which the planner relies.

2.8.2 Composition, Progress, and Recovery

With state estimates and action models available, a planner can select a skill, monitor its intended effects, and plan again from updated conditions. This differs from replaying a fixed order of skills: the next choice depends on the state presented to the planner. If a door closes, the original sequence may no longer be applicable even when the current controller continues to produce valid motor commands. Recovery requires recognizing the changed condition and selecting behavior that can restore progress.

The policy and planner can operate at different timescales. The policy uses frequent observations to adjust movement; the planner reasons over task-relevant conditions to decide which skill should run. This separation permits reuse across goals, provided the skill’s initiation conditions, effects, and state reports remain meaningful in the new composition. It does not remove the need for a suitable symbolic vocabulary or an action model. Learning a controller, learning a predicate evaluator, discovering useful predicates, and learning operator preconditions and effects are distinct problems.

Similarly, zero-shot composition concerns what is new at the level of the requested behavior. The constituent skills can be trained from demonstrations while their requested combination has no end-to-end demonstration. This is different from performing a skill that the robot has never learned. The distinction connects the two sources of generality introduced earlier: learning broadens the available behavior, while planning can arrange that behavior into additional tasks.

Learned knowledge and behavior become useful to planning when the robot can relate them to explicit requirements, conditions, and changes in the world. The same connections make it possible to examine whether the planner’s reasoning applies to the task and environment at hand. Constructing these connections is therefore central to advancing capability and assurance together.

The first such connection is needed before a plan can be constructed. A person’s instruction must become a task that the planner can reason about, with its referents identified and its temporal and spatial requirements made explicit. Chapter 3 develops that connection without assuming a complete semantic map in advance. Its language-to-plan framework then provides the basis for examining how the world representation and the skill library can themselves be extended in Chapters 4 and 5.

Chapter 3

Grounding Language in the Real World

This chapter is adapted from *Verifiably Following Complex Robot Instructions with Foundation Models*, coauthored by Benedict Quartey, Eric Rosen, Stefanie Tellex, and George Konidaris (Quartey et al., 2025). The authorial “we” refers to this collaborative research.

The usefulness of a general-purpose robot depends in part on the range of tasks a person can ask it to perform. An instruction may identify an object through its relationship to another, require several actions in a particular order, or prohibit behavior that would otherwise be feasible. The preceding background chapter established how symbolic planning can express and reason about such requirements. The difficulty is to obtain the right requirements from the person’s words and connect them to the environment in which the robot must act. A plan can satisfy a specification while pursuing the wrong object or the wrong interpretation of the task.

This chapter develops the first connection in the thesis argument: *grounded language understanding* at the interface of intent. LIMP brings the breadth of foundation models into a language-to-plan framework in which task requirements and their spatial grounding can be inspected. The instruction determines which semantic information the robot constructs from its observations, allowing the planner to work beyond a prespecified set of named landmarks. This establishes a basis for expressive instruction following while retaining two assumptions: the mapped environment remains static, and suitable navigation and manipulation skills are already available. Chapters 4 and 5 take up these assumptions in turn.

3.1 From an Instruction to a Grounded Plan

The instruction in Figure 3.1 illustrates why language understanding and spatial grounding must be developed together: “*Bring the green plush toy to the whiteboard in front of it, watch out for the robot in front of the toy.*” Recognizing a plush toy or a whiteboard is insufficient. The robot must




Figure 3.1: LIMP executing the instruction “Bring the green plush toy to the whiteboard in front of it, watch out for the robot in front of the toy”. The robot dynamically detects and grounds open-vocabulary referents with spatial constraints to construct an instruction-specific semantic map, then synthesizes a task and motion plan to solve the task. In this example, the robot navigates from its start location (yellow, A), to the green plush toy (green, B), executes a pick skill then navigates to the whiteboard (blue, C), and executes a place skill. Note that the robot has no prior semantic knowledge of the environment.

distinguish the intended instances through their spatial relationships, determine the order of navigation and manipulation, and interpret the other robot as a constraint on the route. The meaning of the request therefore depends on both the structure of the instruction and the arrangement of the scene.

Language-guided navigation systems connect descriptions to object goals and landmarks (Gadre et al., 2023; Shah et al., 2023), while mobile-manipulation systems extend this connection to picking and placing (Yenamandra et al., 2023; Chen et al., 2023a; Liu et al., 2024b). LIMP addresses the additional problem of resolving descriptive spatial references together with temporal constraints. Directly generating a sequence of subgoals or a control program does not, by itself, establish that these constraints are respected (Liang et al., 2023). Translating the request into temporal logic makes the requirements available to a planner, but the symbols must still refer to the right things. Earlier language-to-logic systems ground those symbols in prebuilt semantic maps or landmark descriptions (Liu et al., 2023; Berg et al., 2020). LIMP instead constructs the required semantics from the instruction and the robot’s retained observations.

Language Instruction grounding for Motion Planning (LIMP) combines language translation, spatial grounding, and task and motion planning. An LLM translates the instruction into a temporal-logic specification with composable descriptions of its referents. Vision-language models and spatial reasoning locate those referents in a metric map, and the planner uses the resulting specification and semantic information to construct navigation and manipulation behavior. The three contributions are a modular language-to-plan framework, a spatial grounding method for open-vocabulary objects with

relational descriptions, and a planning procedure that turns task-progress requirements into geometric goal and avoidance regions. Together, they make the interpretation and planned realization of an instruction available for inspection before execution.

3.2 Related Work

NLMap (Chen et al., 2023a) uses pretrained vision-language models to make a scene representation queryable through open-vocabulary descriptions. Object availability and location then ground an LLM-based planner. LIMP connects such grounding to an explicit representation of spatial relations and temporal constraints, allowing these requirements to guide task and motion planning.

Lang2LTL (Liu et al., 2023) uses entity extraction and replacement to translate language instructions into temporal logic. Lang2LTL-2 (Liu et al., 2024a) extends this approach with visual and textual landmark descriptions and spatial-relation reasoning. Both ground symbols using a supplied semantic map or landmark database. LIMP instead constructs instruction-specific semantic information from the robot’s retained observations.

Action-Oriented Semantic Maps (AOSMs) (Rosen et al., 2023) construct navigation abstractions that connect portable manipulation skills in new environments. Their representation is organized around the supplied manipulation skills and their symbolic abstractions. LIMP develops a complementary connection: an instruction determines which open-vocabulary referents and spatial relationships must be grounded for planning.

3.2.1 Temporal Specifications for Instruction Following

The temporal-logic semantics and automata construction in Section 2.2.2 provide the basis for LIMP’s task representation. The choice of LTL allows ordering, eventual achievement, and avoidance requirements to be expressed together (Pnueli, 1977; Menghi et al., 2021). Recent approaches translate natural language into such specifications (Liu et al., 2023; Berg et al., 2020; Pan et al., 2023; Cosler et al., 2023; Fuggitti and Chakraborti, 2023; Chen et al., 2023b); LIMP couples this translation to the resolution of open-vocabulary spatial descriptions.

Temporal task structure can also support learning. Quartey et al. (2023) combine temporal-logic task representations with language-model-derived object embeddings to generate contextually related auxiliary tasks that can be learned from shared experience.

The planner uses the specification to synthesize behavior rather than checking an independently generated plan afterward (Vardi, 1996; Kress-Gazit et al., 2009; Colledanchise et al., 2017). The resulting correctness claim is relative to the translated specification, its grounding, and the modeled actions. Inspection of the translation and the referent map addresses whether these inputs express the user’s intended task; successful physical execution is a further requirement.

3.3 Problem Definition

Given a natural language instruction l , our goal is to synthesize and sequence navigation and manipulation behaviors to produce a policy that satisfies the temporal and spatial constraints in l . Spatial constraints determine task success based on the sequence of robot poses traversed during execution; temporal constraints determine the sequencing of these spatial constraints as a function of task progression. We assume a robot with an RGB-D camera has already navigated a space, capturing images and camera poses. From these observations, the robot builds a metric map m (e.g., point cloud, 3D voxel grid) of the environment, defining the space of possible $SE(3)$ poses P and enabling robot localization (i.e., estimating $p_{\text{robot}} \in P$). Unlike previous work leveraging temporal logic (Liu et al., 2023), we do not assume access to a semantic map with prespecified object locations or predicates. LIMP uses the foundation-model interfaces introduced in Section 2.1: a task-agnostic vision-language model σ for text-conditioned detection or segmentation, and an LLM ψ for instruction translation.

Navigation: Navigation is formalized as an object-goal oriented continuous path planning problem, where the goal is to generate paths to a goal pose set $P_{\text{goals}} \subset P$ while staying in feasible regions ($P_{\text{feasible}} \subset P$) and avoiding infeasible regions ($P_{\text{infeasible}} = P_{\text{feasible}}^C$). Infeasible regions include environment obstacles as well as dynamically determined semantic regions that violate constraints in the instruction l .

Manipulation: Manipulation behaviors use the option formulation from Section 2.3.2, with an object parameter θ . For $O_\theta = (I_\theta, \pi_\theta, \beta_\theta)$, initiation and termination depend on the robot pose and object attributes, and the policy is conditioned on the target object. The initiation set I_θ denotes the global reference frame robot positions and object-centric attributes—such as object size—that determine if the option policy π_θ can be executed on the object θ . To execute a manipulation skill on an object, an object-goal navigation behavior must first be executed to bring the robot into proximity with the object. We assume access to a library of these manipulation skills and demonstrate our approach on multi-object goal navigation and open-vocabulary mobile pick-and-place (Yenamandra et al., 2023; Yokoyama et al., 2024).

3.4 Language Instruction Grounding for Motion Planning

The framework in Figure 3.2 separates instruction translation, referent grounding, and plan synthesis so that each representation can be inspected at the point where it enters the next computation. The following sections describe these mechanisms. Appendix A provides the prompts, spatial comparators, planning algorithms, skill implementations, and evaluation details accompanying the study.

3.4.1 Language Instruction Module

The language instruction module uses an LLM ψ to translate a natural language instruction l into a linear temporal logic specification φ_l with a novel composable syntax for referent disambiguation. Translation proceeds through two stages of in-context learning. The first stage prompts ψ to translate




Figure 3.2: [A] LIMP translates natural language instructions into temporal logic expressions, where open-vocabulary referents are applied to predicates that correspond to robot skills—note the context-aware resolution of the phrase “blue one” to the referent “blue_sofa”. [B] Vision-language models detect referents, while spatial reasoning disambiguates referent instances to generate a 3D semantic map that localizes instruction-specific referents. [C] Finally, the temporal logic expression is compiled into a finite-state automaton, which a task and motion planner uses with dynamically-generated task progression semantic maps to progressively identify goals and constraints in the environment, and generate a plan that satisfies the high-level task specification.

l into a conventional LTL formula ϕ_l where propositions refer to open-vocabulary objects. The second stage takes l and ϕ_l as input and prompts ψ to generate a new formula φ_l with predicate functions corresponding to parameterized robot skills.

We define three predicate functions—**near**, **pick**, and **release**—for the primitive navigation and manipulation skills required for multi-object goal navigation and mobile pick-and-place. Predicate functions in φ_l are parameterized by *Composable Referent Descriptors* (CRDs), our novel propositional expressions representing specific referent instances by chaining comparators that encode descriptive spatial information. For example, the instruction “the yellow cabinet above the fridge that is next to the stove” can be represented with the CRD:

$$\text{yellow_cabinet} :: \text{isabove}(\text{fridge} :: \text{isnextto}(\text{stove})). \quad (3.1)$$

This specifies that there is a fridge *next to* a stove, and the desired yellow cabinet is *above* that fridge. CRDs are constructed from a set of 3D spatial comparators (Jatavallabhula et al., 2023) defined in our prompting strategy.

LIMP accepts varied phrasing and composable spatial descriptions rather than requiring a fixed instruction template (Yenamandra et al., 2023; Liu et al., 2024b). Its ability to resolve a particular instruction nevertheless depends on the translation and grounding models. The LLM ψ directly samples the entire LTL formula φ_l with predicate functions parameterized by CRDs using appropriate spatial comparators based on the instruction’s context. Figure 3.3 illustrates the result of our two-stage prompting strategy.

Inspecting the Language Translation: Errors in referent extraction or temporal structure can change the task presented to the planner. LIMP’s symbol verification node (Figure 3.2) presents the extracted CRDs and the implied task structure to the user, then reprompts the LLM when the user supplies a correction. A deterministic translation of the specification’s automaton into English statements supports this inspection without requiring the user to read the formula directly. Human

Input Instruction :	“ Bring the green plush toy to the whiteboard in front of it, watch out for the robot in front of the toy”
Conventional LTL:	$\mathcal{F}(\text{green_plush_toy} \wedge \mathcal{F}(\text{whiteboard} \wedge \neg \text{robot}))$ ϕ_l
Our LTL Syntax:	$\mathcal{F}(A \wedge \mathcal{F}(B \wedge \mathcal{F}(C \wedge \neg D \wedge \mathcal{F}E)))$ φ_l
A:	near[green_plush_toy]
B:	pick[green_plush_toy]
C:	near[whiteboard::isinfrontof(green_plush_toy)]
D:	near[robot::isinfrontof(green_plush_toy)]
E:	release[green_plush_toy, whiteboard::infrontof(green_plush_toy)]

Figure 3.3: An instruction is first translated into a conventional LTL formula ϕ_l that loosely captures the desired temporal occurrence of referent objects, then into our LTL syntax φ_l with predicate functions that temporally chain required robot skills parameterized by composable referent descriptors.




Figure 3.4: [A] Our spatial grounding module leverages a VLM to detect all referent occurrences from prior observations of the environment. [B] An initial semantic map with all detected referent instances is generated by backprojecting pixels in segmented referent masks unto the 3D map. [C] Each referent’s spatial comparators is resolved with respect to the origin coordinate frame of reference. [D] Failing instances are filtered out to obtain a Referent Semantic Map (RSM) that localizes the exact referent instances described in the instruction.

review of translated constraints also appears in related approaches (Yang et al., 2024); here, it provides an opportunity to correct the interpretation before planning. The evaluation in Tables 3.1 and 3.2 measures the initial translations without this human correction step. Appendix A.6 reports additional results with interactive reprompting.

3.4.2 Spatial Grounding Module

This module detects and localizes specific object instances referenced in an instruction. The spatial grounding module extracts CRDs from the translated formulas and uses the vision-language models OWL-ViT (Minderer et al., 2022) and SAM (Kirillov et al., 2023) to detect and segment all referent occurrences from the robot’s prior observations of the environment. It backprojects pixels in these segmentation masks into the 3D map using the calibrated geometry described in Section 2.4, creating




Figure 3.5: [A] A given instruction translated into our LTL syntax φ_1 can be compiled into an equivalent finite-state automaton that captures the temporal constraints of the task. A path through this automaton is selected with a strategy that incrementally picks the next progression state from the initial state to the accepting state. The robot then executes the manipulation options and navigation behaviors dictated by this high-level task plan. [B] To execute navigation objectives our approach generates a task progression semantic map (TPSM) that augments the environment with state transition constraints, localizing goal (yellow) and avoidance (red) regions. Generated TPSMs are converted into 2D obstacle maps for constraint-aware continuous path planning.

an initial semantic map of the detected instruction referents. From the example in Figure 3.3, occurrences of *green-plush-toy*, *whiteboard*, and *robot* are detected, segmented, and backprojected onto the map (Figure 3.4[a&b]).

To obtain the specific object instances described in the instruction, the module evaluates the 3D spatial comparators in each referent’s CRD—recall that CRDs are propositional expressions and can be evaluated as true or false. We define eight spatial comparators (*isbetween*, *isabove*, *isbelow*, *isleftof*, *isrightof*, *isnextto*, *isinfrontof*, *isbehind*) to reason about spatial relationships based on backprojected 3D positions. Since all backprojected positions are relative to an origin coordinate system, our spatial comparators are resolved from the perspective of this origin position as shown in Figure 3.4[c]. This type of relative frame of reference (FoR) when describing spatial relationships between objects, in contrast to an absolute or intrinsic FoR, is dominant in English (Majid et al., 2004), and is a logical choice for our work.

Using the 3D position of each referent’s center mask pixel as its representative position, the module resolves a referent with a spatial description by applying the appropriate spatial comparator to all detected pairs of the desired referent and comparison landmark objects. This filtering process yields a Referent Semantic Map (RSM) that localizes specific object instances described in the instruction as shown in Figure 3.4[d].

VLM Verification: Object-detector misclassifications are the main source of error in this module. LIMP provides 3D visualizations through which users can inspect referent localization; the reported system does not implement interactive correction of these misclassifications.

3.4.3 Task and Motion Planning Module

Finally, our TAMP module synthesizes and sequences navigation and manipulation behaviors to produce a plan that satisfies the temporal and spatial constraints expressed in the given instruction.

Progressive Motion Planner (PMP): The TAMP module compiles the supported co-safe LTL specification with parameterized robot skills into a finite-state automaton (Figure 3.5[a]) to guide task and motion planning. A path from the initial to the accepting state in this automaton is a high-level task plan that interleaves navigation and manipulation objectives required to satisfy the instruction. The planner selects such a path with a strategy that incrementally selects the next progression state until the accepting state is reached, ensuring the plan obeys all temporal subgoal objectives. As shown in Figure 3.5[a], automaton states are connected by transition edges representing the logical expressions required for transitions. For each transition, our algorithm executes the necessary low-level behaviors: for manipulation subgoals, it executes the appropriate parameterized skill; for navigation subgoals, it dynamically generates Task Progression Semantic Maps (TPSMs) to localize goal and constraint regions and performs continuous path planning using the Fast Marching Tree algorithm (FMT*) (Janson et al., 2015).

Task Progression Semantic Maps (TPSM): A TPSM augments a 3D scene with navigation constraints specified by logical state transition expressions, enabling goal-directed, constraint-aware navigation. Regions of interest in a TPSM are defined using a nearness threshold specifying proximity to an object. This threshold can be set globally or included in the language instruction module’s prompting strategy, allowing an LLM to infer its value based on the instruction. Like our spatial grounding module, TPSMs are agnostic to temporal logic representations and can be used with various planning approaches for semantic constraint-aware motion planning. For the ground mobile robot used in the evaluation, the module converts 3D TPSMs into 2D geometric obstacle maps, where constraint regions are treated as obstacles (Figure 3.5[b]). The spatial representation also permits direct three-dimensional search, although the reported evaluation concerns ground mobile manipulation.

3.5 Evaluation

Our evaluations test the hypothesis that translating natural language instructions into LTL expressions and dynamically generating semantic maps enables robots to accurately interpret and execute instructions in large-scale environments without task-specific training. We focus on three key questions: (1) Can our language instruction module interpret complex, ambiguous instructions? (2) Can our spatial grounding module resolve specific object instances described in instructions? (3) Can our TAMP algorithm generate constraint-satisfying plans?

To answer these questions, we compare LIMP with two baselines: an LLM task planner (NLMap-Saycan (Chen et al., 2023a)) and an LLM code-writing planner (Code-as-Policies (Liang et al., 2023)), representing language-model task sequencing and code generation for robot instruction following. Both baselines use the same LLM (GPT-4-0613), prompting structure, and in-context learning examples as

our language instruction module. In Code-as-Policies, in-context examples are converted into language model-generated program (LMP) snippets (Liang et al., 2023). To ensure competitive performance, we integrate our CRD syntax, spatial grounding module, and low-level robot control into these baselines, allowing them to query object positions, use our FMT* path planner, and execute manipulation skills.

We also evaluate ablations of our two-stage language instruction module due to its importance in instruction following. In our full approach, the first stage prompts an LLM to generate a conventional LTL formula ϕ_l from instruction l by dynamically selecting relevant in-context examples from a standard dataset (Pan et al., 2023) based on cosine similarity. Our first ablation selects in-context examples randomly; and the second ablation skips this stage entirely, directly sampling our LTL syntax with parameterized robot skills φ_l from l .

We conduct a large-scale evaluation across five real-world environments on a diverse task set of 150 instructions from multiple prior works (Chen et al., 2023a; Liu et al., 2024b, 2023). This task set consists of 24 tasks with fine-grained object descriptions (NLMD), 25 tasks with complex language (NLMC), 25 tasks with simple structured phrasing (OKRB), 37 tasks with complex temporal structures (CT) and 39 tasks with descriptive spatial constraints and temporal structures (CST). Below are examples from each task category illustrating the variety in complexity:

```

1 NLMD: Put the brown multigrain chip bag in the woven basket
2 NLMC: I like fruits, can you put something I would like on the yellow sofa for me
3 OKRB: Move the soda can to the box
4 CT: Visit the purple door elevator, then go to the front desk and then go to the kitchen table, in
      addition you can never go to the elevator once you have seen the front desk
5 CST: I have a white cabinet, a green toy, a bookshelf and a red chair around here somewhere. Take
      the second item I mentioned from between the first item and the third. Bring it the cabinet
      but avoid the last item at all costs.

```

To evaluate instruction understanding, we introduce performance metrics: **referent resolution accuracy**, **avoidance constraint resolution accuracy**, and **spatial relationship resolution accuracy**. These metrics utilize the word error rate (WER), widely used in speech recognition to quantify the difference between a reference and a hypothesis transcription by computing the minimal number of substitutions, deletions, and insertions needed to transform the hypothesis into the reference. WER is calculated as $WER = \frac{S+D+I}{N}$, where S is substitutions, D deletions, I insertions, and N is the total number of words in the reference. In our work:

- **Referent resolution accuracy** compares extracted referents in the generated LTL formula to ground truth referents.
- **Avoidance constraint resolution accuracy** compares referents to avoid in the LTL formula (denoted by the unary negation operator) against ground truth avoidance referents.
- **Spatial relationship resolution accuracy** compares generated Composable Referent Descriptors (CRDs) in the LTL specification with ground truth CRD expressions.

We also define **temporal alignment accuracy** and **planning success rate**. A plan is temporally aligned if the sequence of subgoals matches the instructor’s intention, and successful if it satisfies all spatial and temporal constraints specified in the instruction. Achieving a high plan success rate

Table 3.1: Performance comparison of one-shot instruction understanding and spatial resolution.

Approach	Referent Resolution	Avoidance Constraint Resolution	Spatial Relationship Resolution
	Accuracy (Average WER) ↓	Accuracy (Average WER) ↓	Accuracy (Average WER) ↓
NLMap-Saycan	0.09	0.12	0.05
Code-as-Policies	0.22	0.24	0.06
Limp Single Stage Prompting	0.09	0.11	0.05
Limp Two Stage Prompting [Random Embedding]	0.08	0.11	0.03
Limp Two Stage Prompting [Similar Embedding]	0.07	0.04	0.03

Table 3.2: Performance comparison of one-shot temporal alignment and plan success rate.

Approach	Temporal Alignment Accuracy (% of instructions) ↑					Planning Success Rate (% of instructions) ↑				
	NLMD	NLMC	OKRB	CT	CST	NLMD	NLMC	OKRB	CT	CST
NLMap-Saycan	88%	96%	100%	32%	41%	75%	96%	100%	35%	38%
Code-as-Policies	58%	68%	100%	35%	38%	46%	68%	100%	38%	38%
Limp Single Stage Prompting	79%	64%	100%	68%	74%	63%	60%	100%	62%	62%
Limp Two Stage Prompting [Random Embedding]	83%	68%	100%	76%	85%	79%	68%	100%	57%	72%
Limp Two Stage Prompting [Similar Embedding]	88%	80%	100%	76%	92%	79%	76%	100%	65%	79%

is challenging, requiring accurate referent and avoidance constraint resolution, spatial grounding, and temporal alignment. We report average word error rates for each baseline in Table 3.1 and the percentage of successful and temporally accurate plans in Table 3.2.

3.6 Discussion

Beyond the verification benefits of symbolic planning, LIMP outperforms baselines in most task sets, notably in complex temporal planning and constraint avoidance. While NLMap-Saycan and Code-as-Policies effectively generate sequential subgoals, they struggle with strict temporal constraints—for example, avoiding a specific referent while approaching another. LIMP incorporates the instruction’s avoidance regions into geometric path planning, accounting for its stronger performance on the CT and CST task sets. As shown in Table 3.2, LIMP underperforms only against NLMap-Saycan in the NLMC task category. This task set, introduced in the same paper as the baseline (Chen et al., 2023a) (which outperforms LIMP), includes instructions with implicit details such as: “I like fruits, can you put something I would like on the yellow sofa for me.” NLMap-Saycan is better suited to infer and generate plans with possible fruit options, whereas our few-shot LTL translation process is not designed for this.

3.7 From Grounded Intent to Continued Autonomy

LIMP connects expressive instructions to explicit task constraints and the scene geometry needed for planning. Instruction-conditioned semantic maps let the robot work beyond prespecified landmarks, while inspection exposes the translated task and its referents to the user.

This result depends on the task and world models. The planner handles co-safe specifications, whose satisfaction has a finite good-prefix witness (Section 2.2.2). Grounding remains sensitive to perception and spatial interpretation, and plan synthesis does not guarantee physical success or optimality. Suitable manipulation skills and a static mapped environment are still assumed.

These assumptions motivate the next contributions. Successful manipulation can invalidate the representation used to plan it; later tasks must therefore be grounded in a world that accounts for earlier actions. Chapter 4 develops this evolving representation. Chapter 5 then addresses how additional behaviors can be learned together with the symbolic information needed to compose them.

Chapter 4

Evolving World Representations for Planning

This chapter presents joint work by Benedict Quartey and George Konidaris, developed in *REALM: Action-Conditioned Spatial World Models for Long-Horizon Mobile Manipulation*. The authorial “we” refers to this collaboration.

Understanding an instruction establishes what the robot should accomplish; it does not keep the world used to interpret that instruction current. Chapter 3 connected language to task constraints and grounded those constraints in a captured environment. That connection becomes vulnerable as soon as manipulation changes the scene. An object can keep its identity while acquiring a new location, and the next request may depend on both facts. For autonomy to extend across successive tasks, the representation supporting one plan must account for its consequences before it supports another.

This chapter develops *evolving representations of the world*, addressing the world interface of trust. REALM extends the language-to-plan framework with a persistent spatial representation from which grounding, planning geometry, and visual inspection draw their information. A planned manipulation predicts a change to that representation; local visual reconciliation produces a successor for the next instruction. The central contribution is continuity between tasks: object identity, geometry, and the information supplied to the planner advance together. A sequential study examines this continuity in a captured scene using simulated manipulation plans and synthetic post-action imagery, distinguishing the predicted world from evidence of what a physical robot has actually changed.

4.1 Planning Across World Changes

Consider the sequence in Fig. 4.1. The robot first retrieves a blue ball from beneath a table and leaves it beside a conifer. A later instruction refers to that same ball and asks the robot to place it on the table. The first plan may be perfectly grounded and feasible, but its success invalidates the map from




Figure 4.1: **A robot can invalidate the map it just used.** Instruction 1 moves the ball from under the table to beside the conifer. If instruction 2 is grounded against the original world W_0 , the ball is still represented under the table. REALM instead promotes the evolved world W_1 , preserves the ball’s identity, regenerates the representations used for grounding and planning, and carries the task forward to W_2 . The images are Gaussian renders from the simulated garden sequence; the final panel illustrates the planned terminal configuration. The displayed instructions are abbreviated; Section 4.5 gives the evaluated wording.

which it was produced. If the next instruction is grounded against the original reconstruction, the robot searches beneath the table for an object that its own action moved. This is not principally a failure of language understanding or motion planning, but a failure to represent the world as something that evolves under robot intervention. Reconstructing the scene from scratch does not completely solve the problem: the moved object may be occluded, its former location may expose unseen background, and a new detection need not preserve the identity used by the preceding plan. The robot must connect what it intended to change, what it subsequently observes, and what its next task is allowed to believe.

The problem is to preserve this continuity without reconstructing the entire scene from scratch after each action. The functional view of world models in Section 2.5 identifies the capabilities required: grounding in the environment, prediction of task-relevant change, and an account of the intervention that causes it. REALM supplies these capabilities through an explicit Gaussian scene and mechanisms for changing it under planned manipulation. The Gaussian representation provides editable geometry and rendered views; the transition and update mechanisms make it useful as a spatial world model.

REALM maintains an explicit Gaussian world as the persistent base representation and constructs the semantic, geometric, and visual abstractions needed for task-and-motion planning at runtime from the current world state. Grounded objects, task-progression semantics, digital-twin geometry, collision clouds, and rendered observations are instruction-specific abstractions of that state. For a planned manipulation, REALM first predicts its geometric consequence using an analytic forward model, then

reconciles its visual consequences by constructing registered post-action evidence, and finally promotes a successor from which the next instruction is grounded and planned.

This sequence separates geometry predicted from the plan, appearance inferred from visual evidence, and the successor accepted for reuse after validation. *Predicting physical change* applies a forward model based on the object’s terminal pose in the validated plan. *Reconciling visual consequences* renders the prediction from calibrated viewpoints and uses a pretrained image-generation foundation model to synthesize appearance changes within geometry-derived masks—physical observations can complement or replace these synthetic re-observations through the same interface. *Committing the updated world* uses the resulting RGB-D evidence to locally update the affected Gaussians, validates the result, and promotes the accepted successor. Stable ownership binds grounded objects to the Gaussian rows that move, while immutable history preserves the worlds produced over time.

The sequential study in Figure 4.1 follows one object through this process in a captured real scene, using synthetic post-action RGB-D. The ball begins under the table in W_0 . Instruction 1 derives its grounding and planning abstractions from W_0 and predicts the ball beside the conifer. After reconciliation and promotion, instruction 2 derives a new grounding, twin, collision cloud, and plan from W_1 , associates the ball with the same identity at its evolved location, and predicts W_2 on the table. Across both transitions, REALM transforms the same 254 object-owned Gaussians, preserves every non-moved row exactly during prediction, improves agreement with synthetic targets on held-out views after reconciliation, and recovers every ancestor exactly. The chapter makes four contributions:

1. an action-conditioned spatial world model that evolves one persistent, renderable 3D state authority from which instruction-specific abstractions used for language grounding and planning can be generated;
2. stable object ownership that gives robot actions explicit support inside a monolithic Gaussian scene and preserves object identity across tasks;
3. a prediction, reconciliation, and promotion process that distinguishes geometric predictions based on planned actions from appearance inferred through visual evidence while retaining an immutable world history; and
4. a sequential study with simulated mobile-manipulation plans that regenerates the next instruction’s abstractions from the world produced by the preceding plan rather than from the initial capture.

4.2 Related Work

Semantic scenes and dynamic memory. Hydra builds hierarchical metric-semantic scene graphs online (Hughes et al., 2022), and ConceptGraphs constructs open-vocabulary object-centric graphs for perception and planning (Gu et al., 2024). F3RM distills language-aligned features into 3D fields for few-shot manipulation (Shen et al., 2023); D³Fields combines dynamic geometry, semantic descriptors, and instance structure for zero-shot rearrangement (Wang et al., 2025b). Dynamic Scene

Representation explicitly studies permanence, amodal completion, and spatiotemporal continuity for predictive manipulation (Xu et al., 2021). DynaMem updates an online point-cloud memory as objects move, appear, and disappear (Liu et al., 2025a), and DovSG updates an open-vocabulary scene graph for long-term mobile manipulation (Yan et al., 2025). REALM shares the aim of persistent semantic state, but couples each object identity to the rows transformed by a specified planned action and to the history consumed by later planning.

Gaussian worlds and continual reconstruction. 3D Gaussian Splatting provides an explicit scene representation that is both differentiable and efficiently renderable (Kerbl et al., 2023). CL-Splats identifies changed regions, optimizes them locally, and retains temporal reconstruction history (Ackermann et al., 2025). Gaussian digital twins can also connect semantic reconstruction to collision geometry and physical pick-and-place planning (Sun et al., 2026). GaussMemory learns task-dependent retention and update in a persistent Gaussian memory (Hu et al., 2026). REALM builds on these representational capabilities, placing them inside a robot-facing state lifecycle in which a plan predicts a successor and promotion determines what the next task sees.

Action-conditioned world models. Predictive world models for manipulation take many forms. Physically Embodied Gaussian Splatting couples particles and Gaussians to predict and visually correct physical dynamics (Abou-Chakra et al., 2025). DreMa constructs compositional Gaussian and physics worlds for imagined robot data (Barcellona et al., 2025). GWM and MRO-GWM learn action-conditioned future Gaussian states for policy learning and model-predictive control (Lu et al., 2025; Kreber et al., 2026); GaussianDream decodes future Gaussians as dense supervision for vision-language-action learning (Zhang et al., 2026); GAF represents motion and interaction through a 4D Gaussian action field (Chai et al., 2026); and ParticleFormer predicts multi-object, multi-material point-cloud dynamics (Huang et al., 2025). REALM occupies a complementary regime: the task planner supplies a validated rigid-object outcome, and the world model turns that outcome into a persistent, inspectable, evidence-reconciled state for future language-guided behavior.

Across these literatures, rich spatial representations, action prediction, and language-grounded planning are often developed as separate capabilities. The central role of REALM is to close the loop among them: the representation used to interpret an instruction is updated under its planned action and becomes the representation used to interpret what comes next.

4.3 One World, Many Instruction-Specific Views

Using the Gaussian representation introduced in Section 2.6, let a scene contain rows

$$g_i = (\mu_i, s_i, q_i, \alpha_i, c_i), \quad i \in \{1, \dots, N_t\}, \quad (4.1)$$

where $\mu_i \in \mathbb{R}^3$ is the mean, $s_i \in \mathbb{R}^3$ the componentwise log-scale (rather than the positive scale used in Chapter 2), q_i a unit quaternion specifying orientation, $\alpha_i \in [0, 1]$ the opacity, and c_i the spherical-harmonic appearance coefficients. Gaussian parameters are indexed within a particular revision; their

time index is suppressed for readability. A stable object registry contains identities $o \in \mathcal{O}_t$, poses $X_t^o \in \text{SE}(3)$, and an ownership function

$$\omega_t(i) \in \{0\} \cup \mathbb{N}^+, \quad (4.2)$$

where zero denotes static or unowned scene support and a positive value identifies one movable object. The registry $\mathcal{O}_t$ stores these ownership assignments together with each object record. A pose $X_t^o = (R_t^o, t_t^o)$ maps object-local coordinates into the common world frame: $x^{\text{world}} = R_t^o x^{\text{object}} + t_t^o$.

Let t index planned manipulation transitions. Prediction and reconciliation create distinct revisions at the same logical time; their parent links distinguish them. We represent a persistent world revision as

$$W_t = (\mathcal{G}_t, \mathcal{O}_t, H_t, M_t). \quad (4.3)$$

H_t links the state to its immutable parent, while M_t records the action, evidence provenance, and validation information that explain the transition. This state is the persistent authority. It is not itself a fixed planner input: an instruction ℓ_t asks REALM to construct the abstractions needed at that moment,

$$\mathcal{A}(W_t, \ell_t) = (G_t, S_t, D_t, C_t), \quad (4.4)$$

where G_t contains grounded objects, S_t is the task-progression semantic state, D_t is task-relevant digital-twin geometry, and C_t is the planning collision representation. Grounding views rendered from W_t support the construction of G_t ; they are not a second world model. Planning and subsequent visual inspection are

$$\pi_t = \Pi(\ell_t; G_t, S_t, D_t, C_t), \quad (4.5)$$

$$V_t^{\text{insp}} = \mathcal{R}_{\text{plan}}(W_t, \pi_t). \quad (4.6)$$

Thus inspection renders are produced from a candidate plan rather than used to create it. The instruction determines which abstractions are materialized, but every abstraction, plan, and inspection view remains bound to the same W_t .

Suppose the validated plan π_t yields a manipulation transition a_t with released object pose $\hat{X}_{t+1}^o$. The transition a_t includes the planned endpoint. For an accepted reconciled successor, REALM evolves the world in three stages:

$$\begin{aligned} \widehat{W}_{t+1} &= \text{Predict}(W_t, a_t), \\ B_{t+1} &= \text{Reconcile}(W_t, \widehat{W}_{t+1}), \\ W_{t+1} &= \text{Commit}(\widehat{W}_{t+1}, B_{t+1}). \end{aligned} \quad (4.7)$$

Predict transforms the known object support according to the validated endpoint; Reconcile renders that prediction and constructs registered post-action evidence B_{t+1} ; and Commit uses that evidence to update only affected Gaussian support, validates the result, and grants the successor current authority. The labels *known*, *inferred*, and *trusted* in the method diagram refer to these roles: an explicit planned endpoint, visually inferred consequences, and a successor accepted by the representation checks. They do not establish that the predicted endpoint was physically realized.

For calibrated view k , let κ_k contain the camera intrinsics and pose in the world frame. The renderer produces predicted RGB $\widehat{I}$ and depth $\widehat{Z}$; the image-editing model E produces RGB only:

$$(\widehat{I}_{t+1}^{(k)}, \widehat{Z}_{t+1}^{(k)}) = \mathcal{R}(\widehat{W}_{t+1}, \kappa_k), \quad (4.8)$$

$$I_{t+1, \text{syn}}^{(k)} = E(I_t^{(k)}, \widehat{I}_{t+1}^{(k)}, \mathcal{M}_{t \rightarrow t+1}^{(k)}, r_o), \quad (4.9)$$

$$Z_{t+1, \text{syn}}^{(k)} = \widehat{Z}_{t+1}^{(k)}. \quad (4.10)$$

Here $I_t^{(k)}$ is the parent RGB used for editing, $\mathcal{M}_{t \rightarrow t+1}^{(k)}$ contains source-removal and target-insertion masks, and r_o is an object-appearance reference. The registered bundle is

$$B_{t+1} = \{(I_{t+1}^{(k)}, Z_{t+1}^{(k)}, \mathcal{M}_{t \rightarrow t+1}^{(k)}, \kappa_k)\}_{k=1}^{n_v}, \quad (4.11)$$

where n_v counts views, separately from the N_t Gaussian rows. Each view also records provenance and its fitting or validation role. Physical RGB-D can populate the same bundle, with change masks derived from observations. Equations 4.9–4.10 describe the synthetic branch used here, pairing generated RGB with depth rendered from the geometric prediction.

Commit lifts the registered multi-view evidence into the Gaussian map. Let $\mathcal{L}(\widehat{W}_{t+1}, B_{t+1})$ denote the existing row indices supported as changed by that evidence, $\mathcal{J}_{t+1}^{\text{act}}$ the indices owned by objects changed by the validated action, and $\mathcal{J}_{t+1}^{\text{new}}$ newly introduced support for disoccluded content. The initial active set for fitting is

$$\mathcal{J}_{t+1} = \mathcal{L}(\widehat{W}_{t+1}, B_{t+1}) \cup \mathcal{J}_{t+1}^{\text{act}} \cup \mathcal{J}_{t+1}^{\text{new}}. \quad (4.12)$$

With U_{CL} denoting local fitting based on CL-Splats followed by restoration of action-owned geometry and opacity, and P denoting validation and promotion, the final stage expands to

$$\begin{aligned} \widetilde{W}_{t+1} &= U_{\text{CL}}(\widehat{W}_{t+1}, B_{t+1}; \mathcal{J}_{t+1}), \\ W_{t+1} &= P(\widetilde{W}_{t+1}). \end{aligned} \quad (4.13)$$

Together these two operations implement Commit in Eq. 4.7. Only support in $\mathcal{J}_{t+1}$ and any descendants introduced by local densification may be optimized; unrelated Gaussian support is frozen. Action-owned geometry and opacity are restored to their predicted values, preserving the planned object configuration through visual reconciliation. Equation 4.13 describes successful promotion. A validated prediction may first become current while awaiting reconciliation; if the reconciled candidate fails validation, the preceding current revision is retained (Appendix B.1). After successful promotion, the next instruction calls $\mathcal{A}(W_{t+1}, \ell_{t+1})$ rather than reusing $\mathcal{A}(W_t, \ell_t)$. This regeneration closes the loop from one instruction to the next.

The evolution is governed by six design requirements (Table 4.1). Together they keep the persistent representation local, continuous, reversible, and usable by every task-specific consumer.

4.4 REALM

At each timestep, REALM plays two connected roles. First, it is the base representation from which the robot constructs whatever the current instruction requires. Second, after a plan is produced, it is the

Table 4.1: Design requirements for an evolving robot world.

Requirement	Meaning
Local change	An action changes the owned object without moving unrelated rows.
Stable identity	One object ID connects grounding, prediction, twin geometry, and later reuse.
Persistent history	A successor is added without overwriting the world from which it came.
Split authority	The planned action fixes predicted object geometry; evidence updates appearance and exposed scene content.
Validated promotion	A reconciled candidate replaces the current prediction only after spatial, visual, and lineage checks.
Synchronized views	Grounding, planning, rendering, and twin artifacts are generated from one named state.

model that predicts and absorbs the resulting change. The recurring progression is therefore *represent* $\rightarrow$ *derive task views* $\rightarrow$ *ground and plan* $\rightarrow$ *predict physical change* $\rightarrow$ *reconcile visual consequences* $\rightarrow$ *commit and regenerate*. Figure 4.2 visualizes this progression.

4.4.1 From the Persistent World to the Current Instruction

At mapping time, REALM starts from posed scene images and a renderable Gaussian reconstruction. In the garden study, the captured input is the Mip-NeRF 360 image dataset (Barron et al., 2022); depth used for subsequent view lifting is rendered from the reconstruction. The Gaussian scene is the persistent spatial representation, but it is not the task plan. Given an instruction such as “get the blue ball from under the table and leave it beside the potted lime-green conifer,” REALM first renders views of W_t in which the named referents can be detected. Multi-view detections are lifted into 3D, consolidated into stable object records, and inserted into a task-progression semantic state.


From the same W_t , REALM constructs only the geometric abstractions required by that instruction. Task-relevant surfaces and objects become a digital twin for manipulation; the remaining geometry becomes a collision cloud for navigation and whole-body planning. LIMP translates the language into temporal constraints and produces navigation, pick, and release trajectories from these state-bound abstractions (Quartey et al., 2025). Because the underlying world is renderable, the robot and candidate trajectories can also be composed with Gaussian views from several cameras for visual plan inspection. Grounding views, semantic records, twin geometry, collision geometry, and plan renders serve different consumers, but they all depict the same named world.

Each state stores a checkpoint, object registry, digital twin, planning cloud, optional observation bundle, source plan, transition record, and validation report. Artifact digests bind these products to the revision from which they were derived. The scientific purpose of that binding is consistency: a grounded object cannot come from one world while the collision cloud or plan comes from another.

4.4.2 Grounding Says What the Task Means; Ownership Says What Moves

Grounding identifies the semantic object referred to by the instruction. Ownership connects that object identity to the part of the persistent 3D state that may follow an action. A broad crop first proposes Gaussian rows. For candidate i , REALM measures positive depth-consistent votes p_i , visible

(a) One State Authority, several instruction-specific views



(b) Action-conditioned evolution: predict, reconcile, commit




Figure 4.2: **One persistent world, rebuilt for each instruction.** At timestep t , REALM renders grounding views from the current Gaussian world W_t and constructs grounded objects, task-progression semantics, digital-twin geometry, and collision geometry. Together these instruction-specific abstractions produce a task-and-motion plan, which can then be rendered against the same state for inspection. Its validated endpoint predicts physical change (*known*); rendering and re-observation construct the reconciled visual evidence (*inferred*); and commit locally updates, validates, and grants W_{t+1} authority for future use (*trusted*). The next instruction regenerates its own abstractions from W_{t+1} , not from the original capture.

contradictions n_i , visible views v_i , footprint-supported views f_i , fused-surface distance d_i^S , oriented-box overflow d_i^B , maximum Gaussian scale $\sigma_i^{\max}$ with $\sigma_i = \exp(s_i)$ componentwise, and anisotropy r_i . The selected support is

$$m_i = [p_i \geq \tau_p] \wedge [n_i / \max(v_i, 1) \leq \tau_n] \wedge [f_i \geq \tau_f] \\ \wedge [d_i^S \leq \tau_S] \wedge [d_i^B \leq \tau_B] \wedge [\sigma_i^{\max} \leq \tau_s h_o] \wedge [r_i \leq \tau_r], \quad (4.14)$$

where each bracket is a Boolean test, $\wedge$ requires every test to pass, and h_o is the maximum object half-extent. The distances d_i^S, d_i^B , scales σ_i , and half-extent h_o use the same world-coordinate units; τ_s is dimensionless. The selected checkpoint indices and object-local means are stored with the same digital twin used for planning. Empty, duplicate, overlapping, and insufficient supports are rejected: if a background row is owned by the ball, every future action on the ball would incorrectly move that background as well.

4.4.3 Predict Physical Change (Known from the Plan)

The task-and-motion plan is not merely an output rendered for inspection. Its validated terminal pose provides an intervention prior for the world model. REALM composes the state-bound navigation and manipulation artifacts to obtain the released pose of every changed object. A prediction is accepted only when the plan, trajectories, grounding, checkpoint, and current state agree.




Figure 4.3: **How the world evolves.** REALM predicts the spatial consequence specified by the plan, renders that prediction to reconcile visual consequences as registered multi-view evidence, then commits by locally updating affected Gaussian support, validating the result, and granting it authority for the next instruction. The four panels show one calibrated view from the clean garden run: parent evidence, the predicted Gaussian render, a geometry-constrained synthetic re-observation, and the committed Gaussian world.

For owned Gaussian i , write its object-local mean as $\bar{\mu}_i = (R_t^o)^\top (\mu_i - t_t^o)$. Its predicted world-frame mean is

$$\mu'_i = \hat{R}_{t+1}^o \bar{\mu}_i + \hat{t}_{t+1}^o. \quad (4.15)$$

Its orientation is composed with the relative rotation $\Delta R = \hat{R}_{t+1}^o (R_t^o)^\top$, and its real spherical-harmonic coefficients are rotated consistently through degree three. In covariance form, the geometric update is

$$\Sigma_i = R(q_i) \text{diag}(\exp(2s_i)) R(q_i)^\top, \quad \Sigma'_i = \Delta R \Sigma_i \Delta R^\top. \quad (4.16)$$

This preserves the Gaussian’s principal scales while rotating its axes. The spherical-harmonic rotation similarly carries its represented directional appearance into the new orientation. Scale and opacity remain unchanged. Every row outside the moved object’s support is copied exactly during prediction. The transition simultaneously evolves the object registry, twin, and planning cloud and publishes a reversible delta linking the predicted world to its parent.

The prediction therefore places the planned spatial outcome inside the same renderable representation used by perception. The ball expected beside the conifer is already `blue_ball_obj0`, at the plan’s released pose, before it is detected again. This identity-preserving expectation is what lets later evidence refine the represented consequence of a specified intervention rather than rediscover an unrelated object.

4.4.4 Reconcile Visual Consequences (Inferred from Evidence)

A rigid object transform is geometrically informative but visually incomplete. It does not reveal what lay behind the ball, determine appearance from every camera, or confirm that the successor is consistent across views. REALM therefore renders the predicted checkpoint at calibrated cameras, including full-scene, static-scene, old-object, and new-object RGB, alpha, and metric-depth components. These renders state what the world is expected to look like if the plan reaches its endpoint.

Registered RGB-D evidence enters through the same observation bundle whether it comes from physical sensors or a pretrained image-generation foundation model. Every view carries camera transforms, frame role, metric depth, masks, and provenance. Together these views form B_{t+1} : a registered account of the visual consequences that the rigid prediction alone cannot supply.

The reported study uses geometry-constrained synthetic post-action observations from an image-generation model for this stage. Newly exposed source regions are inpainted, and a canonical object appearance is inserted inside the predicted target silhouette. The Gaussian prediction remains the authority for masks and metric depth; generated color is confined to explicitly marked pixels, and original RGB is restored exactly outside each edit mask. This provides a controlled observation target for exercising the complete prediction–reconciliation loop.

4.4.5 Commit (Trusted for Reuse)

Commit turns the reconciled evidence into the next Gaussian world. Change masks are lifted through calibrated depth into the 3D scene, newly visible background may introduce unowned seed Gaussians, and REALM uses CL-Splats (Ackermann et al., 2025) to optimize only the affected region while unrelated content remains frozen.

Planning and observation have complementary authority. The validated action endpoint determines the transformed object’s geometry and opacity; observation refines appearance and the background revealed by its motion:

$$\underbrace{(\mu, s, q, \alpha)_{\text{owned}}}_{\text{action-conditioned prediction}} \quad \Bigg| \quad \underbrace{(c_{\text{owned}}, \mathcal{G}_{\text{disocc.}})}_{\text{observation-dependent update}} \quad . \quad (4.17)$$

After localized optimization, REALM restores the predicted means, scales, orientations, and opacities of owned rows while retaining reconciled appearance. The resulting candidate preserves the planned object configuration while incorporating appearance and exposed background inferred from the registered targets. This division of responsibility keeps the geometric transition and the visual update consistent within one persistent world.

The locally updated candidate is committed only after its artifacts are intact, tensors remain finite and aligned, ownership is consistent, predicted object geometry is preserved, held-out changed-region error meets its criterion, and the increase in outside-region reconstruction error remains bounded. Every parent edge is also replayed in reverse so that promotion never destroys the worlds that preceded it.

Promotion grants the successor current authority, carrying the world forward. For the next instruction, REALM does not patch the old semantic map or reuse the former planning scene. It renders new grounding views from W_{t+1} , associates detections near the predicted pose with the stable object identity, rebuilds task-progression semantics, constructs a new task-relevant twin and collision cloud, and binds the next plan to those products. The prior trajectory endpoint becomes the next robot start. In this way, the terminal world of one instruction becomes the base representation from which the next instruction is understood.

4.5 Following the World Across Two Instructions

We evaluate REALM through a study over a captured real scene that follows one blue ball across two spatially and temporally constrained mobile-manipulation instructions. The sequence traces the same progression used to explain REALM: construct task views from a persistent state, produce a plan, evolve the state from that plan, and reconstruct the views needed for what comes next. The initial world W_0 is a 6,125,086-row Gaussian splat reconstruction of a captured outdoor garden with calibrated source cameras (Barron et al., 2022).

4.5.1 Timestep 0: Build the First Task from W_0

Instruction 1 asks: “Get the blue ball from under the table and leave it beside the potted lime-green conifer.” Rendered views of W_0 ground the ball (254 rows), table, and conifer; their consolidated 3D records populate the semantic task state. The W_0 digital twin and collision cloud support the navigation, pick, and release plan, while Gaussian views provide several perspectives from which that plan can be inspected. Language grounding and temporal-logic planning follow LIMP (Quartey et al., 2025); Appendix B.3 gives the full specification.

The plan contains a 323-frame pick trajectory and a 249-frame release trajectory. Its validated terminal pose places `blue_ball_obj0` beside the conifer. REALM applies that pose to 254 Gaussian rows owned by the grounded ball to produce the predicted successor $\widehat{W}_1$.

4.5.2 From the First Plan to the Next World

To exercise the visual half of the transition, the study uses registered, geometry-constrained synthetic post-action RGB-D. Newly exposed source regions are inpainted and a canonical object appearance is inserted inside the predicted target silhouette. The Gaussian prediction remains authoritative for masks and metric depth, generated color is confined to marked pixels, and original RGB is restored outside each edit mask. Stage 1 contains 12 source-disocclusion and 9 target-insertion edits. Reconciliation produces a candidate world, and promotion makes the validated result W_1 current.

At this point the first instruction is over, but the representational process is not. W_1 —not W_0 —is now the source from which later grounding, semantics, geometry, and plans are constructed.

4.5.3 Timestep 1: Rebuild the Second Task from W_1

Instruction 2 asks: “Stop by the small garden ornament, collect the blue ball beside the conifer, and put it on the table.” REALM renders new views from W_1 , re-associates the detected ball with `blue_ball_obj0`, and rebuilds the semantic registry, digital twin, and collision cloud around its evolved pose. The new plan begins at instruction 1’s trajectory endpoint, contains a 443-frame pick and a 308-frame place trajectory, and predicts the ball on the table in $\widehat{W}_2$. Stage 2 reconciliation uses 12 source-disocclusion and 12 target-insertion edits before the successor is promoted as W_2 .

Table 4.2: Ownership selector analysis for the garden ball.

Selector	Rows	Prec./IoU	Recall
Legacy positive voting	724	0.387891	1.000
Stricter positive voting	521	0.519845	1.000
Lower depth tolerance	315	0.670623	1.000
OBB filter	281	0.646528	1.000
Footprint filter	648	0.496164	1.000
Scale/anisotropy filter	721	0.432294	1.000
Combined selector	254	0.740647	1.000

The trajectories are planned and simulated directly in the representation rather than executed on hardware. Validation checks kinematics, continuity, collision and contact policy, attachment, release, and artifact integrity.

4.5.4 What the Sequence Must Preserve

We examine four links in this progression. **Grounding-to-action support** asks whether ownership attaches the semantic ball to its own 3D content rather than nearby background. **Local world evolution** measures moved and inactive rows, inverse recovery, and action-owned drift. **Visual reconciliation** measures held-out L1 inside and outside the changed region. **Instruction-to-instruction continuity** asks whether the stable identity, grounding, checkpoint, twin, planning cloud, and robot start for instruction 2 all derive from promoted W_1 .

4.6 Results

4.6.1 Grounding Becomes Action Support Through Ownership

Rendered views first identify the ball as a semantic referent; ownership must then turn that referent into valid action support inside the Gaussian world. The original positive-voting selector assigned 724 rows to the ball, including 470 centers outside its fitted oriented bounding box. These rows contain table and ground support that would move whenever the ball moves. The combined selector reduces the support to 254 rows and raises mask precision/mean IoU from 0.388 to 0.741 while retaining measured recall of 1.0 (Table 4.2). Figure 4.4 shows the consequence: stable ownership makes the rigid transition object-local rather than transporting part of the scene with the object.

4.6.2 Known, Inferred, and Trusted States Remain Local and Reversible

Both predicted transitions move exactly the same 254 owned rows. At stage 1, all 6,124,832 inactive rows remain exactly equal to W_0 . Reconciliation introduces 12,170 unowned disocclusion seeds; the next prediction then preserves all 6,137,002 inactive rows exactly. The action-owned means, scales, orientations, and opacities have zero drift after reconciliation because the update restores them to their predicted values. Stored inverse deltas recover every parent exactly.




Figure 4.4: **Ownership defines what the action moves.** The legacy 724-row selector attaches background support to the ball. The combined 254-row selector remains object-local. Across the seven selector variants, the combined rule improves mask precision/mean IoU from 0.388 to 0.741 while preserving measured recall on the garden ball.

Table 4.3: Spatial and visual continuity across the two world transitions. Row counts describe rigid prediction; image errors compare the predicted and reconciled renders against synthetic targets on validation views. Geometry and opacity drift is measured after restoring the action-owned parameters.

Quantity	Stage 1	Stage 2
Owner rows transformed in prediction	254	254
Non-moved rows exact in prediction	6,124,832	6,137,002
Disocclusion seeds	12,170	3,551
Masked L1 before	0.138485	0.134815
Masked L1 after	0.121646	0.131212
Outside L1 before	0.022277	0.029291
Outside L1 after	0.022299	0.029311
Max owned geometry/opacity drift	0	0

Reconciliation improves agreement with synthetic RGB targets on views held out from fitting. Masked L1 decreases from 0.138485 to 0.121646 at stage 1, a 12.2% reduction, and from 0.134815 to 0.131212 at stage 2, a 2.7% reduction. Outside-region L1 increases by approximately 2.2×10^{-5} and 2.0×10^{-5} , respectively, within the 10^{-4} promotion criterion (Table 4.3). These held-out views provide the visual criteria for promotion: the reconciled state improves agreement with the synthetic targets in the edited region while keeping the increase in outside-region reconstruction error within the prescribed tolerance.

4.6.3 The Next Instruction Rebuilds Its World from W_1

Promotion of W_1 triggers the same abstraction-building process that began at W_0 , now over an evolved world. New rendered views produce 15 ball detections that associate with `blue.ball.obj0`. The selected pose lies 0.01879 m from the action prediction and 3.7198 m from the original under-table pose. From this W_1 pose, REALM rebuilds the grounding record, semantic task state, digital twin, and planning cloud for instruction 2. The pick target is therefore beside the conifer, and the robot start is instruction 1’s terminal state.

After instruction 2, the W_2 checkpoint, twin, semantic registry, and planning cloud agree on the ball’s final table pose. The important continuity is not only that one coordinate changed. The




Figure 4.5: **Sequential world evolution across two instructions.** The ball moves from its observed pose beneath the table to promoted W_1 beside the conifer and then to promoted W_2 on the table. Reconciliation improves held-out changed-region consistency at both stages. Each prediction moves the same 254 owned rows and preserves all other rows exactly; stored history recovers the ancestors. Fifteen detections across rendered views of W_1 associate to one stable object ID.

persistent representation and every instruction-specific abstraction derived from it advance together: the object keeps its identity, grounding searches the updated location, planning receives geometry from that same revision, and history explains how the represented world arrived there.

4.6.4 History Preserves How the Current World Was Reached

The complete sequence contains five distinct states:

$$W_0^{\text{obs}} \rightarrow \widehat{W}_1 \rightarrow \widetilde{W}_1^{\text{syn}} \rightarrow \widehat{W}_2 \rightarrow \widetilde{W}_2^{\text{syn}}.$$

Final validation walks all four parent edges, checks the state-bound artifacts and twin ownership, and reconstructs every ancestor exactly. The archived integration suite records 98/98 passing tests; a later targeted audit records 49/49 passing tests across state, ownership, reconciliation, re-grounding, and orchestration. The history makes prediction, observation, and promotion visible as different events in the robot’s evolving belief.

4.7 Discussion

4.7.1 From Verifiable Plans to an Evolving World

LIMP establishes a path from expressive language to a grounded, verifiable plan. REALM extends that path through time. The plan no longer terminates the representational process; it becomes an

intervention that predicts a new world. That successor is reconciled, promoted, and used to interpret the next instruction. Together, the two systems connect an interpreted instruction to a planned intervention and to the world representation used for the next task.

The world interface concerns the premises of the planner’s reasoning. Even when an instruction has been interpreted correctly, a plan grounded in an obsolete scene can concern the wrong location or geometry. Carrying the representation forward preserves the connection established by language grounding across successive tasks.

4.7.2 One Representation, Four Connected Uses

The Gaussian scene supports grounding, planning geometry, visual inspection, and action-conditioned updating through one persistent representation. These uses depend on different information, but must agree about which world they describe. The sequential study makes that dependency concrete: updating the ball’s rendered appearance would be insufficient if grounding or collision checking continued to use its original location.

That distinction explains why regeneration matters. The semantic task state, digital twin, collision cloud, and plan views are useful precisely because they describe the current world. After promotion, rebuilding them from W_{t+1} carries object identity and geometry forward together. Reusing any one of them from W_t would reintroduce the stale-map failure through a different interface.

4.7.3 From a Renderable Scene to a Spatial World Model

The distinction developed in Section 2.5 is now visible operationally. A Gaussian render provides a view of a scene, whereas REALM’s action-conditioned update produces a successor on which later grounding and planning can operate. Its world-modeling role comes from representing task-relevant change while preserving object identity.

World models are often associated with learned latent dynamics or generated future video. REALM instead models a task-relevant spatial consequence in an explicit 3D representation. The action transition is analytic; visual reconciliation can contain learned components; and the promoted state remains directly renderable and consumable by later grounding and motion planning.

The known–inferred–trusted decomposition makes this hybrid construction explicit and is useful precisely because the planner already knows an important part of the future: where a successfully manipulated rigid object should be released. Prediction records what the plan says should change, reconciliation records what subsequent evidence adds, and promotion records which world later reasoning is allowed to consume.

4.7.4 From an Evolving World to Composable Skills

The garden study shows how the same object can remain available to later grounding and planning after its represented location changes. Its evidence concerns continuity and local reconstruction under planned rigid-object motion. Because the manipulation is simulated and the post-action imagery is

synthetic, agreement with the updated scene does not establish successful physical execution. In a physical deployment, new observations would need to correct predictions when a grasp fails, an object slips, or the environment changes independently of the robot.

An evolving representation still leaves the planner dependent on supplied skills. Learning additional policies broadens what the robot can do, but composition also requires information about when those policies can be used and what conditions their actions are expected to produce. Chapter 5 develops this symbolic interface through joint learning of actions and predicted predicate trajectories.

Chapter 5

Learning Composable Skills

This chapter is adapted from *Jointly Learning Predicates and Actions Enables Zero-Shot Skill Composition*, coauthored by Benedict Quartey, Sebastian Castro, Eric Rosen, Wil Thomason, George Konidaris, and Stefanie Tellex (Quartey et al., 2026). The authorial “we” refers to this collaborative research.

A planner gains new possibilities when it can compose behaviors the robot has learned, rather than depend entirely on a supplied library of controllers. The preceding chapters established how expressive instructions can become grounded plans and how the representation used for planning can evolve across tasks. Both developments still depend on suitable skills being available. Learning those skills from demonstrations offers a way to extend the robot’s repertoire, but a policy that produces movements supplies only part of what composition requires. The planner must also reason about the conditions those movements are expected to produce.

This chapter develops *composable state-reporting skills*, addressing the action interface of trust. PACTS learns a skill’s control policy together with a symbolic trace of its expected effects, connecting continuous behavior to the predicates used in task planning. The resulting interface supports new compositions of learned skills and replanning from updated predicate predictions. The contribution complements LIMP’s grounding of intent and REALM’s evolving world representation by examining how learned behavior can participate in symbolic reasoning. The experiments study this skill interface with a supplied predicate vocabulary and symbolic domain; their relationship to the preceding systems motivates the synthesis in Chapter 6.

5.1 Learning Actions Together with Their Symbolic Outcomes

The generative policies introduced in Section 2.7 can represent several ways to carry out a behavior. For composition, however, plausible movements are insufficient: the symbolic conditions associated with a particular action sequence matter to the next skill. Different ways to manipulate an object can lead to different outcomes, even when both begin from the same observation. A policy and a separately




Figure 5.1: **Predicate-Action Skills**. Conditioned on current observations $\mathbf{o}$, we model skills as a joint generative process over an action trajectory $\mathbf{x}$ and a predicate-belief trajectory $\mathbf{z}$ by learning the coupled distribution $p(\mathbf{x}, \mathbf{z} | \mathbf{o})$. Starting from noise $(\mathbf{x}_T, \mathbf{z}_T)$, our model iteratively refines both modalities to produce paired action–outcome rollouts $(\mathbf{x}_0, \mathbf{z}_0)$. The resulting predicate-belief trajectory $\mathbf{z}_0$ provides an online symbolic interface for monitoring skill execution and planning-based skill composition using off-the-shelf planners.

sampled outcome predictor may therefore choose incompatible alternatives. The distinction between their individual accuracy and their agreement under execution, developed in Section 2.8, motivates learning the two trajectories together.

Predicate-Action Skills (PACTS) represent a skill through a joint distribution over actions and predicted predicate values. A sample contains both a control trajectory and the symbolic trace expected to accompany it (Figure 5.1). This formulation allows the model to represent dependence between the chosen action mode and its expected effects. It also provides a practical interface through which a planner can assess progress, select subsequent skills, and reconsider a sequence as new observations condition later predictions. Whether the predicted effects agree with execution is an empirical question, evaluated separately from predicate classification accuracy.

We instantiate this joint model with diffusion and flow-matching objectives and with UNet and Transformer architectures. The evaluation examines action performance, predicate prediction, and action–predicate coherence on a compositional 2D benchmark and two 3D RoboMimic tasks (Mandlekar et al., 2022, 2023). Planning-based composition is demonstrated in simulation and in a real-world cube-packing task. The contributions are the joint generative skill formulation, a predicate-based interface for composition and monitoring, and a segmentation and labeling toolkit that prepares skill-specific demonstrations for joint learning.

5.2 Related Work

Generative visuomotor policies for imitation learning Behavior cloning learns skills from demonstrated observation–action pairs (Mandlekar et al., 2022), and recent work models *distributions* over action trajectories to capture multiple ways to perform a behavior. Diffusion-based policies (Chi et al., 2025) and conditional flow matching (Chisari et al., 2025) provide two such formulations. OpenVLA and π_0 connect pretrained vision-language representations to robot control (Kim

et al., 2025; Black et al., 2025). Other approaches structure the learning problem through equivariance (Wang et al., 2025a), object-centric representations (Yuan et al., 2022), or affordance-centered task frames (Rana et al., 2025). Diffusion has also been used for planning: Diffuser (Janner et al., 2022) denoises state–action trajectories to enable flexible test-time conditioning, whereas Diffusion Policy (Chi et al., 2025) treats control as learning the distribution over action sequences, conditioning on observations without denoising them for efficient real-time inference. We adopt this conditional-policy perspective, but depart from action-only modeling by learning a coupled generator $p(\mathbf{x}, \mathbf{z} \mid \mathbf{o})$ that jointly samples actions and a low-dimensional predicate-belief rollout $\mathbf{z}$ aligned with PDDL preconditions/effects, yielding an outcome-aware interface for execution-time monitoring and planning-based composition. Vision–language–action models similarly couple control with language semantics, but their semantics are typically part of the conditioning interface rather than an explicit planner-consumable outcome trajectory.

Skills, abstraction, and planning-based composition The symbolic planning framework in Section 2.2 requires both a physical implementation of each operator and a grounding for its predicates. In robotic systems, these abstract actions are typically instantiated as hand-designed controllers, including statecharts (Harel, 1987) and behavior trees (Colledanchise and Ögren, 2018), and predicates are grounded through engineered perception and state estimation modules (Migimatsu et al., 2022). Several frameworks support such symbolic planning architectures (Diab et al., 2020; Martín et al., 2021; Mayr et al., 2023; Paulius et al., 2025; Quartey et al., 2025).

Recent approaches reduce this manual engineering by learning the components that connect behavior to symbolic planning. BLADE learns compositional behaviors and their preconditions and effects from demonstrations and language (Liu et al., 2025b); Pixels-to-Predicates uses pretrained vision–language models to ground predicates for symbolic world-model learning (Athalye et al., 2026). These approaches establish symbolic interfaces for planning, but do not jointly sample an action trajectory and its associated predicate trajectory. In long-horizon manipulation, actions and symbolic outcomes are tightly coupled: different action *modes* correspond to different predicate transitions, and treating them as separate predictions can lead to inference-time inconsistency. PACTS addresses this coupling directly by modeling a skill as a *joint* generative process over action trajectories and predicate-belief trajectories. Joint sampling allows the model to represent the dependence between action modes and symbolic outcomes, which independently sampled heads do not explicitly represent.

5.3 Jointly Modeling Predicates and Actions

5.3.1 Problem definition: joint action–outcome skill model

Using the trajectory conventions from Sections 2.7 and 2.8, we model a closed-loop skill as a conditional distribution over a horizon- H action trajectory $\mathbf{x} = (\mathbf{a}_t, \dots, \mathbf{a}_{t+H-1}) \in \mathbb{R}^{H \times d_a}$, where each $\mathbf{a}_{t+h} \in \mathbb{R}^{d_a}$ is a single control action. To support planning-based skill composition, we additionally model the symbolic outcomes induced by actions using a *predicate-belief trajectory* $\mathbf{z} = (\mathbf{z}_{t+1}, \dots, \mathbf{z}_{t+H}) \in$




Figure 5.2: **Architectures and baselines.** (a) PACTS uses a *single* conditional generative model to sample *joint* action–predicate rollouts, refined together via DDPM denoising or flow matching. To *disentangle* joint modeling from multi-task learning, we include a multi-task baseline (b) that shares the same observation encoder but predicts predicates and actions with *separate* heads trained with a combined loss. Standalone baselines isolate each component: (c) a generative *action-only* generator that generates actions, and (d) a *predicate-only* model predicts predicates from observations.

$[0, 1]^{H \times J}$, where $\mathbf{z}_{t+h,j}$ is a soft prediction that predicate $j \in \{1, \dots, J\}$ will hold at future step $t + h$. These values are called predicate beliefs in this chapter; their range does not establish probabilistic calibration.

Given the current observation $\mathbf{o}_t$, PACTS learns the coupled distribution $p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t)$ so that actions and symbolic outcomes form one generated sample. We implement this by defining the joint trajectory variable $\mathbf{y} \triangleq [\mathbf{x} \mid \mathbf{z}] \in \mathbb{R}^{H \times (d_a + J)}$ and learning a single conditional generative model $p_\theta(\mathbf{y} \mid \mathbf{o}_t)$. This joint modeling choice is the key distinction from decoupled or multi-task baselines: $\mathbf{z}$ is not produced by an independent classifier head, but is part of the generated trajectory sample.

Predicate-belief representation. In the dataset, predicates are provided as binary labels in $\{0, 1\}^J$. We treat predicate channels as continuous during generation by linearly normalizing labels to $[-1, 1]$ and learning to generate them jointly with actions. At inference, the model maps the generated predicate channels back to $[0, 1]$, yielding soft beliefs that can be thresholded when discrete predicate evaluations are required (e.g., checking preconditions/effects during planning).

5.3.2 DDPM formulation

The DDPM construction in Section 2.7.2 is applied here to the joint variable $\mathbf{y} = [\mathbf{x} \mid \mathbf{z}]$. Following Diffusion Policy (Chi et al., 2025), the generator conditions on $\mathbf{o}_t$ without generating future images or a full physical state trajectory. It does predict future symbolic conditions through the predicate channels. Let $q(\mathbf{y}^k \mid \mathbf{y}^{k-1})$ denote the forward noising process with a standard noise schedule $\{\alpha_k\}_{k=1}^K$. We train a noise-prediction network ϵ_θ , conditioned on $\mathbf{o}_t$, to predict the injected noise:

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E}_{\mathbf{o}_t, \mathbf{y}^0, k, \epsilon} \left[\left\| \epsilon_\theta(\mathbf{o}_t, \mathbf{y}^k, k) - \epsilon \right\|^2 \right], \quad (5.1)$$

where $\mathbf{y}^k$ is obtained by noising the clean joint trajectory $\mathbf{y}^0 = [\mathbf{x} | \mathbf{z}]$ at *denoising step* k . At inference, the model samples $\mathbf{y}^K \sim \mathcal{N}(0, I)$ and iteratively denoises it to obtain $\mathbf{y}^0$.

5.3.3 Conditional Flow Matching formulation

The joint object $\mathbf{y} = [\mathbf{x} | \mathbf{z}]$ can also be learned via *conditional flow matching* (CFM) (Tong et al., 2024; Lipman et al., 2023; Black et al., 2025), which targets the same conditional distribution $p_\theta(\mathbf{y} | \mathbf{o}_t)$ but uses a different objective and sampler. Let $\mathbf{y}^1 \sim \mathcal{N}(0, I)$ and use the linear interpolation from Section 2.7.3 between data $\mathbf{y}^0$ and noise $\mathbf{y}^1$ as

$$\mathbf{y}_\tau = \tau \mathbf{y}^1 + (1 - \tau) \mathbf{y}^0, \quad \tau \in [0, 1], \quad (5.2)$$

with constant target velocity $\mathbf{u}_\tau = \frac{d}{d\tau} \mathbf{y}_\tau = \mathbf{y}^1 - \mathbf{y}^0$. A conditional vector field $v_\theta(\mathbf{o}_t, \mathbf{y}_\tau, \tau)$ is trained by regression:

$$\mathcal{L}_{\text{CFM}} = \mathbb{E}_{\mathbf{o}_t, \mathbf{y}^0, \mathbf{y}^1, \tau} \left[\|v_\theta(\mathbf{o}_t, \mathbf{y}_\tau, \tau) - (\mathbf{y}^1 - \mathbf{y}^0)\|^2 \right]. \quad (5.3)$$

At inference, samples are obtained by integrating from noise ($\tau=1$) to data ($\tau=0$), yielding $\mathbf{y}^0$.

DDPM and flow matching instantiate the same modeling choice—learning the coupled action–outcome distribution $p_\theta(\mathbf{x}, \mathbf{z} | \mathbf{o}_t)$ —but provide complementary training lenses with different sampling trade-offs.

5.3.4 Architecture

As shown in Fig. 5.2(a), PACTS factorizes into (i) an observation encoder and (ii) a conditional trajectory generator. The observation encoder maps the current observation $\mathbf{o}_t$ into a latent embedding used to condition the generator. The generator implements either a DDPM noise-prediction network $\epsilon_\theta(\mathbf{o}_t, \mathbf{y}^k, k)$ (Sec. 5.3.2) or a CFM velocity field $v_\theta(\mathbf{o}_t, \mathbf{y}_\tau, \tau)$ (Sec. 5.3.3), operating on the joint trajectory $\mathbf{y} = [\mathbf{x} | \mathbf{z}]$.

We instantiate the trajectory generator with two backbones: (i) a temporal UNet that applies 1D convolutions over the horizon dimension, and (ii) a Transformer that treats time steps as tokens and models long-range temporal dependencies via self-attention. In both cases, conditioning enters through (a) the observation embedding and (b) the diffusion/flow time parameter (discrete k for DDPM, continuous τ for CFM), injected via standard conditional modulation.

Like prior trajectory-generating visuomotor policies, PACTS can be executed in a receding-horizon manner: at each control cycle, the policy conditions on the current observation $\mathbf{o}_t$, samples a length- H action chunk $\mathbf{x}$ (and its paired predicate-belief trajectory $\mathbf{z}$), executes a short prefix of $\mathbf{x}$, then re-observes and re-samples.

5.4 Planning-based Skill Composition with Predicate Beliefs

A key benefit of modeling skills as joint action–outcome rollouts is that the generated predicate-belief trajectory provides a symbolic interface for composition. This interface is compatible with formal

Algorithm 1 Planning with Predicted Predicate Beliefs

Require: Skill library $\mathcal{S}$ with PACTS $\{p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}, s)\}_{s \in \mathcal{S}}$; planning domain $\mathcal{D}$; PDDL problem instance $\mathcal{P}$ with initial predicate state $\hat{\mathbf{p}}_0$ and goal g

- 1: $t \leftarrow 0$; $\hat{\mathbf{p}}_t \leftarrow \hat{\mathbf{p}}_0$; observe $\mathbf{o}_t$
- 2: **while** goal g not satisfied under $\hat{\mathbf{p}}_t$ **do**
- 3: $\pi \leftarrow \text{PLAN}(\mathcal{D}, \hat{\mathbf{p}}_t, g)$ {symbolic plan in predicate space}
- 4: $s \leftarrow \pi[0]$ {next planned skill}
- 5: **while not** EFFECTSSATISFIED($s, \hat{\mathbf{p}}_t$) **and not** TERMINATESKILL(s) **do**
- 6: Sample rollout $(\mathbf{x}, \mathbf{z}) \sim p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t, s)$
- 7: Execute first action from $\mathbf{x}$ and observe $\mathbf{o}_{t+1}$
- 8: $\hat{\mathbf{p}}_{t+1} \leftarrow \text{THRESH}(\mathbf{z}_{t+1})$ {predicted next predicate state}
- 9: $t \leftarrow t + 1$
- 10: **end while**
- 11: **end while**

robot planning pipelines that represent tasks using symbolic or temporal-logic structure (Garrett et al., 2021; Paulius et al., 2025; Quartey et al., 2025). We assume a PDDL symbolic planning domain specifying skill-level operators with preconditions and effects over predicates, and a PDDL problem instance specifying an initial predicate state and a goal. Given a library of learned skills, an off-the-shelf planner synthesizes a sequence of skills to achieve a goal. During execution, PACTS produces online predicate-belief estimates that serve as the planner state, enabling the system to (i) check whether a skill’s intended effects are satisfied and (ii) replan from an updated symbolic state.

Predicate beliefs as an execution interface At each decision point, the robot observes $\mathbf{o}_t$ and (for the selected skill s) samples a joint rollout $(\mathbf{x}, \mathbf{z}) \sim p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t, s)$. The belief trajectory $\mathbf{z} \in [0, 1]^{H \times J}$ provides a compact prediction of how symbolic conditions are expected to evolve as the skill unfolds. The system obtains a discrete predicate valuation by thresholding beliefs (e.g., $\mathbf{z}_{t+h,j} > 0.5$), and uses the predicted next-step predicates $\hat{\mathbf{p}}_{t+1} = \text{THRESH}(\mathbf{z}_{t+1})$ as the symbolic state for effect checking and replanning. Although $\hat{\mathbf{p}}_{t+1}$ is obtained from the model’s prediction $\mathbf{z}_{t+1}$ (made when sampling at time t , before observing $\mathbf{o}_{t+1}$), beliefs are re-grounded at every step because the next rollout is conditioned on the latest observation $\mathbf{o}_{t+1}$.

Plan–execute–check–replan Algorithm 1 summarizes our composition loop. The planner replans after each executed skill, using the updated predicted predicate state. It selects the next skill from this symbolic state, and execution proceeds until the predicted predicates satisfy the skill’s effect conditions or another termination condition is met.

5.4.1 Skill segmentation and labeling toolkit.

To make real-world deployment tractable, we introduce a skill segmentation and labeling toolkit that converts monolithic demonstrations into training-ready, skill-centric datasets for joint $(\mathbf{x}, \mathbf{z})$ learning. The toolkit operates natively on LeRobot (Cadene et al., 2024) episodic datasets (state/action logs with video observations) and provides three core capabilities: (i) sparse keyframe annotation of




Figure 5.3: **Evaluation environments.** We evaluate PACTS across (a) **PushBarrier**, a controlled 2D compositional environment for systematic ablations and planner-based skill sequencing; (b) **Kitchen** and (c) **Coffee Preparation** from RoboMimic/MimicGen, which test joint action–predicate modeling under realistic 3D visual manipulation with simulator predicate oracles; and (d) a **real-world cube packing** task, where a robot packs four colored cubes into a container.

grounded predicates defined directly from the PDDL domain/problem files, (ii) automatic conversion of sparse annotations into dense per-timestep predicate labels via forward-fill interpolation, and (iii) skill-level segmentation by identifying goal-achieving predicate configurations and backward-filling executing phases to extract contiguous per-skill segments.

5.5 Evaluation

We evaluate PACTS across three tasks in simulation and a real-world multi-step task. Our experiments are designed to answer three questions:

1. How does joint predicate–action generation affect action performance and predicate classification accuracy?
2. How do modeling choices—DDPM denoising versus conditional flow matching, and UNet versus Transformer backbones—impact performance?
3. Can the learned predicate-belief interface support planning-based recomposition and replanning-based recovery on novel goals without retraining?

5.5.1 Tasks and datasets

We evaluate PACTS in three settings that trade off control, realism, and compositional complexity (see Fig. 5.3).

2D PushBarrier (compositional benchmark). We introduce a multi-step 2D manipulation task adapted from the Push-T environment (Chi et al., 2025). The workspace contains two objects (a T block and an L block) that must each be pushed from their respective start regions to corresponding goal regions. Successful completion requires not only reaching the goal regions but also achieving the correct final *orientation/alignment* of each object with its goal marker. A barrier separates the start and goal regions, and the robot must first press a button to open the barrier before either object can be

transferred to the goal side. This environment provides a controlled setting with clear compositional structure and is our primary benchmark for systematic evaluations and ablations.

We use PushBarrier in two complementary evaluation modes. To address Questions (1) and (2) in an end-to-end setting, we train *single* long-horizon policies on the full task and compare task performance and predicate classification accuracy against our baselines. To address Question (3) in a compositional setting, we construct a library of short-horizon skills and a task-relevant predicate set, segment demonstrations into per-skill datasets, and train a separate PACTS model for each skill. We then evaluate planning-based compositions that require sequencing skills, including *novel skill sequences* not observed during training.

RoboMimic (3D manipulation). To evaluate under realistic visual and geometric complexity, we additionally evaluate on two RoboMimic (Mandlekar et al., 2022) tasks with MimicGen (Mandlekar et al., 2023) demonstration data: **Kitchen** and **Coffee Preparation**. These tasks involve long-horizon manipulation under visually grounded observations and serve as a complementary test of whether joint modeling improves or maintains performance beyond the controlled 2D benchmark. For each task, we define a predicate set aligned with task structure and implement predicate oracles inside the simulator that compute ground-truth predicate values from simulator state. These predicate labels are used for both training and evaluation, enabling direct comparisons of predicate classification accuracy across methods. We use the base-difficulty (d0) variants, converted to absolute action space, with 200 demonstrations per task: `kitchen_d0` and `coffee_preparation_d0`.

Real-world demonstration We include a qualitative real-world demonstration on a cube-packing task. A tabletop robot must pick and place four colored cubes from fixed start regions into a container. We collect *monolithic*, long-horizon demonstrations that complete the full task (packing all four cubes), then use our segmentation and labeling toolkit to decompose each episode into skill-centric sub-demonstrations corresponding to per-cube packing behaviors (e.g., `pack_red`, `pack_green`, etc.). We train a separate PACTS model for each subskill and compose them with an off-the-shelf PDDL planner. This enables recompositions beyond the original demonstrations, such as executing individual packing skills in isolation or packing only a specified subset of cubes (e.g., “pack red and yellow”), providing a concrete test of zero-shot compositionality in the real world.

5.5.2 Baselines and ablations

We design baselines to isolate the benefit of modeling the joint action–outcome distribution $p_{\theta}(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t)$ and sampling paired action–predicate rollouts in PACTS. We begin with baselines that model each conditional in isolation (action-only and predicate-only). We then include a shared-encoder multi-task baseline that predicts actions and predicates with separate heads trained end-to-end, which isolates the effect of *joint sampling* from standard multi-task representation sharing.

Action-only generative policy (Fig. 5.2(c)). A diffusion/flow policy that models only $p(\mathbf{x} \mid \mathbf{o}_t)$ by generating an action trajectory $\mathbf{x}$ without predicate channels. This represents the standard action-only generative-policy baseline and tests whether adding outcome modeling degrades (or improves) action quality.

Standalone predicate predictor (Fig. 5.2(d)). A predicate-only model that predicts predicate beliefs from observations, modeling $p(\mathbf{z} \mid \mathbf{o}_t)$. This quantifies predicate prediction performance in isolation and mirrors common pipelines where predicates are treated purely as a discriminative output.

Multi-task baseline (Fig. 5.2(b)). To disentangle joint modeling from multi-task learning, we include a baseline with a shared observation encoder and two output heads trained end-to-end with a combined loss. This corresponds to learning a factorized conditional model

$$p(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t) \approx p(\mathbf{x} \mid \mathbf{o}_t) p(\mathbf{z} \mid \mathbf{o}_t), \quad (5.4)$$

where the action head is a conditional diffusion/flow policy trained to model $p(\mathbf{x} \mid \mathbf{o}_t)$ (using the same DDPM/CFM objectives as Sec. 5.3.2–5.3.3), and the predicate head is a discriminative predictor trained with binary cross-entropy on ground-truth predicates. The combined objective is

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{action}} + \mathcal{L}_{\text{predicate}}. \quad (5.5)$$

Although the encoder is shared, the two heads produce outputs independently at inference, so this baseline cannot encourage action–predicate consistency at sampling time.

Objective and backbone sweeps. All methods (PACTS and baselines) are evaluated across both training objectives (DDPM and CFM) and both backbone choices (UNet and Transformer), so differences can be attributed to joint modeling rather than a particular generative objective or architecture.

5.5.3 Metrics

We report the following metrics.

Policy evaluation We evaluate methods by rolling out policies in the environment and reporting task performance. In PushBarrier we report *Max Reward* ($\uparrow$) and *Steps to Max Reward* ($\downarrow$). In RoboMimic, we follow the benchmark protocol and report rollout success/reward statistics. Results are mean $\pm$ standard error over seven seeded rollouts.

Predicate classification evaluation We compare predicted predicates to ground-truth oracle labels and report macro *F1*, *recall*, *precision*, and *accuracy*. For belief outputs $\mathbf{z} \in [0, 1]^{H \times J}$, we threshold at 0.5 to obtain $\hat{\mathbf{p}} \in \{0, 1\}^J$. Predicate metrics are *macro-averaged across predicates* so that each predicate contributes equally, preventing frequent predicates from dominating the score.

Table 5.1: **Single-policy rollout evaluation on PushBarrier.** We compare PACTS to action-only, predicate-only, and shared-encoder multi-task baselines when learning a *single* policy for the full task. This isolates whether joint action–outcome modeling degrades action performance and whether it improves predicate prediction in the absence of any skill segmentation or planning. Policy metrics are mean±standard error over seven seeded rollouts

Setting (Objective + Backbone)	Approach	Policy Evaluation	Predicate Classification Evaluation				
		Max Reward $\uparrow$	F1 $\uparrow$	Recall $\uparrow$	Precision $\uparrow$	Accuracy $\uparrow$	Observations
DDPM + UNet	Action-Only Generator + Predicate-Only Predictor	0.83 $\pm$ 0.09	0.9476	0.9547	0.9471	79.40 %	2403
	Multi-task Baseline (Shared Encoder)	0.77 $\pm$ 0.13	0.9764	0.9712	0.9834	89.39 %	2705
	Predicate Action Skill (Pacts)	1.00 $\pm$ 0.00	0.9885	0.9924	0.9849	93.78 %	2234
DDPM + Transformer	Action-Only Generator + Predicate-Only Predictor	0.45 $\pm$ 0.10	0.9364	0.9355	0.9561	84.18 %	2800
	Multi-task Baseline (Shared Encoder)	0.33 $\pm$ 0.08	0.8684	0.8705	0.8679	85.43 %	2800
	Predicate Action Skill (Pacts)	0.58 $\pm$ 0.15	0.9872	0.9925	0.9822	91.98 %	2606
CFM + UNet	Action-Only Generator + Predicate-Only Predictor	0.31 $\pm$ 0.11	0.8297	0.7985	0.8775	82.11 %	2800
	Multi-task Baseline (Shared Encoder)	0.45 $\pm$ 0.11	0.9551	0.9595	0.9559	84.88 %	2784
	Predicate Action Skill (Pacts)	0.56 $\pm$ 0.12	0.9819	0.9945	0.9712	92.18 %	2800
CFM + Transformer	Action-Only Generator + Predicate-Only Predictor	0.62 $\pm$ 0.12	0.8969	0.9194	0.8869	75.15 %	2612
	Multi-task Baseline (Shared Encoder)	0.45 $\pm$ 0.12	0.8914	0.9387	0.8573	81.43 %	2800
	Predicate Action Skill (Pacts)	0.72 $\pm$ 0.12	0.9813	0.9932	0.9706	92.21 %	2759

Table 5.2: **Single-policy rollout evaluation on RoboMimic (Kitchen_d0 + Coffee_d0).** We evaluate PACTS and baselines on two RoboMimic manipulation tasks under realistic visual and geometric complexity, training a *single* policy per task (no planning or composition). Results are averaged across Kitchen and Coffee Preparation, weighting each environment equally. Policy metrics report mean±SE over seven seeded rollouts, with standard errors propagated across environments. Predicate classification metrics are macro-averaged across predicates within each environment, then averaged across environments.

Setting (Objective + Backbone)	Approach	Policy	Predicate Classification				
		Max Reward $\uparrow$	F1 $\uparrow$	Recall $\uparrow$	Precision $\uparrow$	Accuracy $\uparrow$	Observations
DDPM + UNet	Action-Only Generator + Predicate-Only Predictor	0.93 $\pm$ 0.02	0.9194	0.9095	0.9479	83.98 %	10476
	Predicate Action Skill (Pacts)	0.96 $\pm$ 0.02	0.9340	0.9271	0.9537	90.74 %	10246
DDPM + Transformer	Action-Only Generator + Predicate-Only Predictor	0.97 $\pm$ 0.02	0.9451	0.9360	0.9643	88.36 %	10087
	Predicate Action Skill (Pacts)	0.93 $\pm$ 0.04	0.9185	0.9050	0.9531	88.03 %	10103
CFM + UNet	Action-Only Generator + Predicate-Only Predictor	0.99 $\pm$ 0.01	0.9467	0.9407	0.9574	88.26 %	10243
	Predicate Action Skill (Pacts)	0.91 $\pm$ 0.03	0.8884	0.8824	0.9296	82.47 %	10543
CFM + Transformer	Action-Only Generator + Predicate-Only Predictor	0.97 $\pm$ 0.02	0.9324	0.9206	0.9612	86.07 %	9927
	Predicate Action Skill (Pacts)	0.89 $\pm$ 0.03	0.9386	0.9287	0.9587	89.78 %	10369

Predicate–action coherence To test whether predicted outcomes match the outcomes induced by executed actions, we compute horizon-indexed coherence curves over open-loop action chunks. At each prediction cycle τ , the policy predicts $\hat{\mathbf{z}}_{\tau+1:\tau+H}$ and executes its generated action chunk; we then compare $\hat{\mathbf{z}}_{\tau+h}$ to oracle predicates $\mathbf{p}_{\tau+h}$ for $h \in \{1, \dots, H\}$. We report CROSSCONS(h) (thresholded agreement, $\uparrow$), CROSSBRIER(h) (MSE on beliefs, $\downarrow$), and CROSSNLL(h) (BCE on beliefs, $\downarrow$), along with majority/base-rate reference baselines.

5.6 Results and Discussion

Single-policy analysis. We first study PACTS in the simplest setting: training a *single* long-horizon policy per task. This isolates the impact of jointly generating action–outcome rollouts on policy performance and predicate prediction independent of planning, or composition. The study tests whether joint learning can *add a planner-aligned outcome interface while preserving closed-loop control* and obtain an outcome trace that is (i) coherent under execution and (ii) suitable for downstream composition.



Figure 5.4: **Predicate–action coherence of predicted belief rollouts.** We measure how well PACTS’s predicted predicate-belief rollout matches the predicate outcomes realized after executing its open-loop action chunks. The x-axis reports lookahead h steps into the chunk. We plot (left) $CROSSCONS(h)$ (thresholded agreement; higher is better), (middle) $CROSSBRIER(h)$ (Brier score; lower is better), and (right) $CROSSNLL(h)$ (log loss; lower is better). Horizontal lines are dataset-statistics baselines: constant $\hat{p}=0.5$ (dotted) and a base-rate/majority predictor from label frequencies (dashed). These baselines ignore both observations and planned actions; outperforming them shows that the predicted rollouts carry information beyond fixed label frequencies. Across objectives and backbones, PACTS maintains high agreement and low losses over the horizon.




Figure 5.5: **Skill recomposition and adversarial recovery in PushBarrier.** Left to right: starting from the initial state, the planner synthesizes a sequence for a *novel* subgoal (`open_door` → `align.t.block`). After `open_door` succeeds, an adversarial perturbation (*Gremlin*) closes the door, which invalidates the current plan while aligning the *T* block. Using online predicate-belief estimates to check effects, the system triggers replanning, re-invokes `open_door`, and then resumes `align.t.block` to reach the goal.

PushBarrier (2D compositional benchmark). Table 5.1 reports single-policy performance and predicate classification on PushBarrier across objectives and backbones. Across all settings, PACTS achieves strong task reward while consistently improving predicate prediction and policy performance relative to baselines. Notably, the shared-encoder multi-task baseline often attains strong predicate scores, but its factorized heads yield weaker control; in contrast, PACTS samples *paired* action–predicate rollouts from $p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t)$ and remains competitive (often best) on reward. This supports our primary modeling claim: adding outcome channels does *not* require trading off control quality for symbolic prediction. Instead, learning the coupled object $p_\theta(\mathbf{x}, \mathbf{z} \mid \mathbf{o}_t)$ provides a useful inductive bias, encouraging the policy to represent *which* action modes correspond to *which* symbolic outcomes rather than learning actions and predicates as loosely related outputs.

Objective/backbone choices shift absolute numbers but not the conclusion: joint modeling remains strong across DDPM/CFM and UNet/Transformer. Consistent with prior work, UNet-style backbones are a reliable default, whereas Transformer variants can be more tuning-sensitive (Chi et al., 2025); in our evaluation, all models are trained for the same number of epochs without per-setting hyperparameter sweeps. We therefore view DDPM vs. CFM and UNet vs. Transformer primarily as implementation trade-offs (e.g., sampling speed and tuning sensitivity), with the primary differentiator being the *joint* modeling target.

RoboMimic (Kitchen_d0 + Coffee_d0). Table 5.2 reports single-policy evaluation averaged across both RoboMimic tasks. The action-only baseline remains strong on return across several objective/backbone choices, underscoring that explicit outcome modeling is not strictly necessary to achieve high reward when training end-to-end on a fixed task distribution. Our objective, however, is planning-based composition, which requires a *planner-aligned outcome interface* that stays coupled to the policy’s sampled actions and remains meaningful under execution—without sacrificing closed-loop control in realistic 3D manipulation. In this setting, PACTS remains competitive in return and frequently improves predicate quality, with the strongest overall trade-off in the DDPM+UNet configuration.

Predicate–action coherence. An outcome interface is only useful for composition if it remains *coupled* to what the policy actually does. Per-step predicate metrics alone (Tables 5.1–5.2) do not test whether predicted outcomes match what the model’s *own* generated actions induce. Figure 5.4 therefore evaluates predicate–action coherence: whether predicted belief rollouts agree with oracle predicate outcomes *under the model’s executed actions* across lookahead horizons. We report CROSSCONS(h) and two proper scoring rules, CROSSBRIER(h) and CROSSNLL(h), alongside constant and base-rate/majority reference baselines computed from dataset label statistics (i.e., priors). Across objectives and backbones, PACTS achieves near-perfect CROSSCONS and substantially lower CROSSBRIER/CROSSNLL than these reference baselines across the full horizon, indicating that predicted belief rollouts track outcomes realized under execution rather than reflecting label priors. This provides empirical support for using $\mathbf{z}$ as an outcome trace associated with the sampled action chunk: its predictions agree with predicate outcomes under execution.

Planning-based composition and recovery. We demonstrate planning-based recombination of learned skills and replanning-based recovery under perturbations using the predicate-belief interface. Figure 5.5 shows a PushBarrier run with two key properties. First, the goal is *novel*: it specifies only a subgoal of the full demonstrated behavior (e.g., achieve `align_t_block` without completing the entire monolithic task). Second, execution is perturbed adversarially: the door is closed mid-execution during `align_t_block`. By checking symbolic effects using updated predicate beliefs, the system detects that the current plan is no longer applicable, triggers replanning, and recomposes a corrected sequence (`open_door` $\rightarrow$ `align_t_block`) to reach the goal.

This illustrates two benefits of joint action–outcome skills: (i) *Recomposition*: skills expose planner-aligned predicates, allowing an off-the-shelf planner to synthesize sequences for goals not demonstrated end-to-end; (ii) *Recovery*: predicate beliefs make failures legible at the symbolic level, enabling replanning-based correction without retraining.

5.7 Limitations

PACTS inherits several limitations from predicate and skill learning. First, performance depends on the choice and coverage of the predicate set: incomplete or mis-specified predicates that are misaligned with task structure can make planning brittle and limit the expressivity of the outcome interface. Second, while joint modeling couples action and outcome prediction, it does not eliminate errors from perception and state aliasing—especially under occlusions and long horizons—and execution can still fail under stochasticity, partial observability, or distribution shift. Finally, our composition results assume a manually defined symbolic domain (e.g., PDDL operators and goals); automatically discovering predicates, operators, or abstractions remains an open research direction. A promising avenue is to leverage pretrained vision–language models to propose and evaluate candidate predicates directly from demonstrations (Athalye et al., 2026), then distill a compact planner-compatible set for downstream planning.

5.8 From Composable Skills to Scaling Trusted Autonomy

PACTS connect learned control to the symbolic conditions through which a planner composes behavior. Joint generation makes the action trajectory and its predicted predicate trace parts of the same skill model. The evaluations test control performance, agreement between predicted and executed outcomes, and composition for goals without end-to-end demonstrations of the requested sequence. The contribution is a way to acquire control and state reporting together within an explicit symbolic planning domain.

Together, the technical chapters connect grounded intent, an evolving world representation, and learned behavior to symbolic reasoning. The concluding Chapter 6 develops their shared principle of constructing and maintaining task-relevant abstractions, and examines how these interfaces can support continued learning, observation, and recovery in a physical closed-loop system.

Chapter 6

Scaling Trusted Autonomy

The ambition of generally capable robotics is to make useful behavior possible without anticipating every task or engineering every response in advance. As the range of requests, environments, and learned abilities grows, so does the difficulty of establishing why a robot’s behavior should be relied upon. A system may understand more language while leaving its interpretation difficult to examine, acquire more skills while obscuring their effects, or construct a detailed scene whose contents become outdated through its own actions. The central challenge is to develop the means of assessing behavior alongside the capability to produce it.

This thesis has approached that challenge through the connections between learned models and explicit reasoning. The preceding chapters develop grounded task specifications, persistent spatial representations, and learned skills with symbolic state reports. Together, they support a broader position: the structure required for planning need not all be fixed before the robot encounters a task. Some of that structure can be constructed from language and observations when it becomes relevant; some can be learned with behavior; and some must be revised as the world changes. Understanding how to construct and maintain these connections is central to scaling trusted autonomy.

6.1 From Greater Capability to Trusted Autonomy

The introduction established two requirements for robots that operate in human environments. People should be able to specify useful goals without prescribing each movement, and they should have a basis for assessing whether the resulting behavior satisfies their requests. Foundation models make a much broader range of language and perceptual knowledge available to the robot. Symbolic planning provides a means of expressing requirements and reasoning over combinations of actions. Their complementary strengths become useful together when the robot can connect the particular request and environment to the representations on which planning depends.

The contributions develop this connection along three dimensions of capability. LIMP expands the instructions that can be grounded into explicit temporal and spatial requirements without a complete prebuilt semantic map. REALM carries a persistent spatial state across planned changes, allowing

subsequent instructions to obtain their grounding and planning abstractions from an evolved world. PACTS connect learned control to predicted symbolic conditions, enabling a planner to compose trained skills for requested goals without demonstrations of the entire sequence. These are different forms of expansion: across descriptions, across successive world states, and across combinations of behavior. They make the question of scaling concrete without equating it with a measured scaling law or unrestricted competence.

Each expansion also exposes something that the robot must be able to assess. An expressive interpretation must identify the intended task and referents. A skill sequence must have effects that support the planner’s next decision. A persistent representation must distinguish a predicted change from one supported by observation. The scientific value of the three interfaces lies in making these dependencies explicit. They identify what connects a plan that is correct within its model to a task that is correctly understood and carried out in the world.

This perspective changes where the burden of engineering lies. A planning system needs more than an algorithm for searching a supplied domain; it needs a way to obtain and maintain the domain information that makes its reasoning applicable. The thesis develops parts of that process through learning, grounding, and world evolution. Its unifying question is therefore how a robot can acquire the structure needed for a particular decision while retaining enough information to reconsider that structure when circumstances change.

6.2 Constructing the Abstractions a Task Requires

The world contains far more detail than a planner can use directly. Even a familiar room admits many descriptions: objects have appearances, spatial relationships, possible uses, and properties that matter only for particular activities. An abstraction makes planning tractable by retaining distinctions relevant to a decision and setting others aside. The difficulty is deciding those distinctions before knowing what the robot will be asked to do. A representation sufficient for one instruction may omit the very attribute or relationship needed to understand the next.

The approach developed in this thesis is to defer task-specific representational commitments until the task provides a reason to make them. The robot retains observations or a richer spatial representation from which it can construct the semantic and geometric descriptions required for grounding and planning. The resulting model is bounded and explicit, but it need not be the only model that the robot can construct from the available information. This is the sense in which *delaying closure* offers a path toward planning in open-world settings: the robot constructs a model for the current task without treating it as a complete description of everything it might later need to know.

LIMP makes this principle concrete at the interface between language and the mapped environment. An instruction determines which referent descriptions and spatial relationships the system resolves using retained observations. The semantic map is consequently shaped by the task rather than limited to an exhaustive catalogue of named landmarks prepared beforehand. Foundation models supply the breadth of language and visual concepts used in this construction, while explicit task

structure and spatial reasoning connect their outputs to the planner. The important combination is semantic flexibility with a representation whose interpretation and consequences can be examined.

REALM extends the principle through time. If the underlying scene is fixed, newly constructed abstractions can still describe an obsolete world. REALM therefore maintains a persistent Gaussian scene with object identity and an explicit history of action-conditioned changes. Grounded objects, planning geometry, collision representations, and rendered views are derived from the state associated with the current task. Their continuity comes from that common state, while their form depends on the question they must answer. The predictive transitions and subsequent use for decisions give the scene its world-model function; the Gaussian representation alone does not provide it.

PACTS addresses a complementary source of rigidity: the connection between a learned controller and the symbolic conditions used to compose it. Its action trajectories and predicate-belief trajectories are learned jointly, so the skill exposes predicted state in a form the planner can use. The predicate vocabulary and operator model are supplied, and the constituent skills are trained before their composition. PACTS therefore does not demonstrate runtime discovery of the planning domain. It shows how part of the interface that would otherwise require separate engineering can be acquired with behavior, making those learned abilities available for new combinations at runtime.

Deferring these commitments does not remove the need for prior structure. LIMP depends on captured observations, spatial representations, and available skills; REALM depends on a reconstructed scene and specified transition mechanisms; PACTS depends on demonstrations and a symbolic domain. Nor does constructing a model at runtime change the logical semantics of the planner that consumes it. A finite or fixed map is not inherently closed-world; that assumption enters when what it contains is treated as a complete account of the relevant world. The broader contribution concerns when and how the robot obtains such an account, and how it can replace it as new tasks or evidence make different distinctions necessary.

This principle also reveals a further requirement. Information is relevant to a task when it can affect the task’s interpretation, feasibility, or outcome, whether or not the instruction mentions it. The robot may need to account for an obstacle, an unstable support, or a property that becomes apparent only during interaction. Retaining richer observations creates an opportunity to recover omitted information, but does not ensure that it was observed or can be recovered correctly. Trusted runtime abstraction therefore requires a way to assess whether the constructed model is adequate for the decision it supports. The trust interfaces make that requirement more precise by locating the connections through which an inadequate model can mislead the robot.

6.3 Trust Across Intent, Action, and World

The three questions introduced at the beginning of the dissertation can now be considered in light of the technical results. Did the robot understand what the user meant? Did its actions produce the intended effects? Does its internal world still match reality? Figure 6.1 brings the corresponding contributions together in this conceptual order. The chapter sequence followed the assumptions exposed by LIMP; returning to Intent, Action, and World shows how those contributions bear on the successive decisions involved in carrying out a request.




Figure 6.1: The three trust interfaces and their complementary contributions. LIMP exposes a grounded task and motion plan; PACTS jointly generates actions and predicted predicate trajectories; REALM supplies an evolving spatial representation for subsequent tasks.

At the intent interface, assurance begins with making the interpreted request explicit. LIMP represents temporal requirements in a co-safe task specification and connects its propositions to spatial referents and feasible motion. Its evaluation across 150 instructions in five real-world environments tests the combination of language interpretation, grounding, and planning. The results in Tables 3.1 and 3.2 show the benefit of this structure particularly for complex temporal requirements and avoidance constraints. The weaker performance on instructions with implicit preferences is also informative: a formal specification can expose and preserve an interpretation without ensuring that it is the interpretation the user intended.

This distinction gives human inspection a clear role. LIMP presents its interpreted task structure for review and supports reprompting after a user’s correction; the reported translation evaluation measures the initial output without that correction. Inspection offers a means of identifying errors before the planner acts on them. It is distinct from checking whether a plan satisfies the resulting formal specification. The latter establishes a relationship between a model and a plan, conditional on the interpretation, grounding, and action assumptions that made the model meaningful.

At the action interface, a planner needs to know how execution bears on its symbolic state. PACTS learns to generate a proposed action trajectory together with the predicate conditions expected along it. The composition and perturbation examples show how this interface supports new sequences of

trained skills and replanning when updated beliefs invalidate a plan. The action–predicate evaluation in Chapter 5 goes beyond measuring predicate classification alone: it compares predicted beliefs with outcomes under executed action chunks. This directly tests the connection that matters for composition, within the evaluated tasks and horizons.

The same distinction between representation and evidence applies here. Joint sampling makes actions and predicate predictions parts of one learned model, but the architecture does not prove that they agree with physical outcomes. The reported state is refreshed through observation-conditioned rollouts, yet predicted future predicates are not independent measurements of successful execution. The contribution is an interface through which learned behavior can participate in symbolic reasoning and through which its predicted effects can be evaluated. Reliable use depends on how well those predictions track the behavior and observations encountered during operation.

At the world interface, the question extends beyond the current skill to the state on which later tasks will depend. REALM’s two-instruction study shows that a persistent representation can carry object identity and planned change forward while regenerating the next task’s grounding and planning geometry. Object-local transitions preserve unaffected Gaussian content, and the recorded history permits exact recovery of ancestor states. The representation and image-consistency results in Section 4.6 establish specific properties of that update process.

That boundary is essential to the world interface’s role in continued autonomy. A planned endpoint predicts where an object should be if the action succeeds. A consistent rendering shows what the resulting representation implies. An observation supplies evidence about what the robot actually changed. Keeping these roles distinct allows an evolving world model to support scrutiny rather than merely accumulate plausible predictions. REALM provides an explicit structure for carrying changes and their provenance forward; physical observations would have to test those changes during deployment.

Taken together, the results provide different kinds of evidence at complementary interfaces. Specification satisfaction, empirical agreement between predicted and executed outcomes, and consistency of world-state updates answer different questions. Their value comes from connecting them without allowing one to stand in for another. A plan can satisfy the wrong specification, accurate skill-level predictions can be grounded in the wrong object, and an internally consistent world can depart from reality. The thesis makes these dependencies available for examination and develops mechanisms for several of them. Establishing assurance across their composition is the next scientific step.

6.4 Toward Continued Autonomy

The contributions suggest a robot whose understanding, behavior, and world representation can develop together over successive tasks. An instruction would determine the abstractions needed for planning; learned skills would expose the conditions relevant to their execution; and subsequent observations would inform the state used for the next request. Realizing that loop poses questions about how the components should share evidence, recognize inadequacy, and change their decisions. The following directions arise from those questions and from the assumptions of the demonstrated systems.

6.4.1 Recognizing When an Abstraction Is Insufficient

Runtime construction makes it possible to build different models for different tasks. The next challenge is to determine when the chosen model leaves out something consequential. Suppose an instruction identifies a book to be moved from a table to a sofa. Correctly identifying those objects may be sufficient for understanding the request, while planning and execution also depend on objects obstructing the grasp or the stability of the intended placement. A description adequate at the intent interface may therefore be insufficient at the action interface. The robot needs a way to discover and incorporate distinctions as they become relevant across these uses.

One research direction is to connect planning and execution failures to targeted expansion of the abstraction. Repeatedly failing to find a feasible grasp could prompt the robot to inspect nearby geometry; an unexpected effect could reveal a property absent from its action model. Foundation models could propose additional descriptions or relationships, while geometric checks and observations would test whether those proposals explain the difficulty. The key question is how to select an informative revision without rebuilding an exhaustive model of the environment. Success would involve both improved task decisions and a justified account of why the new information was required.

This question also applies to the meaning of absent information. Failure to observe a relevant object or condition should be distinguished from evidence that it is absent. A system using bounded classical planning models could manage that distinction around the planner, acquiring information or revising the domain before making a consequential commitment. Evaluating this ability requires tasks with deliberately omitted but relevant entities, hidden conditions, and changed instructions. Such tests would measure whether the robot recognizes the limits of its abstraction, rather than only whether it solves tasks already expressible within it.

6.4.2 Connecting Predictions to Observations

Continued autonomy depends on deciding which predicted changes deserve to become part of the robot’s working account of the world. The intended effect of a skill provides a useful expectation: it identifies what might have moved and where new evidence may be needed. It also creates a risk if the system promotes that expectation without checking it. A failed grasp, slip, or unmodeled interaction could otherwise affect every later grounding and planning decision. The action and world interfaces must therefore connect predicted effects to observations that can confirm or contradict them.

PACTS and REALM provide complementary starting points for this connection. Predicate beliefs describe changes in conditions used by the planner, while a persistent spatial representation records objects and their geometric relationships. An integrated system could use a reported effect to direct re-observation and then compare the observed scene with the predicted transition. Agreement would support reuse of the updated state; disagreement would require reconsidering the skill outcome, the object association, or the world update. The challenge is to relate symbolic and geometric evidence without assuming that either is an infallible authority.

Observation should itself become a purposeful decision. The robot need not reconstruct its entire environment after every action, but it needs evidence about changes on which later decisions depend.

Selecting viewpoints, inspecting an occluded object, or asking a person to resolve an ambiguous referent can each reduce a different source of uncertainty. An account of the value of such information should consider both its cost and the consequences of acting without it. This would extend run-time abstraction from deciding what to represent to deciding what must be observed before that representation can be relied upon.

6.4.3 Expanding Skills and Their Symbolic Meaning

Composition extends the usefulness of existing skills, but its reach is bounded by the abilities and conditions available to the planner. PACTS assumes a supplied predicate vocabulary and symbolic operators. Extending this work toward continued learning requires a way to introduce new behaviors together with the descriptions that make them usable alongside existing skills. A controller alone cannot tell the planner when the skill applies, which conditions it may establish, or how those conditions relate to the rest of the domain.

Demonstrations and interactions could provide evidence for learning these relationships. A pre-trained model might propose candidate predicates or describe a skill’s intended effect, but those descriptions would need to be grounded and tested against execution. Different controllers that claim to establish the same condition should support a consistent interpretation of that condition. Conversely, persistent failures may indicate that an operator omits a prerequisite or that one predicate conflates states requiring different behavior. Learning the interface would therefore involve revising symbolic distinctions as well as fitting better action and state predictions.

The central evaluation question is whether each addition expands the planner’s useful choices while preserving the meaning of existing ones. This calls for testing both new compositions and previously achievable tasks after the domain changes. It also requires relating predicate predictions to the actions actually executed, especially when new skills alter observation distributions or introduce unfamiliar interactions. PACTS’s coherence evaluation provides one starting point for this assessment; broader coverage, confidence calibration, and validation under distribution shift would be needed to support sustained expansion.

6.4.4 Managing Uncertainty Across the Interfaces

Uncertainty has different meanings at different points in the autonomy loop. The robot may be unsure which object a phrase refers to, whether a skill will establish a condition, or whether a scene update reflects a physical change. A single confidence score can obscure those differences. For uncertainty to inform behavior, the system needs to associate it with a particular proposition or prediction and identify what decision depends on it.

The interfaces also share sources of error. A mistaken object association can influence the task specification, skill input, and world update simultaneously. Agreement among these components may then reflect their common mistake. Multiplying independently reported confidences would be unjustified without accounting for such dependencies, and a predicate-belief score is not automatically

a calibrated probability. The research problem is to preserve enough information about the origin and meaning of uncertainty to determine which additional evidence could resolve it.

Such an account would make recovery more selective. Ambiguity in intent may call for clarification, uncertainty about an effect may call for observation, and an infeasible plan may call for a different skill or abstraction. The robot should also be able to recognize when none of its available actions has adequate support and seek assistance. Evaluations should examine whether these decisions improve task outcomes and reduce unsupported commitments, alongside the time, observations, and human intervention they require. This would make assurance part of the robot’s decision process rather than a judgment applied only after execution.

6.5 Concluding Perspective

Generally capable robots will encounter tasks and circumstances that cannot all be described in advance. The work in this dissertation shows how planning can draw on learned semantic knowledge, persistent spatial representations, and learned behavior to address parts of that variability. Its central position is that *scaling trusted autonomy requires unifying grounded language understanding, composable state-reporting skills, and evolving representations of the world*. Each contribution makes explicit a connection needed to interpret a request, reason about behavior, or carry useful state into the next task.

The broader lesson concerns how that structure is obtained. By deferring task-specific commitments, constructing relevant abstractions when they are needed, and maintaining the information from which they are derived, a robot can extend the reach of explicit planning beyond a single world description fixed in advance. Learning can help supply the missing semantics and behavior; reasoning can make their consequences available for inspection and testing. Continued autonomy will depend on carrying this relationship further, so that the robot can recognize when its abstractions are inadequate and use experience to revise them. Progress toward generally capable, verifiable robots requires both a growing range of useful behavior and a growing ability to establish when that behavior can be relied upon.

Appendix A

Supplementary Methods for LIMP

A.1 Overview

This appendix supplements Chapter 3 with prompting details, spatial comparators, planning algorithms, robot skill implementations, and evaluation materials for LIMP. The authorial “we” refers to the collaboration credited at the beginning of that chapter. Prompts and evaluation instructions are reproduced as used in the study; project demonstrations are available at robotlimp.github.io.

A.2 Language Instruction Module

We implement a two-stage prompting strategy in our language instruction module to translate natural language instructions into LTL specifications. The first stage translates a given instruction into a conventional LTL formula, where propositions refer to open-vocabulary objects. For any given instruction, we dynamically generate K in-context translation examples from a standard dataset (Pan et al., 2023) of natural language and LTL pairs, based on cosine similarity with the given instruction. Here is the exact text prompt used:

Listing A.1: Base prompt used to obtain a conventional LTL formula from a natural language query

```
1 You are a LLM that understands operators involved with Linear Temporal Logic (LTL),  
   such as F, G, U, &, |, ~, etc. Your goal is translate language input to LTL  
   output.  
2 Input:<generated_example_instruction>  
3 Output:<generated_example_LTL>  
4 ...  
5 Input:<given_instruction>  
6 Output:
```

The second stage takes the given instruction and the LTL response from the first stage as input to generate a new LTL formula with predicate functions that correspond to parameterized robot skills. Skill parameters are instruction referent objects expressed in our novel Composable Referent Descriptor (CRD) syntax. CRDs enable referent disambiguation by chaining comparators that encode

descriptive spatial information. We define eight spatial comparators and provide their descriptions as part of the second stage prompt. The prompt supplies descriptions of the supported comparators and examples of how they can be composed. Here is the exact prompt used:

Listing A.2: Second stage prompt to output our LTL syntax with CRD parameterized robot skills

```

1 You are an LLM for robot planning that understands operators involved with Linear
  Temporal Logic (LTL), such as F, G, U, &, |, ~, etc. You have a finite set of
  robot predicates and spatial predicates, given a language instruction and an LTL
  formula that represents the given instruction, your goal is to translate the ltl
  formula into one that uses appropriate composition of robot and spatial predicates
  in place of propositions with relevant details from original instruction as
  arguments.
2 Robot predicate set (near,pick,release).
3 Usage:
4 near[referent_1]:returns true if the desired spatial relationship is for robot to be
  near referent_1.
5 pick[referent_1]:can only execute picking skill on referent_1 and return True when
  near[referent_1].
6 release[referent_1,referent_2]:can only execute release skill on referent_1 and return
  True when near[referent_2].
7 Spatial predicate set (isbetween,isabove,isbelow,isleftof,isrightof,isnextto,
  isinfrontof,isbehind).
8 Usage:
9 referent_1::isbetween(referent_2,referent_3):returns true if referent_1 is between
  referent_2 and referent_3.
10 referent_1::isabove(referent_2):returns True if referent_1 is above referent_2.
11 referent_1::isbelow(referent_2):returns True if referent_1 is below referent_2.
12 referent_1::isleftof(referent_2):returns True if referent_1 is left of referent_2.
13 referent_1::isrightof(referent_2):returns True if referent_1 is right of referent_2.
14 referent_1::isnextto(referent_2):returns True if referent_1 is close to referent_2.
15 referent_1::isinfrontof(referent_2):returns True if referent_1 is in front of
  referent_2.
16 referent_1::isbehind(referent_2):returns True if referent_1 is behind referent_2.
17 Rules:
18 Strictly only use the finite set of robot and spatial predicates!
19 Strictly stick to the usage format!
20 Compose spatial predicates where necessary!
21 You are allowed to modify the structure of Input_ltl for the final Output if it does
  not match the intended Input_instruction!
22 You should strictly only stick to mentioned objects, however you are allowed to
  propose and include plausible objects if and only if not mentioned in instruction
  but required based on context of instruction!
23 Pay attention to instructions that require performing certain actions multiple times
  in generating and sequencing the predicates for the final Output formula!
24 Example:
25 Input_instruction: Go to the orange building but before that pass by the coffee shop,
  then go to the parking sign.
26 Input_ltl: F (coffee_shop & F (orange_building & F parking_sign ))
27 Output: F ( near[coffee_shop] & F ( near[orange_building] & F near[parking_sign] ))

```

```

28 Input_instruction: Go to the blue sofa then the laptop, after that bring me the brown
    bag between the television and the kettle on the left of the green seat, I am
    standing by the sink.
29 Input_ltl: F ( blue_sofa & F ( laptop & F ( brown_bag & F ( sink ) ) ) )
30 Output: F ( near[blue_sofa] & F ( near[laptop] & F ( near[brown_bag::isbetween(
    television,kettle::isleftof(green_seat))] & F ( pick[brown_bag::isbetween(
    television,kettle::isleftof(green_seat))] & F ( near[sink] & F ( release[brown_bag
    ,sink] ) ) ) ) ) )
31 Input_instruction: Hey need you to pass by chair between the sofa and bag, pick up the
    bag and go to the orange napkin on the right of the sofa.
32 Input_ltl: F ( chair & F ( bag & F ( orange_napkin ) ) )
33 Output: F ( near[chair::isbetween(sofa,bag)] & F ( near[bag] & F ( pick[bag] & F (
    near[orange_napkin::isrightof(sofa)] ) ) ) )
34 Input_instruction: Go to the chair between the green laptop and the yellow box
    underneath the play toy
35 Input_ltl: F ( green_laptop & F ( yellow_box & F ( play_toy & F ( chair ) ) ) )
36 Output: F ( near[chair::isbetween(green_laptop,yellow_box::isbelow(play_toy))] )
37 Input_instruction: Check the table behind the fridge and bring two beers to the couch
    one after the other
38 Input_ltl: F ( check_table & F ( bring_beer1 ) & F ( bring_beer2 ) & F ( couch ) )
39 Output: F ( near[table::isbehind(fridge)] & F ( pick[beer] & F ( near[couch] & F (
    release[beer,couch] & F ( near[table::isbehind(fridge)] & F ( pick[beer] & F (
    near[couch] & F ( release[beer,couch] ) ) ) ) ) ) ) ) ) ) )
40 Input_instruction: <given_instruction>
41 Input_ltl: <stage1_ltl_response>
42 Output:

```

A.2.1 Interactive Symbol Verification

Verifying sampled LTL formulas is essential, as such we implement an interactive dialog system that presents users with extracted referent composable referent descriptors (CRDs) in sampled formulas as well as the implied task structure—encoded in the sequence of state-machine transition expressions that must hold to progressively solve the task. We translate the task structure into English statements via a simple deterministic strategy that replaces logical connectives and skill predicates from the formula with equivalent English phrases. Users can verify a formula as correct or provide corrective statements which are used to reprompt the LLM to obtain new formulas. Below is the exact prompt used for regenerating formulas.

Listing A.3: Corrective reprompting prompt used to obtain new LTL formulas

```

1 There was a mistake with your output LTL formula: Error with <verification_type>!
    Consider the clarification feedback and regenerate the correct output for the
    Input_instruction. Make sure to adhere to all rules and instructions in your
    original prompt!
2 previous_output:<last_response>
3 error_clarification: <given_error_clarification>
4 correct_output:

```

To illustrate, the instruction “*Bring the green plush toy to the whiteboard in front of it*” yields the interactive Referent and Task Structure Verification dialog below:

Listing A.4: Interactive referent and task structure verification dialog.

```

1 *****
2 Instruction Following
3 *****
4 Input_instruction: "Bring the green plush toy to the whiteboard in front of it"
5 Sampled LTL formula: F(A & F(B & F(C & FD)))
6     A: near[green_plush_toy]
7     B: pick[green_plush_toy]
8     C: near[whiteboard::isinfrontof(green_plush_toy)]
9     D: release[green_plush_toy, whiteboard::isinfrontof(green_plush_toy)]
10
11 *****
12 Referent Verification
13 *****
14 I extracted this list of relevant objects based on your instruction:
15     * whiteboard::isinfrontof(green_plush_toy)
16     * green_plush_toy
17 Does this match your intention? (y/n)
18
19 *****
20 Task Structure Verification
21 *****
22 Based on my understanding here is the sequence of subgoal objectives needed to satisfy
    the task:
23 Subgoal_1:
24     Logical Expression: A&!B
25     English translation: I should be near the [green_plush_toy] and not have picked up
        the [green_plush_toy]
26 Subgoal_2:
27     Logical Expression: B&!C
28     English translation: I should have picked up the [green_plush_toy] and not be near
        the [whiteboard::isinfrontof(green_plush_toy)]
29 Subgoal_3:
30     Logical Expression: C&!D
31     English translation: I should be near the [whiteboard::isinfrontof(green_plush_toy
        )] and not have released the [green_plush_toy] at the [whiteboard::isinfrontof(
        green_plush_toy)]
32 Subgoal_4:
33     Logical Expression: D
34     English translation: I should have released the [green_plush_toy] at the [
        whiteboard::isinfrontof(green_plush_toy)]
35 Does this match your intention? (y/n)

```

A.3 Spatial Grounding Module

The spatial grounding module detects and localizes specific instances of objects referenced in a given instruction by first detecting, segmenting and back-projecting all referent occurrences and then filtering based on the descriptive spatial details captured by each referent’s composable referent descriptor (CRD). We use the Owl-Vit model (Minderer et al., 2022) to detect bounding boxes of open-vocabulary referents and SAM (Kirillov et al., 2023) to generate masks from detected bounding boxes. To illustrate referent filtering via spatial information, consider an example scenario where the goal is to resolve the composable referent descriptor below:

$$\text{whiteboard} :: \text{isinfrontof}(\text{green_plush_toy}). \quad (\text{A.1})$$

Let $W = \{w_1, w_2, \dots, w_n\}$ and $G = \{g_1, g_2, \dots, g_m\}$ represent the set of representative 3D positions of detected whiteboards and green_plush_toys respectively. The cartesian product of these sets enumerates all possible pairs (w, g) for comparison.

$$W \times G = \{(w, g) \mid w \in W, g \in G\} \quad (\text{A.2})$$

The ‘isinfrontof(w, g)’ comparator is applied to each pair, yielding a subset S that contains only those ‘whiteboards’ that satisfy the ‘isinfrontof’ condition with at least one ‘green_plush_toy’.

$$S = \{w \in W \mid \exists g \in G \text{ such that } \text{isinfrontof}(w, g) \text{ is true}\} \quad (\text{A.3})$$

A.3.1 3D Spatial Comparators

Our 3D spatial comparators enable Relative Frame of Reference (FoR) spatial reasoning between referents, based on their backprojected 3D positions. Threshold values in the spatial comparators give users the ability to specify the sensitivity or resolution at which spatial relationships are resolved, we keep all threshold values fixed across all experiments. Below is a description of each spatial comparator.

Listing A.5: Implementation description of 3D spatial comparators

1. `isbetween(referent_1_pos, referent_2_pos, referent_3_pos, threshold)`: Returns true if referent_1 is within ‘threshold’ distance from the line segment connecting referent_2 to referent_3, ensuring it lies in the directional path between them without extending beyond.
2. `isabove(referent_1_pos, referent_2_pos, threshold)`: Returns true if the z-coordinate of referent_1 exceeds that of referent_2 by at least ‘threshold’.
3. `isbelow(referent_1_pos, referent_2_pos, threshold)`: Returns true if the z-coordinate of referent_1 is less than that of referent_2 by more than ‘threshold’.
4. `isleftof(referent_1_pos, referent_2_pos, threshold)`: Returns true if the y-coordinate of referent_1 exceeds that of referent_2 by at least ‘threshold’, indicating referent_1 is to the left of referent_2.
5. `isrightof(referent_1_pos, referent_2_pos, threshold)`: Returns true if the y-coordinate of referent_1 is less than that of referent_2 by more than ‘threshold’, indicating referent_1 is to the right of referent_2.

6. `isnextto(referent_1_pos, referent_2_pos, threshold)`: Returns true if the Euclidean distance between `referent_1` and `referent_2` is less than `'threshold'`, indicating they are next to each other.
7. `isinfrontof(referent_1_pos, referent_2_pos, threshold)`: Returns true if the x-coordinate of `referent_1` is less than that of `referent_2` by more than `'threshold'`, indicating `referent_1` is in front of `referent_2`.
8. `isbehind(referent_1_pos, referent_2_pos, threshold)`: Returns true if the x-coordinate of `referent_1` exceeds that of `referent_2` by at least `'threshold'`, indicating `referent_1` is behind `referent_2`.

A.4 Task and Motion Planning Module

We present pseudocode for our Progressive Motion Planner (Algorithm A.1) and our algorithm for generating Task Progression Semantic Maps (Algorithm A.2). Algorithm A.2 generates a TPSM $\mathcal{M}_{\text{tpsm}}$ by integrating an environment map ($\mathcal{M}$) and a referent semantic map ($\mathcal{M}_{\text{rsm}}$) given a logical transition expression ($\mathcal{T}$), a desired automaton state ($\mathcal{S}'$), and a nearness threshold (θ). The algorithm first initializes $\mathcal{M}_{\text{tpsm}}$ with a copy of $\mathcal{M}$ and extracts relevant instruction predicates from $\mathcal{T}$. For each predicate (parameterized skill), the algorithm identifies satisfying referent positions in $\mathcal{M}_{\text{rsm}}$, generates a spherical grid of surrounding points within a radius θ , and assesses how these points affect the progression of the task automaton towards $\mathcal{S}'$. These points demarcate regions of interest, and are assigned a value of 1 if they cause the automaton to transition to the desired state, -1 if they lead to a different automaton state or violate the automaton, and 0 if they do not affect the automaton. The points are then integrated into $\mathcal{M}_{\text{tpsm}}$, yielding a semantic map that identifies goal and constraint violating regions.

Algorithm A.1 Progressive Motion Planning Algorithm

```
1: procedure PMP( $X_{start}, \varphi, \mathcal{M}, \mathcal{M}_{rsm}, \theta$ )
  Input:
     $X_{start}$ : Start position in the environment.
     $\varphi$ : CRD syntax LTL formula specifying task objectives.
     $\mathcal{M}$ : Environment map.
     $\mathcal{M}_{rsm}$ : Referent semantic map.
     $\theta$ : Nearness threshold.
  Output:
     $\Pi$ : Generated task and motion plan.
2:  $\mathcal{A} \leftarrow \text{ConstructAutomaton}(\varphi)$ 
3:  $\text{path} \leftarrow \text{SelectAutomatonPath}(\mathcal{A})$  {Task plan}
4: while  $\Pi.\text{status}$  is active do
5:   while  $\mathcal{A}.\text{state} \neq \text{path}.\text{acceptingState}$  do
6:      $\mathcal{S}, \mathcal{T}, \mathcal{S}' \leftarrow \mathcal{A}.\text{GetTransition}(\text{path}.\text{currentStep})$ 
7:      $\text{objective} \leftarrow \text{NextObjectiveType}(\mathcal{T})$ 
8:     if  $\text{objective} = \text{"skill"}$  then
9:        $\Pi.\text{UpdateWithSkill}(\mathcal{T})$ 
10:       $\mathcal{A}.\text{UpdateAutomatonState}(\mathcal{S}')$ 
11:     else if  $\text{objective} = \text{"navigation"}$  then
12:        $\mathcal{M}_{params} \leftarrow \mathcal{M}, \mathcal{M}_{rsm}, \mathcal{T}, \mathcal{A}, \mathcal{S}', \theta$ 
13:        $\mathcal{M}_{tpsm} \leftarrow \text{GENERATETPSM}(\mathcal{M}_{params})$ 
14:        $\mathcal{O} \leftarrow \text{GenerateObstacleMap}(\mathcal{M}_{tpsm})$ 
15:        $\text{plan} \leftarrow \text{FMT}^*(X_{start}, \mathcal{O})$  {Path plan}
16:       if  $\text{plan}.\text{exists}$  then
17:          $\Pi.\text{UpdateWithPlan}(\text{plan})$ 
18:          $X_{start} \leftarrow \text{plan}.\text{endPosition}$ 
19:          $\mathcal{A}.\text{UpdateAutomatonState}(\mathcal{S}')$ 
20:       else
21:          $\Pi, \mathcal{A}, \text{path} \leftarrow \text{Backtrack}(\Pi, \mathcal{A}, \text{path})$ 
22:       end if
23:     end if
24:   end while
25: end while
26: return  $\Pi$ 
27: end procedure
```

Algorithm A.2 Task Progression Semantic Mapping Algorithm

```

1: procedure GENERATETPSM( $\mathcal{M}, \mathcal{M}_{rsm}, \mathcal{T}, \mathcal{A}, \mathcal{S}', \theta$ )
  Input:
     $\mathcal{M}$ : Environment map.
     $\mathcal{M}_{rsm}$ : Referent semantic map.
     $\mathcal{T}$ : Automaton transition expression.
     $\mathcal{A}$ : Task Automaton.
     $\mathcal{S}'$ : Desired State.
     $\theta$ : Nearness threshold.
  Output:
     $\mathcal{M}_{tpsm}$ : Task Progression Semantic Map.
2:  $\mathcal{M}_{tpsm} \leftarrow \text{Copy}(\mathcal{M})$ 
3:  $\mathcal{P} \leftarrow \text{ExtractRelevantPredicates}(\mathcal{T})$ 
4: for  $p$  in  $\mathcal{P}$  do
5:    $\mathcal{R} \leftarrow \text{QueryPositions}(\mathcal{M}_{rsm}, p)$ 
6:   for  $r$  in  $\mathcal{R}$  do
7:      $G \leftarrow \{g \mid g = r + \delta, \|\delta\| \leq \theta\}$  {spherical grid of surrounding points}
8:     for  $g$  in  $G$  do
9:        $\mathcal{Q} \leftarrow \text{TruePredicatesAt}(g, \mathcal{M}_{rsm}, \theta)$ 
10:       $\mathcal{S}_{next} \leftarrow \text{ProgressAutomaton}(\mathcal{A}, \mathcal{Q})$ 
11:      if  $\mathcal{S}_{next} = \mathcal{S}'$  then
12:         $g.value \leftarrow 1$  {Goal value}
13:      else if  $\text{IsUndesired}(\mathcal{S}_{next})$  then
14:         $g.value \leftarrow -1$  {Avoidance value}
15:      else
16:         $g.value \leftarrow 0$ 
17:      end if
18:    end for
19:     $\text{AddPoints}(\mathcal{M}_{tpsm}, G)$ 
20:  end for
21: end for
22: return  $\mathcal{M}_{tpsm}$ 
23: end procedure

```

A.5 Robot Skills

We define three predicate functions: **near**, **pick** and **release** for the navigation, picking and placing skills required for multi-object goal navigation and mobile pick-and-place. As described in Chapter 3, we formalize navigation as continuous path planning problems and manipulation as object parameterized options. Chapter 3 discusses navigation, so here we focus on the pick and place manipulation skills.

Pick Skill: Once the robot has executed the near skill and is at the object to be manipulated, it takes a photo of the current environment to detect the object using the Owl-Vit model. The path planner uses the backprojected object position to compute yaw angles that orient the robot toward the target. After detecting the object, the robot obtains a segmentation mask from the detected bounding box using the Segment Anything model and computes the center pixel of this mask. It passes this pixel to the Boston Dynamics grasping API to compute a grasp motion plan.

Release Skill: We implement a simple routine for the release skill which takes two parameters: the object to be placed and the place receptacle. Once a navigation skill gets the robot to the place receptacle, the robot gently moves its arm up or down to release the grasped object, based on the place receptacle 3D position. More expressive placement strategies could extend the use of LIMP’s spatial grounding of instruction-specific receptacles. Demonstrations are available on the [website](#).

A.6 Evaluation and Baseline Details

All computation including planning, loading and running pretrained visual language models was done on a single computer equipped with one NVIDIA GeForce RTX 3090 GPU. We leverage GPT-4-0613 as the underlying LLM for our instruction understanding module and all our baselines. In all experiments we set the LLM temperature to 0, however since deterministic greedy token decoding is not guaranteed with GPT4, we perform three (3) queries for each instruction and evaluate on the most recurring response (mode response).

We compare LIMP with baseline implementations of NLMaP-Saycan (Chen et al., 2023a) and Code-as-policies (Liang et al., 2023). Both baselines use the same GPT-4 LLM, prompting structure, and in-context learning examples as our language understanding module. We integrate our composable referent descriptor syntax, spatial grounding module and low-level robot control into these baselines as APIs. This enables baselines to execute plans by querying relevant object positions, using our FMT* path planner to find paths to said positions and executing manipulation options.

Figure A.1 shows qualitative planning results. Figure A.2 reports task success across successive chat turns with the interactive verification and reprompting procedure described in Section A.2.1.




Figure A.1: [A] Sample generated plan for a multi object-goal navigation task. [B] Sample generated plan for a mobile pick-and-place task.




Figure A.2: Our interactive reprompting strategy implemented in the symbol verification node regenerates corrective formulas that improve plan success rates with minimal chat turns.

A.6.1 NLMMap-Saycan Implementation Prompt

Listing A.6: Exact prompt to implement NLMMap-Saycan LLM planner

```
1 You are an LLM for robot planning that understands logical operators such as &, |, ~ ,  
   etc. You have a finite set of robot predicates and spatial predicates, given a  
   language instruction, your goal is to generate a sequence of actions that uses  
   appropriate composition of robot and spatial predicates with relevant details from  
   the instruction as arguments.  
2 Robot predicate set (near,pick,release).  
3 Usage:  
4 near[referent_1]:returns true if the desired spatial relationship is for robot to be  
   near referent_1.  
5 pick[referent_1]:can only execute picking skill on referent_1 and return True when  
   near[referent_1].  
6 release[referent_1,referent_2]:can only execute release skill on referent_1 and return  
   True when near[referent_2].  
7 Spatial predicate set (isbetween,isabove,isbelow,isleftof,isrightof,isnextto,  
   isinfrontof,isbehind).  
8 Usage:  
9 referent_1::isbetween(referent_2,referent_3):returns true if referent_1 is between  
   referent_2 and referent_3.  
10 referent_1::isabove(referent_2):returns True if referent_1 is above referent_2.  
11 referent_1::isbelow(referent_2):returns True if referent_1 is below referent_2.  
12 referent_1::isleftof(referent_2):returns True if referent_1 is left of referent_2.  
13 referent_1::isrightof(referent_2):returns True if referent_1 is right of referent_2.  
14 referent_1::isnextto(referent_2):returns True if referent_1 is close to referent_2.  
15 referent_1::isinfrontof(referent_2):returns True if referent_1 is in front of  
   referent_2.  
16 referent_1::isbehind(referent_2):returns True if referent_1 is behind referent_2.  
17 Rules:  
18 Strictly only use the finite set of robot and spatial predicates!  
19 Strictly stick to the usage format!  
20 Compose spatial predicates where necessary!  
21 You should strictly stick to mentioned objects, however you are allowed to propose and  
   include plausible objects if and only if not mentioned in instruction but  
   required based on context of instruction!  
22 Pay attention to instructions that require performing certain actions multiple times  
   in generating and sequencing the predicates for the final Output!  
23 Example:  
24 Input_instruction: Go to the orange building but before that pass by the coffee shop,  
   then go to the parking sign.  
25 Output:  
26 1. near[coffee_shop]  
27 2. near[orange_building]  
28 3. near[parking_sign]  
29 Input_instruction: Go to the blue sofa then the laptop, after that bring me the brown  
   bag between the television and the kettle on the left of the green seat, I am  
   standing by the sink.  
30 Output:  
31 1. near[blue_sofa]
```

```

32 2. near[laptop]
33 3. near[brown_bag::isbetween(television,kettle::isleftof(green_seat))]]
34 4. pick[brown_bag::isbetween(television,kettle::isleftof(green_seat))]]
35 5. near[sink]
36 6. release[brown_bag,sink]
37 Input_instruction: Hey need you to pass by chair between the sofa and bag, pick up the
    bag and go to the orange napkin on the right of the sofa.
38 Output:
39 1. near[chair::isbetween(sofa,bag)]
40 2. near[bag]
41 3. pick[bag]
42 4. near[orange_napkin::isrightof(sofa)]
43 Input_instruction: Go to the chair between the green laptop and the yellow box
    underneath the play toy
44 Output:
45 1. near[chair::isbetween(green_laptop,yellow_box::isbelow(play_toy))]]
46 Input_instruction: Check the table behind the fridge and bring two beers to the couch
    one after the other
47 Output:
48 1. near[table::isbehind(fridge)]
49 2. pick[beer]
50 3. near[couch]
51 4. release[beer,couch]
52 5. near[table::isbehind(fridge)]
53 6. pick[beer]
54 7. near[couch]
55 8. release[beer,couch]
56 Input_instruction: <given_instruction>
57 Output:

```

A.6.2 Code-as-Policies Implementation Prompt

Listing A.7: Exact prompt to implement Code-as-policies planner

```

1 ##Python robot planning script
2 from robotactions import near, pick, release
3 spatial_relationships = [
4 "isbetween", #referent_1::isbetween(referent_2,referent_3):returns true if referent_1
    is between referent_2 and referent_3.
5 "isabove",   #referent_1::isabove(referent_2):returns True if referent_1 is above
    referent_2.
6 "isbelow",   #referent_1::isbelow(referent_2):returns True if referent_1 is below
    referent_2.
7 "isleftof",  #referent_1::isleftof(referent_2):returns True if referent_1 is left of
    referent_2.
8 "isrightof", #referent_1::isrightof(referent_2):returns True if referent_1 is right of
    referent_2.
9 "isnextto",  #referent_1::isnextto(referent_2):returns True if referent_1 is close to
    referent_2.

```

```

10 "isinfrontof", #referent_1::isinfrontof(referent_2):returns True if referent_1 is in
    front of referent_2.
11 "isbehind" #referent_1::isbehind(referent_2):returns True if referent_1 is behind
    referent_2.]
12 ##Rules:
13 ##Strictly only use the finite set of robot and spatial predicates!
14 ##Strictly stick to the usage format!
15 ##Compose spatial predicates where necessary!
16 ##You should strictly stick to mentioned objects, however you are allowed to propose
    and include plausible objects if and only if not mentioned in instruction but
    required based on context of instruction!
17 ##Pay attention to instructions that require performing certain actions multiple times
    in generating and sequencing the predicates for the final Output!
18 # Go to the orange building but before that pass by the coffee shop, then go to the
    parking sign.
19 ordered_navigation_goal_referents = ["coffee_shop", "orange_building", "parking_sign"]
20 for referent in ordered_navigation_goal_referents:
21     near(referent)
22 # Go to the blue sofa then the laptop, after that bring me the brown bag between the
    television and the kettle on the left of the green seat, I am standing by the sink
    .
23 ordered_navigation_goal_referents = ["blue_sofa", "laptop", "brown_bag::isbetween(
    television,kettle::isleftof(green_seat))", "sink"]
24 referents_to_pick= ["brown_bag::isbetween(television,kettle::isleftof(green_seat))"]
25 release_location_referents = ["sink"]
26 picked_item = None
27 for referent in ordered_navigation_goal_referents:
28     near(referent)
29     if referent in referents_to_pick:
30         pick(referent)
31         picked_item = referent
32     if referent in release_location_referents:
33         release(picked_item, referent)
34 #Hey need you to pass by chair between the sofa and bag, pick up the bag and go to the
    orange napkin on the right of the sofa.
35 ordered_navigation_goal_referents = ["chair::isbetween(sofa,bag)", "bag", "
    orange_napkin::isrightof(sofa)"]
36 referents_to_pick= ["bag"]
37 picked_item = None
38 for referent in ordered_navigation_goal_referents:
39     near(referent)
40     if referent in referents_to_pick:
41         pick(referent)
42         picked_item = referent
43 #Go to the chair between the green laptop and the yellow box underneath the play toy
44 near("chair::isbetween(green_laptop,yellow_box::isbelow(play_toy))")
45 #Check the table behind the fridge and bring two beers to the couch one after the
    other
46 for i in range(2):
47     near("table::isbehind(fridge)")

```

```

48     pick("beer")
49     near("couch")
50     release("beer", "couch")
51
52 #<given_instruction>

```

A.6.3 Instruction set

We perform a large-scale evaluation on 150 instructions across five real-world environments. The taskset includes 24 tasks with fine-grained object descriptions (NLMD) from Chen et al. (2023a), 25 tasks with complex language (NLMC) from Chen et al. (2023a), 25 tasks with simple structured phrasing (OKRB) from Liu et al. (2024b), 37 tasks with complex temporal structures (CT) from Liu et al. (2023), and an additional 39 tasks we propose that have descriptive spatial constraints and temporal structures (CST).

Listing A.8: Nlmap Detailed Object Tasks (NLMD)

```

1: put the red can in the trash bin
2: put the brown multigrain chip bag in the woven basket
3: find the succulent plant
4: pick up the up side down mug
5: put the apple on the macbook with yellow stickers
6: use the dyson vacuum cleaner"
7: bring the kosher salt to the kitchen counter
8: put the used towels in washing machine
9: move the used mug to the dish washer
10: place the pickled cucumbers on the shelf
11: find my mug with the shape of a donut
12: put the almonds in the almond jar
13: fill the zisha tea pot with a coke from the cabinet
14: take the slippery floor sign with you
15: take the slippers that have holes on them to the shoe rack
16: find the mug on the mini fridge
17: bring the mint flavor gum to the small table
18: find some n95 masks
19: grab the banana with most black spots
20: fill the empty bottle with lemon juice
21: throw away the apple that's about to rot
22: throw away the rotting banana
23: take the box of organic blueberries out of the fridge
24: give a can of diet coke to the toy cat

```

Listing A.9: Nlmap Complex Language Understanding Tasks (NLMC)

```

1: I opened a pepsi earlier, bring an open can to the orange table
2: I spilled my coke, can you put a replacement on the kitchen counter
3: I spilled some coke on the television, go and bring something to clean it up
4: I accidentally dropped that jalapeno chips after eating it. Would you mind throwing
   it away

```

- 5: I like fruits, can you put something I would like on the yellow sofa for me
- 6: There is a green counter, a yellow counter, and a table. visit all the locations
- 7: There is a green counter, a trash can, and a table. visit all the locations
- 8: Redbull is my faviorite drink, can you put one on the desk please
- 9: Would you bring a coke can to the door for me
- 10: Please, move the pepsi to the red counter
- 11: Can you move the coke can to the orange counter
- 12: Would you throw away the bag of chips for me
- 13: Put an energy bar and water bottle on the table
- 14: Bring a lime soda and a bag of chips to the sofa
- 15: Can you throw away the apple and bring a coke to the bed
- 16: Bring a 7up can and a tea to the office desk
- 17: Move the multigrain chips to the table and an apple to the yellow counter
- 18: Move the lime soda, the sponge, and the water bottle to the table
- 19: Bring two sodas to the table
- 20: Move three cokes to the trash can
- 21: Throw away two cokes from the counter
- 22: Bring two different sodas to the cabinet, there is a coke, pepsi, soup, tea and 7 up in the fridge
- 23: Bring an apple, a coke, and water bottle to the sofa
- 24: I spilled my coke on the table, throw it away and then bring something to help clean
- 25: I just worked out, can you bring me a drink and a snack to recover, i am on the sof

Listing A.10: Ok-Robot Tasks (OKRB)

- 1: Move the Takis on the desk to the nightstand
- 2: Move the soda can to the box
- 3: Move the purple shampoo to the red bag
- 4: Move the white meds box to the trash bin
- 5: Move the power adapter to the chair
- 6: Move the blue gloves to the sink
- 7: Move the McDonalds paper bag to the stove
- 8: Move the herbal tea can to the box
- 9: Move the cooking oil bottle to the marble surface
- 10: Move the milk bottle to the chair
- 11: Move the purple shampoo to the white rack
- 12: Move the purple lightbulb box to the sofa chair
- 13: Pick up purple medicine, drop it on chair
- 14: Pick up eyeglass case, drop it on chair
- 15: Pick up grey rag, drop it in sink
- 16: Pick up golden can rag, drop it on table
- 17: Pick up red navaratna oil, drop it on table
- 18: Pick up purple shampoo, drop it on green rack
- 19: Pick up taki chips, drop it on countertop
- 20: Pick up bandage box, drop it in dustbin
- 21: Pick up white aerosol, drop it in trash can
- 22: Pick up peanut butter, drop it on countertop
- 23: Pick up blue gloves, drop it in sink
- 24: Pick up brown box, drop it on chair

25: Pick up axe body spray, drop it on shel

Listing A.11: Lang2LTL Complex temporal Tasks (CT)

```
1: go to brown bookshelf, metal desk, wooden desk, kitchen counter, and the blue couch
   in any order
2: move to grey door, then bookrack, then go to the brown desk, then counter, then
   white desk
3: visit brown wooden desk but only after bookshelf
4: go from brown bookshelf to white metal desk and only visit each landmark one time
5: go to brown wooden desk exactly once and do not visit brown desk before bookshelf
6: go to brown desk only after visiting bookshelf, in addition go to brown desk only
   after visiting white desk
7: visit the blue IKEA couch, in addition never go to the big steel door
8: visit white kitchen counter then go to brown desk, in addition never visit white
   table
9: go to the grey door, and only then go to the bookshelf, in addition always avoid
   the table
10: go to kitchen counter then wooden desk, in addition after going to counter, you
    must avoid white table
11: Go to bookshelf, alternatively go to metal desk
12: Go to counter, alternatively go to metal desk
13: Go to the counter, but never visit the counter
14: do not go to the wooden desk until bookshelf, and do not go to bookshelf until
    wooden desk
15: go to brown desk exactly once, in addition go to brown desk at least twice
16: move to couch exactly twice, in addition pass by counter at most once
17: navigate to the counter then the brown desk, in addition after going to the
    counter, you must avoid doorway
18: visit counter at least six times
19: either go to bookshelf then the brown desk, or go to couch
20: navigate to the wooden door, the glass door and the table, kitchen counter, and
    the blue couch in any order
21: go to the painting, then find the kitchen table, then front desk, then staircase
22: navigate to classroom but do not visit classroom before the white table
23: only visit classroom once, and do not visit classroom until you visit elevator
    first
24: Go to the staircase, front desk and the white table in that exact order. You are
    not permitted to revisit any of these locations
25: go to the front desk then the yellow office door, in addition do not visit the
    glass door
26: go to the stairs then the front desk, in addition avoid purple elevator
27: move to elevator then front desk, in addition avoid staircase
28: go to front desk then the cabinet, always avoid the elevator
29: Go to elevator, alternatively go to staircase
30: Visit the elevator exactly once, in addition visit the front desk on at least 2
    separate occasions
31: Go to the office, in addition avoid visiting the elevator and the classroom
32: Visit the front desk, in addition you are not permitted to visit elevator and
    staircase
```

- 33: Visit the purple door elevator, then go to the front desk and then go to the kitchen table, in addition you can never go to the elevator once you have seen the front desk
- 34: Visit the front desk then the sofa then the white table, in addition if you visit the sofa you must avoid the television after that
- 35: Go to the glass door, but never visit the glass door
- 36: do not go to the white table until classroom, and do not go to the classroom until white table
- 37: find the office, in addition avoid visiting the front desk and the classroom and the table

Listing A.12: Complex Spatiotemporal Tasks (CST)

- 1: Go to the red sofa but dont pass by any door then go to the computer, the sofa is behind the pillar and on the right side
- 2: Find the door with posters in front of the yellow pillar then go stand by the tree but avoid any sofa
- 3: There is a cabinet, a television and a tree in this room. I want you to go to the last item I mentioned then pass by the second item but dont go close to the first when you are going to thsecond item.",
- 4: Go to the trash bin, after that go to the large television. Actually can you go to the laptop on the table before doing all that?
- 5: Try to go to the tree without going near the trash bin.
- 6: Find the yellow trash bin and go to it, then go by the white table. But before doing the first thing visit where the fridge is then after the last thing, stop by the red sofa but avoid theoden door with posters",
- 7: Go to the television but avoid the tree.
- 8: Go to any trash bin but not the yellow one
- 9: Visit one of the televisions in this room then go to the fridge to the left of the cabinet
- 10: I need you to first go to the table between the pillar and the door with posters, then go to the tree. You know what, ignore the last thing I asked you to go to, after the first thing go tthe computer next to the green box instead.
- 11: Go to the red sofa then the blue one next to the cabinet but first pass by the book shelf
- 12: Go to water filter, dont pass by any computer to the left of the red sofa
- 13: Visit these things in the following order the microwave, then the black chair infront of the yellow robot then the fridge, but before doing any of this go to the robot
- 14: Go to the cabinet between the blue sofa and the brown box.
- 15: I want you to go the cabinet but dont go near any black chair. Hold on, actually can you pass by the whiteboard before doing all that?
- 16: I need you to go to the red sofa but dont go anywhere near the sink
- 17: You can either go as close as you can to the sink or the plant pot but dont go near the yellow robot.
- 18: Go to the green curtain after that the large television but try not to go near the red sofa
- 19: Go to the sink then the red sofa, then the guitar. After that I want you to return to the first thing you visited.

- 20: I need you to go to the sink then any computer on a table but before you start that can you first go to the bookshelf, when you are done with all that, return to brown box
- 21: There are a couple of things in this room, a coffee machine, a robot, a computer and a plant pot. Visit the second item I mentioned then the first, after that visit the brown box, then go the items you haven't visited yet.
- 22: Avoid any cabinet but go get the mug on the blue sofa and take it to red one
- 23: Please get me my plant pot from between the sofa and the sink, bring it to the white cabinet
- 24: Can you pick up my book from the left side of the television and bring it to me? I am sitting on the couch which is to the right of the door with posters. Make sure to never go near the trhbin while doing all this, it on the yellow table.
- 25: Bring the green plush toy to the whiteboard infront of it, watchout for the robot infront of the toy
- 26: Go get my yellow bag and bring it to the table between the yellow pillar and the wooden door with posters next to the whiteboard
- 27: Find the bottle that is on the table to the left of the computer and bring it to the wardrobe that is next to the glass door
- 28: Go to the cabinet between the blue sofa and the yellow robot or the cabinet to the left of the blue sofa.
- 29: Go take the green toy that is next to the sofa under the poster and bring it to the bookshelf
- 30: I have a white cabinet, a green toy, a bookshelf and a red chair around here somewhere. Take the second item I mentioned from between the first item and the third. Bring it the cabinet butvoid the last item at all costs.
- 31: Hey I am standing by the whiteboard infront of the bookshelf, can you bring me the mug from the table to the left of the fridge?
- 32: Never pass by a robot, but i need you to bring the bag on the table to the cabinet
- 33: Take a soda can to the fridge, you can find one on the table to the left of the blue sofa
- 34: Hey, can you pick up my bag for me? It is under the table infront of the glass door, bring it to the cabinet. Actually you know what, forget what i asked for and bring the green toy instea its at the same place",
- 35: Go to the sink and take the mug beside the coffee maker, drop it off at the red sofa
- 36: Go to the sofa behind the chair next to the orange counter in front of the pillar
- 37: I want you to go visit the chair behind the sofa then the green bag, do that 3 times but during the second time avoid the fridge
- 38: Go bring the plush toy between the sofa and the television to the couch, my cat is on the brown table so please dont pass near the table when you are returning with the toy
- 39: I think I left my wallet on the kitchen counter, go get it i will meet you at the bed. There is a lantern near the bed, make sure you dont hit it

Appendix B

Supplementary Methods for REALM

B.1 State Lifecycle and Promotion

Table B.1 summarizes the logical contents of a world revision. Each stored artifact is addressed relative to one run root and paired with a SHA-256 digest. A revision is registered after its checkpoint, twin, planning cloud, semantic registry, transition delta, and validation report are staged successfully.

Table B.1: Logical contents of a persistent world revision.

Field	Examples	Role
Identity	state ID, logical time, revision type	Distinguishes observed, predicted, reconciled-synthetic, and reconciled-real worlds.
Authority	status, current-state index	Identifies the one state from which subsequent views may be generated.
Persistent state	Gaussian checkpoint, object registry	Stores renderable scene content, stable objects, poses, and ownership.
Derived views	digital twin, planning cloud, ground-ing/render bundle	Supplies instruction-specific geometry and observations bound to the state.
Action evidence	plan, trajectory binding, released pose	Explains the intervention used to predict the successor.
Observation evidence	calibrated RGB-D, masks, provenance	Explains the appearance and disocclusion information used in reconciliation.
History	parent ID, reversible delta	Connects each successor to the world from which it evolved.
Validation	locality, fidelity, authority, integrity	Determines whether a candidate may become current.

After checks on the plan binding, rigid transformation, ownership, and history, a prediction can become the current revision before post-action evidence arrives. This is provisional authority based on the planned outcome, not observation-based validation. Reconciliation creates a candidate at the

same logical timestep without overwriting the prediction. Promotion revalidates the lineage, marks the prediction historical, and atomically grants the reconciled successor current authority. Alternative observation and fit variants can remain as sibling candidates; the demonstrated sequence promotes one mainline.

Promotion of a reconciled candidate proceeds as follows: (1) re-hash registered artifacts; (2) check finite and shape-consistent Gaussian tensors; (3) verify stable owner IDs and agreement between twin indices and checkpoint ownership; (4) replay inverse deltas to reconstruct every ancestor; (5) enforce action-owned geometry and opacity; (6) evaluate held-out changed-region and outside-region criteria; and (7) update the current-state index atomically. Failure leaves the preceding current state unchanged.

B.2 Ownership Diagnostics

The selector in Eq. 4.14 combines complementary evidence. Positive depth-consistent voting locates rows that project through the object; visible contradictions reject support inconsistent with other cameras; covariance-footprint overlap tests whether the rendered row actually covers the mask; surface and oriented-box distances bound the object spatially; and scale and anisotropy prevent scene-scale splats from becoming movable object support.

The legacy ball support illustrates their interaction. Of 724 selected centers, 470 lie outside the fitted ball OBB, 443 lie more than 0.01 world units outside it, and 180 have fewer than three footprint-supported views. Median anisotropy is 95.529 and the maximum is 177,390.8. These rows visibly include table and ground content. The combined selector retains 254 rows and is used throughout both transitions. Exact thresholds and per-row diagnostics are stored in the ownership-quality artifact.

Ownership construction rejects empty or insufficient supports, duplicate and out-of-bounds indices, and overlap between independently movable objects. The current analysis covers one ball; extending ownership to thin, articulated, reflective, transparent, and heavily occluded objects remains a central generalization problem.

B.3 Task Specifications and Plan Evidence

The temporal-logic specification for instruction 1 is

$$\begin{aligned}\varphi_1 &= \mathbf{F}(A \wedge \mathbf{F}(B \wedge \mathbf{F}(C \wedge \mathbf{F}D))), \\ A &= \text{near}[\text{blue_ball} :: \text{isbelow}(\text{table})], \\ B &= \text{pick}[\text{blue_ball} :: \text{isbelow}(\text{table})], \\ C &= \text{near}[\text{potted_lime_green_conifer}], \\ D &= \text{release}[\text{blue_ball}, \text{potted_lime_green_conifer}].\end{aligned}$$

Instruction 2 is encoded as

$$\begin{aligned}\varphi_2 &= \mathbf{F}(A \wedge \mathbf{F}(B \wedge \mathbf{F}(C \wedge \mathbf{F}(D \wedge \mathbf{F}E)))), \\ A &= \text{near}[\text{small_garden_ornament_in_bushes}], \\ B &= \text{near}[\text{blue_ball}], \quad C = \text{pick}[\text{blue_ball}], \\ D &= \text{near}[\text{table}], \quad E = \text{release}[\text{blue_ball}, \text{table}].\end{aligned}$$

Here $\mathbf{F}$ is the temporal “eventually” operator. Its nesting expresses nondecreasing temporal order, and $A, \dots, E$ are local abbreviations defined separately for each specification. Instruction 2’s ball referent is evaluated on views rendered from promoted W_1 . The selected record retains `blue_ball_obj0`; the saved plan binds W_1 , its checkpoint, grounding record, and planning cloud.

The first manipulation has 572 frames (323 pick and 249 release); the second has 751 (443 pick and 308 place). A plan-execution manifest assigns every final plan step a unique artifact and verifies the plan and artifact digests. The central collision checker covers 184 pick frames and all 249 release frames for task 1, and 92 pick frames and all 308 place frames for task 2. Approach, attachment, contact, and release are also checked under their declared policies.

B.4 Observation and Reconciliation Protocol

Each observation frame records intrinsics, a rigid camera transform in the REALM world frame, a fit or holdout role, and hashed RGB, metric depth, and masks. Physical bundle types cannot carry synthetic-pixel masks; synthetic bundles require one for every frame. Implemented physical adapters accept registered robot/OSG frames, Record3D, COLMAP, or an already prepared bundle.

For the reported study, rendered geometry determines source-disocclusion and target-insertion masks. Generated color is allowed only inside those masks, after which source pixels are restored exactly outside them. Target depth remains the predicted metric depth rather than generated depth. Stage 1 contains 21 image-edit requests (12 source and 9 target); stage 2 contains 24 (12 source and 12 target). Requests, model identity, masks, and results are content-addressed and cached.

Reconciliation inserts 12,170 unowned disocclusion seeds at stage 1 and 3,551 at stage 2, activates mask-supported scene regions, and optimizes the selected Gaussians. Densification and pruning are disabled for action-owned rows. Predicted owner means, scales, orientations, and opacities are restored after fitting; spherical-harmonic appearance may change. The demonstrated candidates require non-increasing changed-region L1 and an increase of at most 10^{-4} in outside-region L1, both measured against the same synthetic targets on validation views.

B.5 Complete State Sequence

At W_1 , stable-identity association has 15 detection supports, lies 0.0187933 m from the predicted pose, and lies 3.7198327 m from the old pose. Instruction 2’s pick location equals the re-grounded

Table B.2: The five-state mainline produced by the sequential garden study.

State	Revision	Status	Direct evidence
t0000-observed	observed	historical	6,125,086 Gaussian rows from the captured garden.
t0001-predicted	predicted	historical	254 moved rows; 6,124,832 inactive rows exact.
t0001-reconciled-syn.	reconciled-synthetic	historical	12,170 seeds; improved held-out masked L1; preserved action geometry.
t0002-predicted	predicted	historical	254 moved rows; 6,137,002 inactive rows exact.
t0002-reconciled-syn.	reconciled-synthetic	current	3,551 seeds; improved held-out masked L1; final task views agree.

position. At W_2 , the twin reports final ball position [1.0219336, 0.0055544, 1.4529843] in run world units; checkpoint ownership, twin, registry, and planning cloud agree.

Bibliography

- Abou-Chakra, J., Rana, K., Dayoub, F., and Suenderhauf, N. (2025). Physically embodied Gaussian splatting: A visually learnt and physically grounded 3D representation for robotics. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 513–530. PMLR.
- Ackermann, J., Kulhanek, J., Cai, S., Xu, H., Pollefeys, M., Wetzstein, G., Guibas, L. J., and Peng, S. (2025). CL-Splats: Continual learning of Gaussian splatting with local optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7808–7817.
- Athalye, A., Kumar, N., Silver, T., Liang, Y., Wang, J., Lozano-Pérez, T., and Kaelbling, L. P. (2026). From pixels to predicates: Learning symbolic world models via pretrained VLMs. *IEEE Robotics and Automation Letters*, 11(4):4002–4009.
- Barcellona, L., Zadaianchuk, A., Allegro, D., Papa, S., Ghidoni, S., and Gavves, E. (2025). Dream to manipulate: Compositional world models empowering robot imitation learning with imagination. In *International Conference on Learning Representations*.
- Barron, J. T., Mildenhall, B., Verbin, D., Srinivasan, P. P., and Hedman, P. (2022). Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479.
- Berg, M., Bayazit, D., Mathew, R., Rotter-Aboyoun, A., Pavlick, E., and Tellex, S. (2020). Grounding language to landmarks in arbitrary outdoor environments. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 208–215.
- Black, K., Brown, N., Driess, D., Esmail, A., Equi, M. R., Finn, C., Fusai, N., Groom, L., Hausman, K., Ichter, B., Jakubczak, S., Jones, T., Ke, L., Levine, S., Li-Bell, A., Mothukuri, M., Nair, S., Pertsch, K., Shi, L. X., Smith, L., Tanner, J., Vuong, Q., Walling, A., Wang, H., and Zhilinsky, U. (2025). π_0 : A vision-language-action flow model for general robot control. In *Proceedings of Robotics: Science and Systems*.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., et al. (2021). On the opportunities and risks of foundation models. arXiv preprint arXiv:2108.07258.

- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901.
- Burridge, R. R., Rizzi, A. A., and Koditschek, D. E. (1999). Sequential composition of dynamically dexterous robot behaviors. *The International Journal of Robotics Research*, 18(6):534–555.
- Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., and Leonard, J. J. (2016). Past, present, and future of simultaneous localization and mapping: Towards the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332.
- Cadene, R., Alibert, S., Soare, A., Gallouedec, Q., Zouitine, A., Palma, S., Kooijmans, P., Aractingi, M., Shukor, M., Aubakirova, D., Russi, M., Capuano, F., Pascal, C., Choghari, J., Meftah, K., Ellerbach, M., Moss, J., and Wolf, T. (2024). LeRobot: State-of-the-art machine learning for real-world robotics in PyTorch. <https://github.com/huggingface/lerobot>.
- Chai, Y., Deng, L., Shao, R., Zhang, J., Lv, K., Xing, L., Li, X., Zhang, H., and Liu, Y. (2026). GAF: Gaussian action field as a 4D representation for dynamic world modeling in robotic manipulation. In *2026 IEEE International Conference on Robotics and Automation (ICRA)*. Accepted paper.
- Chen, B., Xia, F., Ichter, B., Rao, K., Gopalakrishnan, K., Ryoo, M. S., Stone, A., and Kappler, D. (2023a). Open-vocabulary queryable scene representations for real world planning. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11509–11522.
- Chen, Y., Gandhi, R., Zhang, Y., and Fan, C. (2023b). NL2TL: Transforming natural languages to temporal logics using large language models. In Bouamor, H., Pino, J., and Bali, K., editors, *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15880–15903, Singapore. Association for Computational Linguistics.
- Chi, C., Xu, Z., Feng, S., Cousineau, E., Du, Y., Burchfiel, B., Tedrake, R., and Song, S. (2025). Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10–11):1684–1704.
- Chisari, E., Heppert, N., Argus, M., Welschehold, T., Brox, T., and Valada, A. (2025). Learning robotic manipulation policies from point clouds with conditional flow matching. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 982–993. PMLR.
- Colledanchise, M., Murray, R. M., and Ögren, P. (2017). Synthesis of correct-by-construction behavior trees. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6039–6046.

- Colledanchise, M. and Ögren, P. (2018). *Behavior Trees in Robotics and AI: An Introduction*. CRC Press.
- Cosler, M., Hahn, C., Mendoza, D., Schmitt, F., and Trippel, C. (2023). nl2spec: Interactively translating unstructured natural language to temporal logics with large language models. In Enea, C. and Lal, A., editors, *Computer Aided Verification*, pages 383–396, Cham. Springer Nature Switzerland.
- Diab, M., Pomarlan, M., Beßler, D., Akbari, A., Rosell, J., Bateman, J., and Beetz, M. (2020). SkillMaN — a skill-based robotic manipulation framework based on perception and reasoning. *Robotics and Autonomous Systems*, 134:103653.
- Emerson, E. A. (1990). Temporal and modal logic. In van Leeuwen, J., editor, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*, pages 995–1072. Elsevier and MIT Press.
- Fikes, R. E. and Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3–4):189–208.
- Fuggitti, F. and Chakraborti, T. (2023). NL2LTL – a Python package for converting natural language (NL) instructions to linear temporal logic (LTL) formulas. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(13):16428–16430.
- Gadre, S. Y., Wortsman, M., Ilharco, G., Schmidt, L., and Song, S. (2023). CoWs on pasture: Baselines and benchmarks for language-driven zero-shot object navigation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23171–23181. IEEE.
- Garrett, C. R., Chitnis, R., Holladay, R., Kim, B., Silver, T., Kaelbling, L. P., and Lozano-Pérez, T. (2021). Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:265–293.
- Garrett, C. R., Lozano-Pérez, T., and Kaelbling, L. P. (2020). PDDLStream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning. *Proceedings of the International Conference on Automated Planning and Scheduling*, 30:440–448.
- Ghallab, M., Howe, A., Knoblock, C., McDermott, D., Ram, A., Veloso, M., Weld, D., and Wilkins, D. (1998). PDDL—the planning domain definition language. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control. Version 1.2.
- Gu, Q., Kuwajerwala, A., Morin, S., Jatavallabhula, K. M., Sen, B., Agarwal, A., Rivera, C., Paul, W., Ellis, K., Chellappa, R., Gan, C., de Melo, C. M., Tenenbaum, J. B., Torralba, A., Shkurti, F., and Paull, L. (2024). ConceptGraphs: Open-vocabulary 3D scene graphs for perception and planning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5021–5028.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. (2017). On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.

- Harel, D. (1987). Statecharts: a visual formalism for complex systems. *Science of Computer Programming*, 8(3):231–274.
- Hartley, R. I. and Zisserman, A. (2004). *Multiple View Geometry in Computer Vision*. Cambridge University Press, second edition.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851.
- Hu, Z., Huang, S., and Ishikawa, M. (2026). GaussMemory: Task-driven 3D Gaussian scene memory for long-horizon robotic manipulation. In *2026 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Accepted paper.
- Huang, S., Chen, Q., Zhang, X., Sun, J., and Schwager, M. (2025). ParticleFormer: A 3D point cloud world model for multi-object, multi-material robotic manipulation. In *Proceedings of the 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 4941–4957. PMLR.
- Hughes, N., Chang, Y., and Carlone, L. (2022). Hydra: A real-time spatial perception system for 3D scene graph construction and optimization. In *Proceedings of Robotics: Science and Systems*.
- Ichter, B., Brohan, A., Chebotar, Y., Finn, C., Hausman, K., Herzog, A., Ho, D., Ibarz, J., Irpan, A., Jang, E., Julian, R., Kalashnikov, D., Levine, S., Lu, Y., Parada, C., Rao, K., Sermanet, P., Toshev, A. T., Vanhoucke, V., Xia, F., Xiao, T., Xu, P., Yan, M., Brown, N., Ahn, M., Cortes, O., Sievers, N., Tan, C., Xu, S., Reyes, D., Rettinghouse, J., Quiambao, J., Pastor, P., Luu, L., Lee, K.-H., Kuang, Y., Jesmonth, S., Joshi, N. J., Jeffrey, K., Ruano, R. J., Hsu, J., Gopalakrishnan, K., David, B., Zeng, A., and Fu, C. K. (2023). Do as I can, not as I say: Grounding language in robotic affordances. In *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 287–318. PMLR.
- Janner, M., Du, Y., Tenenbaum, J., and Levine, S. (2022). Planning with diffusion for flexible behavior synthesis. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 9902–9915. PMLR.
- Janson, L., Schmerling, E., Clark, A., and Pavone, M. (2015). Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *The International Journal of Robotics Research*, 34(7):883–921.
- Jatavallabhula, K., Kuwajerwala, A., Gu, Q., Omama, M., Iyer, G., Saryazdi, S., Chen, T., Maalouf, A., Li, S., Keetha, N., Tewari, A., Tenenbaum, J., Melo, C., Krishna, M., Paull, L., Shkurti, F., and Torralba, A. (2023). ConceptFusion: Open-set multimodal 3D mapping. In *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation.
- Kerbl, B., Kopanas, G., Leimkühler, T., and Drettakis, G. (2023). 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4).

- Kim, M. J., Pertsch, K., Karamcheti, S., Xiao, T., Balakrishna, A., Nair, S., Rafailov, R., Foster, E. P., Sanketi, P. R., Vuong, Q., Kollar, T., Burchfiel, B., Tedrake, R., Sadigh, D., Levine, S., Liang, P., and Finn, C. (2025). OpenVLA: An open-source vision-language-action model. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 2679–2713. PMLR.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. (2023). Segment anything. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3992–4003.
- Konidaris, G., Kaelbling, L. P., and Lozano-Pérez, T. (2018). From skills to symbols: Learning symbolic representations for abstract high-level planning. *Journal of Artificial Intelligence Research*, 61:215–289.
- Kreber, J. U., Mack, L., and Stueckler, J. (2026). Learning action-conditional and object-centric Gaussian splatting world models for rigid objects. arXiv preprint arXiv:2606.01950.
- Kress-Gazit, H., Fainekos, G. E., and Pappas, G. J. (2009). Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics*, 25(6):1370–1381.
- Kupferman, O. and Vardi, M. Y. (2001). Model checking of safety properties. *Formal Methods in System Design*, 19:291–314.
- Liang, J., Huang, W., Xia, F., Xu, P., Hausman, K., Ichter, B., Florence, P., and Zeng, A. (2023). Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500.
- Lipman, Y., Chen, R. T. Q., Ben-Hamu, H., Nickel, M., and Le, M. (2023). Flow matching for generative modeling. In *Proceedings of the 11th International Conference on Learning Representations*.
- Liu, J. X., Shah, A., Konidaris, G., Tellex, S., and Paulius, D. (2024a). Lang2LTL-2: Grounding spatiotemporal navigation commands using large language and vision-language models. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2325–2332.
- Liu, J. X., Yang, Z., Idrees, I., Liang, S., Schornstein, B., Tellex, S., and Shah, A. (2023). Grounding complex natural language commands for temporal tasks in unseen environments. In *Proceedings of The 7th Conference on Robot Learning*, pages 1084–1110. PMLR.
- Liu, P., Guo, Z., Warke, M., Chintala, S., Paxton, C., Shafiullah, N. M. M., and Pinto, L. (2025a). DynaMem: Online dynamic spatio-semantic memory for open world mobile manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13346–13355.
- Liu, P., Orru, Y., Vakil, J., Paxton, C., Shafiullah, N. M. M., and Pinto, L. (2024b). Demonstrating OK-Robot: What really matters in integrating open-knowledge models for robotics. In *Robotics: Science and Systems*. Robotics: Science and Systems Foundation.

- Liu, W., Nie, N., Zhang, R., Mao, J., and Wu, J. (2025b). Learning compositional behaviors from demonstration and language. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 1992–2028. PMLR.
- Lu, G., Jia, B., Li, P., Chen, Y., Wang, Z., Tang, Y., and Huang, S. (2025). GWM: Towards scalable Gaussian world models for robotic manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9263–9274.
- Majid, A., Bowerman, M., Kita, S., Haun, D. B. M., and Levinson, S. C. (2004). Can language restructure cognition? The case for space. *Trends in Cognitive Sciences*, 8(3):108–114.
- Mandlekar, A., Nasiriany, S., Wen, B., Akinola, I., Narang, Y., Fan, L., Zhu, Y., and Fox, D. (2023). MimicGen: A data generation system for scalable robot learning using human demonstrations. In *Proceedings of the 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 1820–1864. PMLR.
- Mandlekar, A., Xu, D., Wong, J., Nasiriany, S., Wang, C., Kulkarni, R., Fei-Fei, L., Savarese, S., Zhu, Y., and Martín-Martín, R. (2022). What matters in learning from offline human demonstrations for robot manipulation. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 1678–1690. PMLR.
- Martín, F., Ginés, J., Matellán, V., and Rodríguez, F. J. (2021). PlanSys2: A planning system framework for ROS2. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9742–9749.
- Mayr, M., Rovida, F., and Krueger, V. (2023). SkiROS2: A skill-based robot control platform for ROS. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6273–6280.
- Menghi, C., Tsigkanos, C., Pelliccione, P., Ghezzi, C., and Berger, T. (2021). Specification patterns for robotic missions. *IEEE Transactions on Software Engineering*, 47(10):2208–2224.
- Migimatsu, T., Lian, W., Bohg, J., and Schaal, S. (2022). Symbolic state estimation with predicates for contact-rich manipulation tasks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 1702–1709.
- Minderer, M., Gritsenko, A., Stone, A., Neumann, M., Weissenborn, D., Dosovitskiy, A., Mahendran, A., Arnab, A., Dehghani, M., Shen, Z., Wang, X., Zhai, X., Kipf, T., and Houlsby, N. (2022). Simple open-vocabulary object detection. In Avidan, S., Brostow, G., Cissé, M., Farinella, G. M., and Hassner, T., editors, *Computer Vision – ECCV 2022*, pages 728–755, Cham. Springer Nature Switzerland.
- Pan, J., Chou, G., and Berenson, D. (2023). Data-efficient learning of natural language to linear temporal logic translators for robot task specification. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11554–11561.

- Paulius, D., Agostini, A., Quartey, B., and Konidaris, G. (2025). Bootstrapping object-level planning with large language models. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16233–16239. IEEE.
- Pnueli, A. (1977). The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, pages 46–57.
- Quartey, B., Castro, S., Rosen, E., Thomason, W., Konidaris, G., and Tellex, S. (2026). Jointly learning predicates and actions enables zero-shot skill composition. In *2026 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Accepted paper.
- Quartey, B., Rosen, E., Tellex, S., and Konidaris, G. (2025). Verifiably following complex robot instructions with foundation models. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE.
- Quartey, B., Shah, A., and Konidaris, G. (2023). Exploiting contextual structure to generate useful auxiliary tasks. In *NeurIPS 2023 Workshop on Generalization in Planning*.
- Rana, K., Abou-Chakra, J., Garg, S., Lee, R., Reid, I., and Suenderhauf, N. (2025). Learning from 10 demos: Generalisable and sample-efficient policy learning with oriented affordance frames. In *Proceedings of the 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 5464–5482. PMLR.
- Rosen, E., James, S., Orozco, S., Gupta, V., Merlin, M., Tellex, S., and Konidaris, G. (2023). Synthesizing navigation abstractions for planning with portable manipulation skills. In *Proceedings of The 7th Conference on Robot Learning*, pages 2278–2287. PMLR.
- Ross, S., Gordon, G., and Bagnell, D. (2011). A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pages 627–635. PMLR.
- Shah, D., Osiński, B., Ichter, B., and Levine, S. (2023). LM-Nav: Robotic navigation with large pre-trained models of language, vision, and action. In *Proceedings of The 6th Conference on Robot Learning*, pages 492–504. PMLR.
- Shen, W., Yang, G., Yu, A., Wong, J., Kaelbling, L. P., and Isola, P. (2023). Distilled feature fields enable few-shot language-guided manipulation. In *Proceedings of the 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 405–424. PMLR.
- Sun, Z., Bao, L., Peng, T., Sun, J., and Zhou, C. (2026). A high-fidelity digital twin for robotic manipulation based on 3D Gaussian splatting. *Robot Learning*, 3(2):0013.
- Sutton, R. S., Precup, D., and Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1–2):181–211.

- Tellex, S., Gopalan, N., Kress-Gazit, H., and Matuszek, C. (2020). Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:25–55.
- Tong, A., Fatras, K., Malkin, N., Huguet, G., Zhang, Y., Rector-Brooks, J., Wolf, G., and Bengio, Y. (2024). Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*.
- Vardi, M. Y. (1996). An automata-theoretic approach to linear temporal logic. In Moller, F. and Birtwistle, G., editors, *Logics for Concurrency: Structure versus Automata*, pages 238–266. Springer, Berlin, Heidelberg.
- Wang, D., Hart, S., Surovik, D., Kelestemur, T., Huang, H., Zhao, H., Yeatman, M., Wang, J., Walters, R., and Platt, R. (2025a). Equivariant diffusion policy. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 48–69. PMLR.
- Wang, Y., Zhang, M., Li, Z., Kelestemur, T., Driggs-Campbell, K. R., Wu, J., Fei-Fei, L., and Li, Y. (2025b). D^3 Fields: Dynamic 3D descriptor fields for zero-shot generalizable rearrangement. In *Proceedings of the 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 272–298. PMLR.
- Xu, Z., He, Z., Wu, J., and Song, S. (2021). Learning 3D dynamic scene representations for robot manipulation. In *Proceedings of the 2020 Conference on Robot Learning*, volume 155 of *Proceedings of Machine Learning Research*, pages 126–142. PMLR.
- Yan, Z., Li, S., Wang, Z., Wu, L., Wang, H., Zhu, J., Chen, L., and Liu, J. (2025). Dynamic open-vocabulary 3D scene graphs for long-term language-guided mobile manipulation. *IEEE Robotics and Automation Letters*, 10(5):4252–4259.
- Yang, Z., Raman, S. S., Shah, A., and Tellex, S. (2024). Plug in the safety chip: Enforcing constraints for LLM-driven robot agents. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14435–14442.
- Yenamandra, S., Ramachandran, A., Yadav, K., Wang, A. S., Khanna, M., Gervet, T., Yang, T.-Y., Jain, V., Clegg, A., Turner, J. M., Kira, Z., Savva, M., Chang, A. X., Chaplot, D. S., Batra, D., Mottaghi, R., Bisk, Y., and Paxton, C. (2023). HomeRobot: Open-vocabulary mobile manipulation. In *Proceedings of The 7th Conference on Robot Learning*, pages 1975–2011. PMLR.
- Yokoyama, N., Clegg, A., Truong, J., Undersander, E., Yang, T.-Y., Arnaud, S., Ha, S., Batra, D., and Rai, A. (2024). ASC: Adaptive skill coordination for robotic mobile manipulation. *IEEE Robotics and Automation Letters*, 9(1):779–786.
- Yuan, W., Paxton, C., Desingh, K., and Fox, D. (2022). SORNet: Spatial object-centric representations for sequential manipulation. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pages 148–157. PMLR.

Zhang, Z., Jiang, Y., Cheng, Q., Li, X., Liu, S., Zhao, D., Luo, P., Zhou, W., and Yu, H. (2026). GaussianDream: A feed-forward 3D Gaussian world model for robotic manipulation. arXiv preprint arXiv:2605.20752.

Zitkovich, B., Yu, T., Xu, S., Xu, P., Xiao, T., Xia, F., Wu, J., Wohlhart, P., Welker, S., Wahid, A., Vuong, Q., Vanhoucke, V., Tran, H., Soricut, R., Singh, A., Singh, J., Sermanet, P., Sanketi, P. R., Salazar, G., Ryoo, M. S., Reymann, K., Rao, K., Pertsch, K., Mordatch, I., Michalewski, H., Lu, Y., Levine, S., Lee, L., Lee, T.-W. E., Leal, I., Kuang, Y., Kalashnikov, D., Julian, R., Joshi, N. J., Irpan, A., Ichter, B., Hsu, J., Herzog, A., Hausman, K., Gopalakrishnan, K., Fu, C., Florence, P., Finn, C., Dubey, K. A., Driess, D., Ding, T., Choromanski, K. M., Chen, X., Chebotar, Y., Carbajal, J., Brown, N., Brohan, A., Arenas, M. G., and Han, K. (2023). RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of the 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 2165–2183. PMLR.