

In-Context Learning for Robots

Methods and Applications

Haojian Huang^{1,2}, Zexi Li^{1,3}, Junhao Guo¹, Yehang Zhang^{1,2}, Wenxuan Peng^{1,4},
Bohan Zhou^{1,3}, Weilin Ruan^{1,3}, Leyi Wu^{1,2}, Chenxu Wang^{1,5}, Jianchong Su^{1,2}, Binghui Xie^{1,3},
Wosong Chen^{1,2}, Yingjie Xu^{1,2}, Tianhao Zhou^{1,2}, Suzeyu Chen^{1,2}, Pukun Zhao¹, Jiaqi He¹,
Xinyi Li^{1,3}, Runze Li⁷, Peiran Dong^{1,3}, Shaoxiang Dang¹, Jing Huang¹, Yingbing Chen¹,
Yifan Chang¹, Tianyi Zhang¹, Shiyuan Deng¹, Haozhi Wang¹, Yangkai Wei¹, Wenqian Li¹,
Han Yang¹, Kaiwen Zhou¹, Huaping Liu⁵, James Cheng³, Rui Shao⁶, Donglin Wang⁷,
Yaochu Jin⁷, Jianye Hao⁸, Ying-Cong Chen^{1,2,*}, Yinchuan Li^{1,*}

¹Knowin AI ²HKUST(GZ) ³The Chinese University of Hong Kong ⁴Tongji University
⁵Tsinghua University ⁶Harbin Institute of Technology, Shenzhen ⁷Westlake University ⁸Tianjin University

* Corresponding authors.



Project Page

GitHub

General-purpose robots must infer what a new task requires and translate that understanding into appropriate physical action. In-context learning (ICL) for robots supports this process by using demonstrations and interaction to direct existing competence with neural parameters held fixed during deployment. We organize this literature review around the interfaces connecting contextual evidence to execution, distinguishing four families: context-conditioned policies, geometric demonstration transfer, world-model-based control, and skill- and agent-based execution. Comparing these interfaces clarifies their transfer assumptions and the roles of training, correspondence, and memory in making context useful. Across manipulation and navigation, we examine how these mechanisms preserve taught requirements as objects, environments, and execution conditions change. This analysis links method design to evaluation practices that distinguish responsiveness to teaching, physical transfer, and benefits from retained experience. The resulting agenda connects compositional task acquisition and faithful transfer with physical recursive self-improvement, in which experience improves the ability to learn subsequent tasks.

✉ CONTACT {huanghaojian, lizexi}@ knowin.ai




Figure 1 Six learning horizons. S1 encodes control explicitly; S2 learns reusable policies. With neural parameters fixed at deployment, S3 adapts to interaction outcomes, while S4 infers task requirements from teaching. S5 improves how subsequent tasks are learned; S6 extends this process through validated knowledge exchange across embodiments.

Contents

1	Introduction	3
2	Foundations: Why Do Robots Need Context?	5
2.1	From competence to contextual adaptation	5
2.2	Contextual and parameter adaptation	6
2.3	Four method families and shared design dimensions	9
2.4	Historical development	11
3	Methods: How Does Context Change Action?	15
3.1	Context-conditioned policies	16
3.2	Geometric demonstration transfer	20
3.3	World-model-based control	23
3.4	Skill- and agent-based execution	26
3.5	Shared mechanisms: correspondence and memory	29
3.6	External knowledge reuse and revision	35
4	Data and Training: How Is Context Use Learned?	38
4.1	Input sources and context modalities	38
4.2	Training units and context dependence	40
4.3	Data coverage and action correspondence	44
4.4	Acquisition cost and contextual capability	47
4.5	Scaling data construction	48
4.6	Connecting teaching data to a pretrained policy	50
5	Grounding: How Does Context Guide Recovery?	51
5.1	Grounding the task in the current physical state	51
5.2	From failure evidence to corrective execution	53
5.3	Applicability and reuse of corrections	54
6	Applications: What Must Context Resolve?	54
6.1	Task specification in manipulation	55
6.2	Navigation and spatial adaptation	57
6.3	Long procedures and irreversible operations	58
6.4	Adaptation to body and dynamics	60
6.5	Task changeover, teaching, and reuse	60
7	Evaluation: When Does Context Help?	61
7.1	Context use and transfer across tasks and environments	62
7.2	Sources of improvement and comparable evidence	65
7.3	Physical realization and learning across attempts	68
7.4	Teaching effort, response delay, and reuse cost	69
8	Future Directions: What Limits Transfer and Improvement?	70
8.1	From context dependence to a reusable learning rule	70
8.2	Instruction tuning for more complex contextual learning	71
8.3	Preserving taught relations across physical change	72
8.4	Retaining teaching while revising obsolete experience	73
8.5	Physical recursive self-improvement	74
9	Conclusion	76
	References	76
A	Complete Reference Index	96
B	Method and Resource Comparisons	97
C	Related Survey Perspectives	100

1 Introduction

A robot may know how to move yet still need evidence about what to do. Demonstrations specify a fold or assembly order; corrections revise the procedure. Interaction reveals friction or misalignment. In navigation, earlier visits reveal locations outside the current view [203]. In-context learning (ICL) for robots studies how such evidence changes deployed behavior without another task-specific update to neural parameters.

Broad pretraining supplies reusable perception and control [118–120] from multi-task and multi-embodiment collections [57, 60]. Yet the current scene can leave several procedures feasible, conceal an earlier event, or reveal little about an unfamiliar material’s response. These are information gaps that broader motor competence alone cannot resolve. Fine-tuning incorporates new evidence into weights; ICL makes it available at the decision. Cross-object transfer tests whether teaching remains useful after replacing either or both interacting objects [346]. The task relation must survive while its physical realization changes.

Two historical roots motivate a broad notion of context: one-shot imitation infers intended behavior, while meta-reinforcement learning infers tasks or dynamics from outcomes [1, 2]. Sequence models retain trajectories and learning histories [5, 6]; multimodal prompting combines task specifications and sensorimotor examples [13, 100]. With broader priors, teaching can direct generalist actions [98, 99], predicted futures [24], or executable procedures [31]. Physical experiments supply evidence for revising execution [88, 89]. Language participates throughout as instructions, examples, corrections, and retained summaries. Figure 1 situates these capabilities within complementary horizons of adaptation and longer-term learning.

Context use depends on the relationships learned during training. GPT-3 demonstrated few-shot performance after autoregressive pretraining [121]; controlled studies identify distributional conditions that promote using examples [111]. For robots, the critical relationship links earlier teaching or interaction to the later action it changes. Retaining a correction can extend that relationship across attempts [47, 49], provided its conditions remain valid. Broader motor competence, more informative teaching, and selective experience reuse therefore address different limits of adaptation.

Existing surveys organize learning from demonstration by teaching interfaces and learned policies, rewards, or plans [122, 123]. ICL and in-context RL reviews examine adaptation through examples and interaction [124, 125], including the validity of context under environmental change [127]. In robotics, human-video reviews compare the information transferred from observation to control [126]; manipulation ICL reviews organize context content, inference targets, adaptation mechanisms, and transfer [372]. VLM-based VLA reviews distinguish monolithic and hierarchical integration of planning and action generation [373]. Complementary reviews emphasize data and evaluation [129], cross-embodiment adaptation [130], predictive control [128, 178], and the acquisition and improvement of executable skills [370].

Our organizing question is *how new evidence resolves what existing competence leaves undetermined, and how that resolution survives physical execution*. The four families place this burden in different intermediates; correspondence, training, and memory determine whether those intermediates preserve the needed information. This connects the reason for using context to the mechanism and limits of transfer. Table C1 in Appendix C compares related reviews, including recent world-action taxonomies [371].

Placing an object at the correct destination can still violate a required handle grasp. We therefore trace both the intended effect and prescribed order or contact through execution, developing three contributions:

- A taxonomy of context-conditioned policies, geometric demonstration transfer, world-model-based control, and skill- and agent-based execution, with shared correspondence and memory.
- An account of how training relationships establish context use and how object substitution, unfamiliar environments, and execution conditions limit transfer, separating broader motor competence from broader ability to learn through teaching.
- A synthesis of reported comparisons and evaluation controls that separates context dependence, transfer, and retained-experience benefits, motivating compositional learning and improved teachability.




Figure 2 Organization of the literature on ICL for robots. Four method families connect contextual evidence to behavior; the other branches organize foundations, data, grounding, applications, evaluation, and research directions. Correspondence and memory mechanisms connect methods across families.




Figure 3 Technical foundations for six learning horizons. S1–S2 establish programmed and learned competence. S3 and S4 adapt execution through interaction history and supplied teaching, respectively, with fixed neural parameters. The S5 and S6 columns collect mechanisms supporting improvements in subsequent learning and knowledge exchange across embodiments, including reusable procedures, shared representations, and policy transfer.

Figure 2 follows this argument: why a robot needs context (Section 2), how context changes action (Section 3), and how that dependence is learned (Section 4). Recovery and applications expose its physical limits (Sections 5–6); evaluation tests its benefit (Section 7); the research agenda asks which limits must be overcome for broader transfer and improved learning (Section 8).

2 Foundations: Why Do Robots Need Context?

A packing robot may already grasp every part yet need a demonstration to establish their order, history to identify completed placements, and contact evidence to accommodate a tighter fit. The same visible scene can therefore require different actions. Useful context resolves this decision-relevant ambiguity through actions the robot can already execute. This distinction motivates the learning horizons, adaptation regimes, and four control interfaces below. Their history traces how robots acquired broader ways to interpret and act on new evidence.

2.1 From competence to contextual adaptation

We use six learning horizons to organize the relationship between prior competence, contextual adaptation, and learning from retained experience. S3 and S4 describe complementary sources of evidence; S5 and S6 express research objectives whose realization can combine several earlier capabilities. Together, the horizons connect distinct adaptation mechanisms to individual and collective learning objectives.

The six horizons distinguish where learning changes the system (Figure 3). Explicit control programming (S1) encodes task logic, control laws, and planning models [134–138]; neural policy learning (S2) acquires reusable behavior through imitation and reinforcement [139, 140]. History-based adaptation (S3) uses interaction memory to infer task or dynamics information. In-context task learning (S4) infers new task requirements from supplied teaching, including instructions, demonstrations, textual examples, and corrections. Both forms of ICL for robots redirect execution with deployed neural parameters fixed.




Figure 4 Fine-tuning and in-context learning. (a) Task data $\mathcal{D}_{\text{task}}$ update θ to θ_{task} . (b) A demonstration supplies $C_t = D^s$ with θ fixed. Both policies use observation–action history h_t to predict the next action block A_t .

In assembly, contact outcomes can reveal a tighter fit (S3), while a demonstration or correction can specify a new part order (S4). A retained correction may guide both interpretation and execution; these roles can share a model and input modality.

Automated practice [114, 141], reward and experiment design [89, 115, 116], and retained guidance [47, 49] support a further objective: physical recursive self-improvement (S5), in which experience improves how the next task is learned. Collective knowledge evolution (S6) extends this learning cycle to a group of robots. Its premise is that an execution contains knowledge another body can use, even when the original motion cannot be copied. Shared progress and skill representations preserve what should be accomplished [29, 30, 79]; pooled policy learning supplies competence across sensing and action spaces [57, 132]; morphology-aware transfer changes the physical realization [131, 133]. These are complementary foundations for exchanging demonstrations and corrections. Whether such exchanges improve the group’s subsequent learning is the S6 research objective in Figure 1.

Physical self-improvement describes the feedback loop through which interaction changes later behavior. It can operate through retained context, executable programs, or neural updates. This is a cross-cutting process: S3 identifies history-based adaptation, S5 concerns improvements in subsequent learning, and S6 adds knowledge exchange between robots. Section 8.5 compares these update objects and their evaluation.

New evidence changes either policy parameters or inference-time inputs (Figure 4). Fine-tuning uses task data $\mathcal{D}_{\text{task}}$ to update pretrained parameters θ to task-adapted parameters θ_{task} . The illustrated ICL case supplies a support demonstration D^s as context C_t , keeping θ fixed. Both use existing perception and motor competence; the storage of task information determines acquisition cost, persistence, and reset behavior.

What context enables a robot to learn. Fixed-parameter adaptation through supplied teaching or interaction is the review’s central subject. Parameter-adaptation methods provide comparisons; pretrained policies and control mechanisms supply the competence on which adaptation operates. New evidence can specify a command, identify a familiar routine, compose known relations, or teach an unfamiliar rule. These behaviors differ in the information acquired from context and the prior competence used to realize it. Section 8.1 develops these distinctions. The initial test is whether changing task-defining evidence changes behavior appropriately under matched execution conditions.

2.2 Contextual and parameter adaptation

Adaptation can reside in contextual state, a retained external artifact, or neural parameters. A common policy interface makes these alternatives comparable; their storage and reset behavior determine how an evaluation can identify the source of change. Table 1 collects the shared notation, separating supplied evidence, execution intermediates, and retained state.

At decision step t , the current attempt has history $h_t = (o_1, a_1, \dots, a_{t-1}, o_t)$. Observation o_t includes the available images, proprioception, force, tactile or inertial measurements, and feedback messages. Action a_t is the applied command; its measured physical effect appears in subsequent observations. Task-relevant context C_t may contain instructions, goal references, multimodal support demonstrations D^s , textual input–output examples, corrections, retrieved episodes, or memory of interaction. Superscript s denotes support evidence;

Table 1 Core notation used throughout the survey. Groups follow the argument from contextual evidence to execution, retained knowledge, and evaluation. Superscripts s and q distinguish support and query; other superscripts identify component roles.

Notation	Meaning
<i>Interaction and contextual evidence (Section 2)</i>	
t, n	Decision step; complete-attempt index.
i, j, ℓ, b	Demonstration element; memory record; skill; motion segment.
o_t, a_t, h_t	Observation; executed command; observation–action history.
$A_t, \hat{a}_{t+\omega t}, H$	Proposed action block; action at offset ω ; block horizon.
D^s, L, d_i^s	Support demonstration; its length; its i th element.
$C_t, C_t^{\text{task}}, r_t$	Policy context; supplied task evidence; interpreted representation.
θ, θ_0	Deployed neural parameters; initialization of the adaptation phase.
$\bar{\theta}, \phi_t$	Frozen backbone; adaptive parameter subset in test-time training.
<i>Context-to-action interfaces (Sections 3.1–3.4)</i>	
π_θ	Conditional distribution of the proposed action block.
E_θ^c, E_θ^m	Context interpreter; demonstration motion-descriptor extractor.
q_θ^a, q_θ^v	Action generators conditioned on context or a predicted future.
$v^+, p_\theta^v, p_\theta^d$	Anticipated future; context-conditioned predictor; action-conditioned dynamics.
g, f_θ, p_θ^g	Execution specification; deterministic or stochastic specification generator.
$\kappa, k_\theta, \kappa_{B_n}$	Deterministic, stochastic, and artifact-dependent executors.
z_ℓ, N_g, ℓ_t	Skill specification; number of skills; active skill index.
$\zeta_i, \mathcal{B}$	Motion descriptor; robot kinematics and control description.
<i>Memory and external revision (Sections 2.2, 3.5–3.6)</i>	
U_θ	Update of a self-contained context from an observed action transition.
m_t, M_t, R_t	Recurrent state; external archive; retrieved records.
x_t, ξ_j, J_t, N_t	Executed transition; archive record; selected indices; archive size.
$B_n, \tilde{B}_{n+1}$	Retained external artifacts; proposed revision.
$\mathcal{T}_n, \chi_n$	Attempt trace with feedback; revision-acceptance indicator.
<i>Training and evaluation (Sections 4, 7)</i>	
$\psi, \tau^q, P_{\text{pair}}$	Task specification; query trajectory; support–query pair distribution.
$I^q, \ell_\theta, \mathcal{L}$	Supervised query positions; per-target action loss; expected training loss.
$S_n, \Delta_n^{\text{reuse}}$	Success/adherence rate; gain from retained evidence at attempt n .

the demonstration-only case in Figure 4 sets $C_t = D^s$. Context can summarize the current history or retain information from earlier attempts. The policy proposes an action block $A_t = (\hat{a}_{t|t}, \dots, \hat{a}_{t+H-1|t})$ of horizon H , where $\hat{a}_{t+\omega|t}$ is the action proposed at decision t for offset ω . With deployed neural parameters θ ,

$$A_t \sim \pi_\theta(\cdot \mid h_t, C_t). \quad (1)$$

Only executed actions enter the subsequent history; new observations can replace the unexecuted part of a proposed block. When C_t itself contains the state needed for future decisions, a memory update incorporates an action and its observed consequence:

$$C_{t+1} = U_\theta(C_t, o_t, a_t, o_{t+1}). \quad (2)$$

Equation (2) updates a self-contained context state. For context retrieved or compressed from experience, Equation (17) distinguishes the stored memory from the evidence supplied to the policy.

The policy interface is compatible with several architectures. Throughout the survey, θ denotes the collection of deployed neural parameters, including context interpretation, action generation, and any intervening




Figure 5 The workflow of ICL for robots. Prior training supplies competence; deployment acquires evidence, constructs context, and generates actions through four method families. Outcome assessment provides online feedback to execution, while retained experience informs subsequent context construction.

learned model. A subscript θ on a component means that it uses the relevant subset of this collection; different components need not share weights. During fixed-parameter adaptation,

$$\theta_{t+1} = \theta_t = \theta_0, \quad h_{t+1} = (h_t, a_t, o_{t+1}). \quad (3)$$

The subscript 0 denotes the start of the deployment adaptation phase. Observations, retrieved examples, recurrent state, or stored experience change instead. Equations (2) and (3) use one executed action per update; if several actions of A_t are executed before replanning, their observed transitions are incorporated in order. A reward or other feedback signal can be included when available. Updating a neural component at deployment constitutes parameter adaptation, even when the rest of the system stays frozen.

Figure 5 follows task information from prior competence to experience reuse. In a packing task, prior competence supplies grasps and motions. Evidence acquisition provides an instruction or demonstration specifying the part order. Context construction selects the relevant steps and relates them to the current scene. Contextual execution turns that evidence into actions; outcome assessment checks the resulting placements. Experience reuse makes verified outcomes and corrections available to later decisions. The four method families differ in how contextual execution converts evidence into control, while the surrounding operations determine which evidence it receives and whether the resulting behavior succeeds.

A corrected order changes the required actions; recorded placements identify the remaining work. Both can affect a fixed policy, but the retained information and its reset rules determine how long the change persists.

A fixed model can revise a goal, skill sequence, or program after feedback; an executor turns that specification into actions. For example, a language model can use prior errors to adjust a motion primitive’s settings [9]. Retaining the revised program lets later attempts reuse the correction without training neural weights. This stores task information in executable content, whereas fine-tuning stores it in model parameters.




Figure 6 Five locations and timescales of adaptation. In (a–c), supplied context, recurrent state, or retrieved episodes condition a policy with fixed neural parameters. In (d), fast weights change while the base remains fixed; (e) retains updated weights for later tasks. Blue borders and snowflakes mark fixed parameters; red borders mark updated parameters, with flames indicating the update stage.

The surrounding execution system, or *harness*, assembles context, exposes robot tools, executes requests, and returns feedback. Normally, this machinery remains fixed while the supplied evidence and generated requests change. A learning system can also revise its retrieval rules, context-construction code, or executable controller. We call this *model-external adaptation*: experience changes the machinery connecting a model to the robot. Identifying this revised component distinguishes changes in the execution system from changes in the neural model.

Figure 6 distinguishes three carriers used by fixed models: supplied evidence in (a), recurrent state in (b), and retrieved episodes in (c). The lower row introduces parameter changes: fast weights adapt during deployment in (d), and updated weights persist into later tasks in (e). Storage and retention timescale are independent: an external memory can persist without training, and a test-time update can be either reset or retained. Continual learning studies acquisition and retention across changing experience [143].

These distinctions determine the evaluation controls. Removing supplied evidence tests (a); clearing recurrent state tests (b); suppressing episode retrieval tests (c). Restoring pre-adaptation weights tests the contribution of parameter updates in (d–e). Section 7 develops these interventions under matched task conditions, so gains can be attributed to the information or parameters that actually changed.

One-shot specifies the amount of teaching, not its storage mechanism. Visual meta-imitation uses one demonstration for task-specific gradient updates [7]. RoboTTT restricts updates to a neural subset during long-context execution [10]; WAM-TTT adapts video-side fast weights from target-domain human video before rollout [11]. Both follow the selective-update principle in (d), despite different update timing. Thus one demonstration can support either fixed-weight conditioning or parameter adaptation. The next taxonomy distinguishes how the resulting context reaches robot actions.

2.3 Four method families and shared design dimensions

The primary taxonomy follows the *deployed path from context to action*. The four families below are alternative computational interfaces, distinguished by the intermediate that execution consumes. We then separate this classification from the cross-cutting roles, modalities, and storage of evidence, before resolving boundaries between hybrid implementations.

- **Context-conditioned policies** infer actions from task evidence, including demonstrations, language, and interaction history as supported, through a learned generator or local action retrieval.




Figure 7 Four context-to-control mechanisms: (a) action inference, (b) motion transfer, (c) future prediction, and (d) skill/tool selection. Surrounding works locate representative methods; † marks fast-weight adaptation in RoboTTT and WAM-TTT [10, 11].

- **Geometric demonstration transfer** reconstructs a demonstrated motion or contact reference, aligns or retargets it to the current scene, and executes that reference through tracking or replay.
- **World-model-based control** predicts future observations, states, or transitions and uses those predictions to decode, plan, or select actions.
- **Skill- and agent-based execution** infers skill choices, programs, or robot-tool requests for a separate executor; VLM and LLM agents form one branch.

These distinctions concern deployed computation, building on the mappings, system models, and plans distinguished in learning from demonstration [122]. They also preserve the separation between representation, optimization, and training objective emphasized in meta-learning [144]. Transformers, diffusion models, and VLMs each support multiple interfaces. Classification follows their deployed computation: hybrid systems can combine predicted futures with skill execution, or programs with reference tracking.

Four information roles explain what this context contributes. *Task specification* identifies the intended goal, relation, or procedure. *Correspondence* relates demonstrated objects, contacts, phases, or commands to their executable counterparts. *Physical response* describes how the current body and environment respond to action. *Execution state* records the scene, relevant hidden events, and completed prerequisites. A demonstration can specify a procedure and its contact geometry; an interaction can update both physical response and execution state. These roles guide evidence selection within every method family.

Language can state a goal, illustrate a convention, revise a procedure, or summarize an earlier event. Images and sensor readings can supply the same information roles, including through numerical states serialized as text. The learning question concerns what the evidence lets the robot infer; its encoding determines how that information reaches the model. Thus modality describes how evidence is encoded, while the learning mechanism explains how that evidence changes a decision.

Data acquisition and adaptation describe two further dimensions. *Teleoperation*, *UMI*, *egocentric video*, and *simulation* identify how records are collected (Section 4.1); visual, proprioceptive, contact, and action channels specify their sensorimotor contents. Language can accompany any source as instructions, examples, feedback, or memory; it is a cross-cutting context modality. *Contextual adaptation* changes the information or external artifacts used by fixed neural components, whereas *parameter adaptation* updates a neural component. Application demands describe what new evidence must resolve: task intent, unfamiliar spatial structure, physical response, or their combination. Thus a method has a control interface, uses evidence with identifiable roles, and operates under a stated update regime.

These dimensions let the same demonstration support different control computations. It can condition action inference, define a motion reference, constrain a predicted future, or specify a skill request (Figure 7). What distinguishes the families is the intermediate required by execution. How that evidence was acquired or remembered remains a separate design choice, which is why retrieval and memory can support several families.

Storage adds another independent choice. A context window, recurrent state m_t , or external archive M_t can supply the same task information with different access costs and persistence. Static API documentation establishes the command interface; task-specific examples and interaction supply changing evidence. For later comparisons with test-time training, write $\theta_t = (\bar{\theta}, \phi_t)$: $\bar{\theta}$ is the frozen backbone and ϕ_t the adaptive neural subset. This follows the update-regime distinction in Section 2.2.

The boundary between policy conditioning and geometric transfer is the object being executed. Matching points to condition a diffusion policy remains context-conditioned policy inference; transferring a recorded object-relative trajectory and tracking it is geometric demonstration transfer. Similarly, estimating current dynamics for an action policy differs from predicting future consequences for control. RMA and the physical-context inference in ICWM illustrate the former [39, 41]; DOMINO and Zero-WAM illustrate the latter [19, 24]. An auxiliary future-prediction head contributes to training; a deployed future predictor contributes to world-model-based control when its output guides action generation or selection.

The policy and agent interfaces can use the same VLM backbone. A VLM that emits motor-action tokens can implement a context-conditioned policy. A VLM that selects robot tools, supplies their arguments, and interprets returned execution feedback implements skill- and agent-based execution. GPT-Policy’s Astra-based configuration belongs to this agent branch [94]. Demonstrations and interaction feedback supply the task-specific evidence for its tool choices and arguments.

Memory connects these interfaces across time. Alongside its form, function, and dynamics [145], a retained fact carries physical conditions that determine when it remains applicable to the current robot and scene. Section 3.5 examines access and update mechanisms, and Section 5 examines how execution validates or revises retained evidence.

Intent adaptation changes the desired relation; *physical adaptation* changes its realization under new dynamics or embodiment. *Joint adaptation* requires both. Navigation additionally needs spatial knowledge even when the goal and dynamics remain unchanged. Section 6 organizes applications by these information needs and their implications for long procedures and task changeover. Any of the four computational families can serve several demands.

Table 2 links each control intermediate to its transfer requirement, then separates the mechanisms shared across interfaces. Matching a handle establishes correspondence; that match may condition an action generator or define an explicit motion reference. The execution interface determines the transfer assumption carried into the final motion. Organizing methods by that interface connects shared evidence-processing mechanisms to the physical behavior they support.

2.4 Historical development

The historical account follows an expansion in what must be inferred from context. Task-family adaptation learned reusable responses to bounded variation; sequence interfaces made more of the preceding

Table 2 Two levels of the taxonomy. (A) Four method families differ in their execution intermediate and a key requirement for transfer. (B) Shared mechanisms establish correspondence and manage evidence within or across these interfaces. A system can combine several interfaces and shared mechanisms.

A. Control interfaces: how context determines actions		
Method family	Intermediate	Transfer requirement
Context-conditioned policies	Action distribution	Action inference preserves the taught distinction
Geometric demonstration transfer	Motion or contact reference	Matched interaction remains applicable
World-model-based control	Predicted consequences	Predicted evolution is realizable
Skill- and agent-based execution	Execution specification	Available skills preserve task constraints
B. Shared mechanisms: how evidence supports those interfaces		
Mechanism	Information resolved	Example
Geometric and functional correspondence	Where to interact	Handle contact on a new object
Semantic correspondence	Which roles and order	Cup as source; bowl as receiver
Temporal alignment	Which demonstration phase	Current grasp phase
Retrieval	Which experience applies	Relevant past demonstration
Evidence retention	What remains available	Earlier observation or outcome
External knowledge reuse	What guides later attempts	Repair guidance or reusable skill

experience accessible; broader priors increased the behaviors that new teaching could direct. These developments overlap across the four families. The final publication census locates their recent growth within the literature reviewed here.

Task-family adaptation. Early contextual policies amortized adaptation across a task distribution: training learned how to interpret evidence so that deployment could change behavior without relearning the entire controller. In one-shot imitation, paired executions taught how one instance specifies actions in another [1]. In meta-RL, recurrent state carried adaptation across episodes [2]; probabilistic task inference separated evidence interpretation from control [4] and made task uncertainty available to action selection [334]. The resulting adaptation rule inherits the distinctions made available during training. Extending it to human demonstrations requires correspondence across bodies [8]; extending it to changed robot dynamics requires evidence about action response [39]. These extensions concern different unknowns, despite sharing rapid adaptation as their objective.

The four branches in Figure 8 respond to different obstacles in reusing experience. Context-conditioned policies learn which actions an example implies; geometric transfer identifies a relation under which a recorded interaction remains applicable. Predictive control makes consequences available before motion, while skill and program interfaces separate a new procedure from its motor realization. Their histories overlap because the obstacles coexist. Broader priors expand the available behaviors, but also make evidence selection and retention more consequential: a capable controller still needs to follow the current teaching rather than a familiar default.

Online system identification provided an explicit early version of the physical-adaptation idea: recent state–action history identified dynamics parameters for a universal policy [146]. Later methods could learn a latent instead of naming every parameter, but the central problem remained inferring how the present body would respond. The demonstration and interaction strands were already connected by Watch–Try–Learn: one policy attempted a demonstrated task and another used the resulting feedback [147]. A demonstration narrows the task hypothesis, while interaction can resolve information it omits.

Reusable sequence interfaces. A sequence interface addresses a tension between compressing experience and revisiting a decisive event. SNAIL combined temporal convolutions, which aggregate the preceding sequence, with attention, which retrieves relevant earlier information [3]. This made adaptation depend on both

	(a) Context-conditioned policies	(b) Geometric demonstration transfer	(c) World-model-based control	(d) Skill- and agent-based execution
2016–2019	- RL² - UP-OSI - One-Shot IL - SNAIL - TEC - PEARL	—	—	NTP
2020–2024	- TOSIL - RMA - MOSAIC - VIMA - Vid2Robot - ICRT - RoboPrompt - Instant Policy	- FlowControl - DOME - Interaction Warping - RoboTAP - DINOBot - DITTO - ImOP - MAGIC	- DOMINO - MetaDiffuser	- SayCan - Inner Monologue - Code as Policies - ProgPrompt - VoxPoser - XSkill - DROC - ORION
2025–2026	- RICL - MemoryVLA - LocoFormer - MEM - ICWM - MatchingPolicy - S1 - GEN-1.5 - GLOW - RoboTTT	- FUNCTO - ODIL - Annotation-Free OSIL - HRT1 - MT3 - ManiLong-Shot - SemAnCorr - Sparse Meets Dense	- OSVI-WM - Demo-JEPA - ReCAP (WAM) - WorldScape 2.0 - HOST - Zero-WAM - WAM-TTT	- UniSkill - Act-Observe-Rewrite - ASPIRE - ARCHITECT - SHAPER - Zetta - Show-Harness - GPT-Policy - Robo-Harness K1 - RACaP - WAA

Figure 8 Evolution of the four method families by first public release period. NTP: Neural Task Programming.

accumulated state and selective access. Transformers provide query-dependent access through attention [231]. Retaining more history increases the available evidence; attention and downstream computation determine which evidence informs the current decision.

Broader teaching made a fixed task summary less sufficient: two demonstrations can have the same goal yet differ in order, contact, or tool use. Multimodal task specifications placed goals and procedures in the input [100, 148], while trajectory models retained earlier observations and actions as evidence for later decisions [5, 6]. The benefit of a sequence interface was access to distinctions that a single task label could erase. Its cost was an alignment problem: the robot had to recover which part of that evidence mattered at its own execution stage. Temporal video correspondence and causal sensorimotor modeling addressed this pressure through learned representations [13, 14]; spatial representations imposed geometric correspondence before action prediction [16]. These complementary approaches preserve temporal evidence and geometric correspondence at different stages of action generation.

Adaptation over broad priors. Broad pretraining changes what adaptation must accomplish. When perception and a useful range of motion are already available, a new example can specialize an action model [17] or constrain an anticipated robot-domain future [24], rather than supply every ingredient of a new skill. Interaction histories similarly let reusable controllers account for a deployed body’s response [40, 88]. Broader pretraining therefore expands the behaviors available to teaching, but also strengthens familiar continuations that may conflict with a new example. Evaluation must establish whether the example can redirect this competence.

Geometric reuse and learned contextual inference allocate generalization differently. Visual nearest-neighbor methods rely on a representation in which nearby observations admit compatible actions [149]. Explicit alignment instead restores the conditions under which a recorded robot motion or human-demonstrated object trajectory remains usable [34, 35]. When those conditions are insufficient, a video can initialize further interaction learning, as in WHIRL [150]. The same one-example interface can therefore support replay, transformed reference execution, or new policy learning. What distinguishes these routes is the assumption that makes the demonstrated behavior executable in the new setting.

World-model-based control exposes the consequences of behavior before executing it. Contextual dynamics models first made recent transitions useful for evaluating candidate action sequences [19, 20]. Demonstration-




Figure 9 Composition of the reviewed literature. (a) First-publication years for all 412 references. (b) The four-family distribution of 260 method and comparison references. Percentages in (b) use this subset, which includes supporting mechanisms and parameter-adaptation comparisons. The first bin covers publications through 2015; † marks the partial year ending 29 September 2026.

conditioned models then added a second role for context: specifying which future should be realized [21, 24]. This development joins physical prediction with task interpretation. The controller must preserve the procedure inferred from the example while predicting how it can unfold in the robot’s current scene.

Task structure offers a different way to reuse competence: preserve the procedure while delegating its physical realization. Demonstration-derived hierarchies express a new task through existing routines [26]. Language-model programming expands the procedure to include perception calls, arithmetic, and feedback logic [27]; video-derived outlines let visual teaching specify that procedure for an existing planner and executor [31]. The gain is compositional reuse of capabilities. Its condition is that the intermediate preserve the task’s binding details, because the executor cannot recover a contact or ordering requirement that was omitted during interpretation. Section 3 develops this trade-off across the four interfaces.

These developments connect teaching, interaction, and procedural reuse. General-Level distinguishes breadth from synergy in multimodal generalists [151]; for robots, the corresponding learning question is whether combining these capabilities improves subsequent acquisition. LMPC studies responsiveness to later teaching [117], while reward and learning-program search change how new skills are acquired [89, 115]. Section 8.5 develops this question for individual and shared learning.

Review scope and source selection. This review uses a mechanism-oriented narrative synthesis with targeted updates through 29 September 2026. Searches and citation chaining span visual imitation, contextual meta-RL, navigation from video and experience, system identification, example-conditioned programming, interaction memory, and robot foundation models. Sources include arXiv, official proceedings and journal articles, and author or organization releases. Technical claims are traced to these primary sources; release reports are identified as such where their evidence differs from a research paper.

Works enter the synthesis when they explain a route from new evidence to behavior, supply a relevant data or evaluation resource, or establish a training or theoretical foundation. Fixed-weight contextual methods form the central subject; parameter adaptation and prior-competence methods supply explicit comparisons. Versions of the same work share one bibliography entry, using an official publication when available. Separately cited papers and dataset releases remain distinct references. For hybrid systems, the analyzed deployment branch determines the primary family; modality, storage, and parameter updates remain independent attributes.

The reference inventory in Appendix A records this coding and the first public year, using arXiv v1 or an official release. Supplement S1 provides the item-level records and counting rules. The resulting counts describe the composition and publication history of this reviewed corpus.

Publication patterns. The 412 references include 297 first released since 2024 and 197 in the partial 2026 window. The concentration coincides with the recent expansion of generalist policies, multimodal teaching, and embodied agents. Figure 9 separates this complete reference set from the 260 method and comparison references; the latter includes supporting control mechanisms and parameter-adaptation comparisons.

Within the method subset, 123 references concern context-conditioned policies, 24 geometric transfer, 22 world-model-based control, and 91 skill- and agent-based execution. Branch-specific coding places ReCAP's Cosmos Policy branch under predictive world-action modeling [152], ENPIRE under program revision [89], and GPT-Policy under contextual tool use [94].

The recent concentration of publications and the distribution across control interfaces in Figure 9 motivate a closer comparison of how methods use context. The following chapter examines their execution intermediates and the assumptions that support transfer.

Takeaways

- ① Context resolves what the current observation and prior competence leave undetermined: the required task, relevant history, or physical response. Its value depends on whether that information changes a physically executable decision.
- ② Demonstration following and physical adaptation have distinct historical roots. Their convergence joins inference of a new task with inference of an unfamiliar body's response.

3 Methods: How Does Context Change Action?

Context changes action through an execution intermediate: an action distribution, a motion reference, predicted consequences, or a skill or tool specification. The four families are compared by what this intermediate preserves, why existing competence can execute it, and which changes invalidate that assumption (Table 2). Correspondence binds the evidence to the current situation; memory keeps it available; external revision carries useful lessons into later attempts. Together, these mechanisms connect evidence processing, execution, and reuse across the four families.

The factorizations use the context C_t , history h_t , proposed action block A_t , and fixed neural parameters θ defined in Section 2. Each exposes the intermediate consumed by execution, allowing stochastic or deterministic implementations. Table 1 collects the notation, with local symbols defined alongside each equation.

A shared primitive vocabulary. Tables 3, 4, and 5 use one primitive taxonomy. *Source* codes demonstrations or examples (D), retained history (H), instructions or goals (I), explicit feedback (F), and scene or map evidence (E). *Form* codes visual observations or features (V), state/action records (S), geometry (G), and linguistic or symbolic records (L). *Context operation* separates eight operations: *Retrieve* selects stored records; *Attend* makes a query-dependent neural read; *Align* establishes spatial or temporal correspondence; *Encode* builds a learned representation; *Estimate* infers hidden physical response; *Parse* constructs symbolic bindings or procedures; *Verify* checks candidates against feedback or constraints; and *Fit* solves query-specific coefficients or geometry with network weights fixed. The tables code task-relevant contextual evidence and the principal operations that use it during deployment. Commas list evidence types; slashes list operations or carriers used together, without imposing a temporal order.

The resulting *carrier* is an action output, motion reference, predicted future, learned skill, program/procedure, or tool/controller call. A carrier identifies what passes between components; the family identifies its role in the complete control path. *Execute* then distinguishes six operations: *Replay* returns recorded controls; *Combine* blends retrieved action continuations; *Decode* generates actions with a learned model; *Track* follows an explicit reference; *Plan* searches for an executable realization; and *Invoke* runs a program or calls a tool.

Evidence and execution roles. Language contributes according to the information it carries, independently of the control interface. In direct action inference, numerical states and action examples can be serialized for a language model, as in RoboPrompt [153]; ICRT learns its sensorimotor interface without linguistic data [13]. For a geometric archive, language can narrow the intended task before geometric matching resolves the motion, as in MT3 [36]. Command-program examples instead teach an executable mapping [27], while textual event memory preserves information missing from recent visual history [44]. These uses resolve task

selection, action interpretation, and retention at different stages. The relevant comparison is what the text contributes to the decision, including whether it supplies a goal, a new convention, or a remembered outcome.

The separation between information and execution also applies to geometry. Matching can condition a learned action generator or supply a retargeted reference to a separate motion executor. Table B1 in Appendix B (p. 97) compares these roles. Geometry can instead supervise further learning: using hand-object motion to guide residual reinforcement learning trains a new policy [188]. The role of geometry becomes clear by identifying what it changes: a policy input, an executable reference, or the policy parameters.

Semantic abstraction preserves object roles, prerequisites, and skill order while allowing grasps or paths to change. Object-centric plans and video outlines expose that structure at different resolutions [31, 165]. Their executor must receive every binding constraint: plate-first teaching specifies order, while a handle-grasp requirement additionally restricts contact. A generic pick-and-place request can erase either distinction. This motivates tracing the taught requirement through the intermediate to the command it changes.

3.1 Context-conditioned policies

Context-conditioned policies map task evidence and robot history to actions. They either reuse recorded actions selected by correspondence or generate actions from an interpreted context representation. We compare these two forms before examining three complementary demands on the representation: execution phase, spatial correspondence, and physical response. The final comparison considers how pretrained competence affects the evidence each form needs.

For a context-conditioned policy, write r_t for the representation of the evidence used at the current decision:

$$r_t = E_\theta^c(h_t, C_t), \quad \pi_\theta(A_t | h_t, C_t) = q_\theta^a(A_t | h_t, r_t). \quad (4)$$

Here E_θ^c interprets context and q_θ^a generates actions from its representation r_t ; both may belong to one network. The factorization requires r_t to retain the context information used by action generation. Autoregressive, diffusion, and flow-based generators can implement this interface. If an example teaches red-then-blue packing, the representation must preserve that order while the generator adapts the reach-and-place motion to the current scene.

Reusing actions through correspondence. Let the nonempty set J_t index the demonstration records selected for the current query, and let $\bar{A}_j$ be the action block stored in record j . A nonparametric action-reuse policy has the form


$$\pi_\theta(\cdot | h_t, C_t) = \sum_{j \in J_t} \beta_{tj} \delta_{\bar{A}_j}, \quad \beta_{tj} \geq 0, \quad \sum_{j \in J_t} \beta_{tj} = 1. \quad (5)$$

Here β_{tj} is the query-dependent weight of record j , and $\delta_{\bar{A}_j}$ is a point mass at its stored action block. Nearest-neighbor replay selects one record with weight one. A policy that averages retrieved controls instead returns $A_t = \sum_{j \in J_t} \beta_{tj} \bar{A}_j$; this is a deterministic action combination rather than sampling a stored block. Such averaging requires compatible continuous action coordinates and separate treatment of discrete commands. Stored blocks must be expressed or retargeted in the receiving robot’s action interface. Equation (4) instead uses a learned generator q_θ^a to transform interpreted evidence into actions. When retrieval supplies either interface, it selects the evidence; only the replay interface directly returns recorded actions.

Action reuse assumes that proximity in a representation implies compatible control. The workflow encodes the current scene, retrieves comparable demonstration states, and returns their recorded actions [149]. Its generalization rests on representation quality and archive coverage. A retrieved demonstration can instead condition a learned generator [17], which must infer how the example applies rather than return its action directly. This introduces a second source of generalization: transforming available experience beyond local replay.

Distribution shift can separate these two sources. In REGENT’s unseen-MuJoCo-embodiment comparison, Retrieve-and-Play outperforms the frozen contextual transformer, while fine-tuning improves the learned agent [154]. The retrieved actions remain useful in a condition where the learned interpreter transfers less

(a) Context selection and interpretation



(b) Task specification




Figure 10 Constructing and interpreting task context. (a) Retrieval selects relevant experience, temporal alignment identifies the current phase, and observed action effects inform physical response. (b) Language specifies an operation, while a visual prompt makes its spatial target explicit. These operations supply task evidence to the execution interfaces compared in Figure 12.

effectively. Archive coverage and learned transformation should therefore be compared separately: retrieval can recover an applicable motion even when a more expressive generator does not preserve its relevance.

Local continuation bridges replay and learned generation. Behavior Predictive Control (BPC) reconstructs the live observation–action history from retrieved trajectory windows, then applies the same coefficients to their action continuations [391]. Unlike the nonnegative mixture in Equation (5), its regularized coefficients may be signed, permitting local extrapolation. The retrieval metric and residual correction are fitted offline; deployment solves the continuation coefficients from the current history while keeping those fitted components fixed. Reliable continuation depends on whether nearby records jointly explain the query with stable coefficients, connecting retrieval quality to local control compatibility.

The use of correspondence determines which component carries motion generalization. A retargeted human motion can serve as a proposal that a pretrained diffusion policy refines, as in DemoDiffusion [155]. The learned prior then contributes to the final motion. Alternatively, registration can directly transfer the next demonstrated end-effector pose, as in IMOP [156]. In that case, geometric applicability carries more of the transfer burden. Section 3.2 develops this second interface, where the matched evidence becomes an explicit reference for execution.

Context interpretation can resolve three kinds of uncertainty before action generation: which experience applies, which execution phase matters, and what the observed transitions reveal about physical response (Figure 10a). Task specification adds a complementary constraint: language can identify an operation, while a visual prompt binds it to an explicit spatial target (panel b). These shared operations organize the evidence available to each method family. Their output can condition direct action inference or supply task information to a geometric reference, a predictive model, or a skill executor. The four families differ in this subsequent context-to-action interface, illustrated in Figure 12.

From task identity to execution correspondence. When retrieved actions do not transfer directly, the learner must interpret what the example specifies. Early demonstration-conditioned policies compressed examples into a task representation. Task-Embedded Control Networks jointly learned a demonstration metric and an embedding-conditioned policy [157]. Such a summary makes few-shot conditioning economical, but the policy must recover the relevant execution phase from its current observation. As variation extends from task identity to ordering, viewpoint, and object placement, identifying the task is no longer enough: the robot must relate what it sees now to a particular part of the example.

Table 3 Context-conditioned policies: retrieval, demonstration interpretation, and physical adaptation (Section 3.1). Neural components remain fixed during the listed deployment branches. Codes: D examples, H history, F feedback; V visual, S state/action, G geometry, L symbolic. Slashes combine operations. †: technical report; −: unspecified. Operations are defined in Section 3.

Method	Source	Form	Context operation	Carrier	Execute
<i>Retrieval and trace guidance</i>					
VINN [149]	D	V,S	Retrieve	Actions	Replay
BPC [391]	D	V,S	Retrieve / Fit	Actions	Combine
RICL [17]	D	V,S	Retrieve / Encode	Actions	Decode
TraceFlow [392]	H,F	V,S	Retrieve / Align	Actions	Decode
<i>Demonstration interpretation</i>					
TOSIL [21, 158]	D	V	Attend	Actions	Decode
MOSAIC [12]	D	V	Attend	Actions	Decode
RoboPrompt [153]	D	G,S	Parse	Actions	Decode
BiCICLe [92]	D	S	Parse	Actions	Decode
Vid2Robot [14]	D	V	Attend	Actions	Decode
ICRT [13]	D	V,S	Attend	Actions	Decode
KAT [15]	D	G,S	Encode	Actions	Decode
Instant Policy [16]	D	G,S	Encode	Actions	Decode
MatchingPolicy [161]	D	G,S	Align / Encode	Actions	Decode
StellaVLA [163]	D	L	Encode	Actions	Decode
SynthICL [77]	D	V	Encode	Actions	Decode
BPP [54]	D	S	Encode	Actions	Decode
<i>Physical-response inference</i>					
RMA [39]	H	S	Estimate	Actions	Decode
Self-Adaptive VLA [393]	H	V,S	Estimate	Actions	Decode
<i>Generalist demonstration interfaces</i>					
GEN-1.5† [98]	D	S	−	Actions	Decode
S1† [99]	D	V	−	Actions	Decode
GLOW† [389]	D	V	Encode	Actions	Decode

Table 3 traces shared context operations into distinct action rules. Retrieval can return actions directly, condition a generator, or constrain a predicted future. Alignment can condition action inference or construct a motion reference. Transfer depends on what the carrier preserves and how execution realizes it; the following comparisons develop those assumptions.

An attention-based interpreter implements a query-dependent read over demonstration elements. Let $D^s = (d_i^s)_{i=1}^L$ be a support demonstration with L elements. Its i th element d_i^s may contain an image or a synchronized observation–action record; k is a summation index over the same elements. One representative alignment is

$$\alpha_{ti} = \frac{\exp s_\theta(h_t, d_i^s)}{\sum_{k=1}^L \exp s_\theta(h_t, d_k^s)}, \quad r_t = \sum_{i=1}^L \alpha_{ti} e_\theta(d_i^s). \quad (6)$$

The score s_θ matches the robot history to an encoded demonstration element $e_\theta(d_i^s)$. Normalized weights α_{ti} select evidence relevant at decision t , allowing the read to change with execution progress. This inference-time read supports image-only and action-bearing demonstrations alike; its training objective is a separate choice considered in Section 4.2. Richer attention or geometric representations can preserve several relationships without imposing a monotone frame match.

Execution-dependent attention addresses the loss of phase information in a single task summary: each action can consult the relevant spatial and temporal evidence [158]. Cross-domain contrastive supervision makes that read more reliable when demonstration and query differ visually [12]. Video cross-attention

and causal sensorimotor sequences expose different evidence to the read, respectively visual progression and observations with recorded actions [13, 14]. Execution mismatch creates a further problem when frame similarity is misleading. RHyme constructs semantically compatible training videos through clip retrieval and composition, then uses learned alignment to condition its policy [80]. These designs address complementary parts of correspondence: when to read, what to compare, and how training makes the comparison meaningful.

The context representation determines which distinctions action generation can recover. RGB retains appearance and temporal evidence but leaves motion implicit; compatible sensorimotor records expose the action itself. Keypoints and spatial graphs supply correspondence before generation [15, 16], reducing the need to infer task geometry from appearance. Motion latents pursue a different compression: cross-domain decoding ties human and robot movement together [159], while hierarchical tokens separate persistent demonstration structure from finer execution decisions [160]. Continuous flow matching generates actions from encoded images, proprioception, and commands [55]. TP-Flow additionally extracts phase-level prototypes from support demonstrations and conditions both the initial flow prior and velocity field on them [386]. Episodic training teaches this dependence; deployment recomputes prototypes with fixed weights. Representation and generator are therefore coupled choices: richer action generation depends on preserving the taught distinction during context encoding.

Temporal selection addresses a different limit from action encoding. A complete trajectory may contain many locally irrelevant stages, so ICI-VLA retrieves phase-aligned micro-demonstrations using a retriever trained with dynamic-time-warping supervision [18]. The deployed policy reads examples selected from observable query inputs; target-action masking during training reduces reliance on exact action-prefix continuation. Phase selection determines which evidence reaches the generator, whereas token prediction or flow matching determines how that evidence becomes an action block.

Spatial correspondence as conditioning. Temporal phase identifies when an example applies; spatial correspondence identifies where its action should be realized. MatchingPolicy makes this second requirement explicit [161]. Foundation-model features establish demonstration–query point matches; tracking propagates them through time; RGB-D lifts them into 3D. A graph diffusion policy then generates gripper motion from these relationships. This workflow assumes that matched task geometry remains informative about the required action, without fixing a trajectory for the robot to replay. Section 3.2 considers the complementary case in which correspondence transforms an explicit executable reference.

Navigation extends this correspondence from object interactions to places. A policy can read successive visual subgoals [355] or attend to an action-bearing path memory [356] to follow new teaching. A scene preview supplies a different spatial reference: NOLO [351] combines a video with optical-flow-derived action labels, the current view, and a goal image to predict actions. Its offline reinforcement learning trains the interpreter; deployment reads a new scene video with weights fixed. The shared computational interface is contextual action inference, while the context can specify a route or reveal an environment.

Identifying physical context from interaction. Temporal and spatial correspondence can identify a desired motion while leaving its physical effect uncertain. Interaction history supplies the complementary evidence when the goal is known but the current body’s response is not. System identification makes this unknown explicit through physical parameters, whereas a learned adaptation latent can retain the response information needed by a policy [39, 146]. ICWM extends physical-context inference to changes in observation and embodiment [41]. These approaches require transitions that reveal how commands affect the present system; a task demonstration alone may leave that uncertainty unresolved. Their representation conditions an action policy directly. Predicting and evaluating candidate futures introduces the additional computational stage considered in Section 3.3.

The observation channels determine whether such a response is identifiable. Self-Adaptive VLA pairs rollouts under injected hardware shifts with compensated expert targets during post-training, then aggregates rollout-derived context tokens to modulate a fixed deployed policy [393]. Under its actuation-bias model, proprioception exposes a mismatch between commands and realized joints; under joint-encoder offsets, readings can remain consistent with commands while the physical pose is wrong, requiring visual evidence. Additional history helps only when it contains a measurement that distinguishes the possible causes. The inferred context calibrates physical response for the specified task.

Connecting pretrained competence to new evidence. Broader robot and language priors change what must be learned for demonstration following. An existing action generator can be conditioned on retrieved executions [17]; grounded spatial descriptions can supply scene information missing from the current visual input [162]. These inputs resolve different unknowns. A demonstration may change the requested procedure, while a scene description may clarify how to execute an already known procedure. The contribution of each channel depends on the decision its information resolves.

Observation-only teaching leaves action correspondence implicit. Behavioral Cloning from Observation estimates actions through inverse dynamics before training a policy [167]. Fixed-weight teaching uses an existing action interface learned from demonstrations or constructed through pretrained reasoning and geometry.

Textual action interfaces reveal how much of this burden can be handled outside a visuomotor backbone. Object poses and keyframe end-effector actions can be supplied as examples to a frozen language model [153]. For bimanual tasks, conditioning one arm’s trajectory on the other’s prediction preserves inter-arm dependence while reducing the joint prediction problem [92]. BiCICLE’s main TWIN experiments use ten demonstrations per task, with iterative refinement and candidate selection adding inference computation. These systems reuse pretrained reasoning, but rely on perception, coordinate conventions, and action representations that make its outputs physically meaningful.

Teaching budget can change the preferred adaptation mechanism. KAT compares a text-pretrained model reading keypoint–action examples with KeyAct-DP, a diffusion policy trained on similar representations [15]. Contextual inference is competitive with few examples; policy training becomes stronger with more. The crossover differs against an image-based diffusion policy, indicating that both representation and adaptation mechanism determine the useful demonstration budget.

Generalist systems expose different teaching representations. BPP and GEN-1.5 accept sensorimotor prompts [54, 98]; S1 accepts video prompts [99]. The GLOW technical report [389] describes a Skill Demo Encoder that converts one human video into reusable skill context. A fixed-weight autoregressive model combines it with current observations, language, robot state, and history to generate end-effector actions. Its teaching pathway belongs to context-conditioned policies (Figures 5 and 7); Harness manages tools and feedback, while KnowinWorld supplies candidate rollouts. The primary family follows the teaching-to-action pathway, with prediction and tool management providing complementary support.

The defining trade-off is between compatibility and transformation. Retrieved actions retain an explicit link to experience but need compatible control coordinates and archive coverage. A learned generator can adapt the example more flexibly, provided its representation preserves the taught distinction. REGENT’s comparison locates this trade-off in the transfer of the learned interpreter under distribution shift [154]. Transfer depends on correspondence identifying usable evidence and on the generator preserving its relevance. Geometric transfer makes that dependency explicit by turning the evidence into a reference that execution must track.

3.2 Geometric demonstration transfer

Geometric demonstration transfer preserves a motion or contact reference and adapts its realization to the query scene. Perception establishes correspondence; alignment or retargeting transforms the reference; tracking or replay executes it. This explicit reference is useful when contact geometry matters, such as preserving a handle grasp. We follow the mechanism from reference construction through spatial alignment and shape changes to reuse across a task repertoire.

IMOP makes the reference local to the next demonstrated action: learned invariant-region correspondences support analytic pose transfer, and a state router selects the relevant demonstration transition [156]. The reference can therefore be a target pose rather than a complete trajectory.

Reconstructing an executable reference. Let D^s denote the supplied demonstration. A geometric transfer computation can be written as

$$\zeta_{1:L} = E_{\theta}^m(D^s), \quad g = \text{Retarget}(\zeta_{1:L}, o_t, \mathcal{B}), \quad A_t = \kappa(g, o_t). \quad (7)$$

Table 4 Geometric demonstration transfer, grouped by the reference preserved (Section 3.2). R-NDF lists descriptor alignment. Codes: D demonstration; V visual, S state/action, G geometry. Slashes combine operations. Tracking, replay, and planning denote different stages of realizing the reference.

Method	Source	Form	Context operation	Carrier	Execute
<i>Visual alignment</i>					
FlowControl [32]	D	V,G	Align	Motion	Track
DOMe [33]	D	V,S	Align	Motion	Track / Replay
DINOBot [34]	D	V,S	Retrieve / Align	Motion	Track / Replay
<i>Object-relative motion</i>					
MT3 [36]	D	G,S	Retrieve / Align	Motion	Track / Replay
Pose-based OSIL [170]	D	G,S	Align	Motion	Track
IMOP [156]	D	G,S	Align	Motion	Track
DITTO [35]	D	V,G	Align	Motion	Plan / Track
HRT1 [168]	D	G	Align / Fit	Motion	Track
<i>Functional correspondence</i>					
R-NDF [346]	D	G	Align / Fit	Motion	Track
Interaction Warping [171]	D	G	Align / Fit	Motion	Track
Part decomposition [347]	D	G	Align	Motion	Track
SemAnCorr [38]	D	G	Align	Motion	Plan / Track
<i>Staged and coordinated execution</i>					
ODIL [37]	D	V,S	Align	Motion	Track / Replay
ManiLong-Shot [93]	D	G,S	Align	Motion	Track
Annotation-free OSIL [177]	D	V,S	Align	Motion	Track / Replay

Here E_θ^m is the motion-descriptor extractor, with m identifying its role. It produces one descriptor ζ_i for each of the L demonstration elements, such as hand poses, object-relative poses, or contact locations. The operator Retarget retargets that sequence to the current scene o_t and robot description $\mathcal{B}$, which specifies kinematics and control constraints. Its output g is a robot motion reference consumed by executor κ . Object-centric transfer and hand-based retargeting differ in what ζ_i preserves: the motion of a manipulated object, or the demonstrator’s grasp and hand motion [35, 168]. Fitting a pose, grasp, or trajectory changes these inferred variables while the estimator’s neural parameters can remain fixed.

For rigid object-relative transfer, this relationship has a particularly simple form:

$${}^W T_{E,b}^q(\sigma) = {}^W T_{O_b}^q ({}^W T_{O_b}^s)^{-1} {}^W T_{E,b}^s(\sigma). \quad (8)$$

Here ${}^W T_X \in SE(3)$ is the rigid pose of frame X in world frame W ; E denotes the end effector, and O_b the reference object for motion segment b . The group $SE(3)$ contains three-dimensional rigid transformations, and the exponent -1 denotes the inverse transform. Superscripts s and q distinguish demonstration and query scenes, and $\sigma \in [0, 1]$ is normalized progress within that segment. The middle and right factors express the demonstrated end-effector motion relative to its reference object; the left factor places it in the query scene. This form assumes that the reference remains fixed during the segment and that its relative motion remains applicable. Shape change, moving references, and embodiment differences require the richer correspondence in Equation (7). Choosing O_b is therefore part of task interpretation, not simply a coordinate convention.

Table 4 traces four ways to preserve an interaction: align its visual reference, transform object-relative motion, match functional parts, or switch references across stages. A richer correspondence is useful when it preserves a constraint that a global pose transform would lose.

From image alignment to object-relative transfer. Geometric reuse begins with an applicability assumption: restoring the relevant relative configuration makes the demonstrated interaction useful again. Image-based approaches enforce this through visual feedback. FlowControl aligns live RGB-D observations to

successive masked reference frames [32]; DOME separates a learned visual servo to a bottleneck pose from subsequent velocity replay [33]. The difference is where feedback acts: throughout reference following or primarily before the interaction segment. Tracking active points and gripper events allows RoboTAP to change the visual reference as a task advances [169]. POIL closes the loop on the transferred object points themselves: functional-part correspondences establish a reference, and point-set feedback converts tracking errors into a rigid-body twist [405]. Using the same points for transfer and execution ties disturbance recovery to the demonstrated interaction.

An object-relative invariant makes applicability more explicit than image agreement. Pose-based transfer preserves a recorded end-effector trajectory [170]; category-level shape fitting extends correspondence to interaction keypoints [171]. DITTO reconstructs manipulated-object motion from an RGB-D human demonstration, transfers that trajectory to the detected objects in the robot scene, and realizes it through grasping and motion planning [35]. Hand-based retargeting instead reconstructs the demonstrator’s grasp and movement before converting them into robot configurations [168]. The embodiment gap thus appears in different places: realizing an object effect in DITTO, or translating a demonstrated hand configuration. In either case, successful transfer requires both a preserved interaction and a feasible realization.

Semantic anchors and functional geometry. Rigid pose transfer preserves a motion only when the source and target support the same local interaction. A different handle or articulated part requires correspondence within the object. SemAnCorr [38] selects semantically consistent surface anchors and propagates them through a functional map. Its manipulation pipeline transfers contact regions and local frames, selects a feasible grasp, and reuses motion relative to that grasp. Contact keyframes and relative motions form its execution specification. Semantic correspondence supplies the local geometry needed to extend the rigid transfer in Equation (8) across object shapes.

Shape variation makes whole-object similarity too restrictive for some interactions. What must remain compatible may be a grasp role and a contact role, rather than the complete tool geometry [172]. A complementary approach starts from global shape and refines local curvature correspondences to preserve the contact used by a motion [173]. Surface interaction functions and screw interpolation connect such functional geometry to a transferred trajectory [174]. These routes differ in how they identify the invariant, but share a control assumption: preserving the relevant interaction structure makes the adapted motion useful. This shifts the transfer problem from matching objects as wholes to matching the parts of their geometry that the task actually uses.

When both interacting objects change, correspondence must preserve their relation as well as locate each part. R-NDF infers local frames on task-relevant parts and aligns them to reproduce a relation on a new object pair [346]. Part decomposition further separates each part’s geometry from its arrangement within the object, selecting the part relationships that explain the demonstration [347]. This reduces interference from irrelevant shape variation while retaining the constraints needed for placement. These object-pair representations extend reference transfer beyond relocating a trajectory around one object; Section 3.5 compares the resulting transfer scope.

Deformable manipulation requires correspondence to remain meaningful as the object changes shape. Sparse Meets Dense extracts task-relevant contact pairs from a demonstration and uses an offline learned dense representation to track them during constrained execution [175]. Its final-state relative-pose annotation supplies interaction geometry; sparse constraints express the task, and dense correspondence maintains that geometry through deformation. This extends the reference from a rigid frame to a changing set of contact relations.

Affordance prediction can supply a related execution interface from a learned prior. VLAff predicts interaction heatmaps, grasp poses, and motion trajectories from an observation and language instruction [176]. Its human videos provide offline affordance supervision. These outputs can ground a contextual planner’s specification, while a demonstration-conditioned system additionally uses the newly supplied example to determine which interaction should transfer.

From one demonstration to a reusable repertoire. A reusable geometric repertoire separates two uncertainties: which recorded interaction applies, and how to enter its starting configuration. DINOBot couples visual retrieval with wrist-view alignment before replay [34]. MT3 scales this decomposition using language filtering, geometric retrieval, pose alignment, and recorded interaction motion [36]. Its comparison of retrieval and behavioral cloning treats alignment and interaction as separate phases; its repertoire reaches 1,000 single-interaction tasks in less than 24 hours of human demonstration time. Multi-stage tasks use a demonstration for each stage. The scalability comes from adding applicable references without retraining a task policy, while success still depends on restoring the conditions under which each interaction can be reused.

Independent object motion and arm coordination expose the limit of one global alignment. A grasp and a later insertion must refer to different objects when the part and receptacle move independently; dual-arm interaction must additionally preserve relative coordination. Multi-stage visual servoing establishes a coordinated reference before ODIL’s replay [37]. Interaction segmentation with invariant-region matching, and VLM-based keyframe selection with repeated alignment, extend the reference-switching principle to longer demonstrations [93, 177]. Table 4 compares how individual systems bind and execute these references, while Table B1 in Appendix B contrasts the computational roles of correspondence. The key decision is when a change of interaction requires a different reference.

Geometric transfer makes its assumption inspectable: the matched interaction must remain applicable after alignment. It shifts the burden from task-specific action generation to correspondence, reference switching, and feasible tracking. Shape variation, lost contact, and independent object motion challenge different parts of this assumption. Predictive control instead represents the consequences execution should produce.

3.3 World-model-based control

World-model-based control uses contextual evidence to predict task-relevant consequences. These predictions can specify a desired evolution for action decoding or evaluate candidate actions under the current dynamics. We compare these complementary uses, then examine demonstration conditioning and the distinction between deployment-time prediction and auxiliary training supervision.

A representative factorization is

$$\begin{aligned} v^+ &\sim p_{\theta}^v(\cdot \mid C_t, h_t), \\ A_t &\sim q_{\theta}^v(\cdot \mid v^+, C_t, h_t), \end{aligned} \tag{9}$$

where v^+ denotes anticipated robot observations or their latent representation; the superscript $+$ identifies a future relative to the current decision. The future distribution p_{θ}^v predicts that evolution, and the future-conditioned decoder q_{θ}^v generates actions from it; v identifies the future-prediction interface. All neural components use fixed parameters from θ in the ICL regime. The two distributions may share a backbone; joint visual-action generation couples them within world-model-based control [178]. For the packing example, v^+ can show the red item placed before the blue one in the robot scene; action decoding must realize that particular sequence. At the distribution level, the law of total probability combines the future-conditioned action distributions. Each is weighted by the probability of its corresponding future:

$$\pi_{\theta}(A_t \mid h_t, C_t) = \int q_{\theta}^v(A_t \mid v^+, C_t, h_t) p_{\theta}^v(v^+ \mid C_t, h_t) dv^+. \tag{10}$$

The integral averages the action distribution over possible futures; discrete futures give a sum. Implementations may sample a future and then an action, or select one future as an approximation. This dependence on a predicted consequence, rather than the choice of sampling algorithm, defines the interface.

Generating a desired future or evaluating an action’s future. World-model-based control has two complementary forms. In Equations (9)–(10), context specifies a desired evolution and an action decoder realizes it. In model-based selection, candidate actions specify possible evolutions, and context determines which consequence is preferred. One planning form is

$$A_t^* \in \arg \max_{\hat{A}_t \in \mathcal{A}_t} \mathbb{E}_{v^+ \sim p_{\theta}^d(\cdot \mid h_t, \hat{A}_t, C_t)} [\mathcal{R}(v^+, C_t)]. \tag{11}$$

The set $\mathcal{A}_t$ contains admissible candidate action blocks, $\hat{A}_t$ denotes a candidate, and A_t^* is a selected maximizer. The action-conditioned dynamics distribution p_θ^d estimates its future v^+ ; d distinguishes forward dynamics from the context-conditioned future generator p_θ^v . The score $\mathcal{R}$ measures how that future satisfies the contextual task requirement. The expectation averages over predicted futures. This form makes the two roles explicit: learned dynamics predicts what an action may cause, whereas teaching determines which consequence is desirable. Optimizing $\hat{A}_t$ leaves the neural parameters θ unchanged.

Unknown dynamics changes which candidate trajectories are credible. One response is to infer dynamics context from recent transitions and evaluate actions with a context-conditioned world model, as in DOMINO [19]. Another is to condition trajectory generation on those transitions and use reward and dynamics guidance during sampling, as in MetaDiffuser [20]. Both specialize prediction through new evidence while leaving the learned parameters fixed. Their difference is where feasibility and preference enter the computation: in evaluation of candidates or in their generation and refinement.

Online fitting offers a complementary response to uncertain dynamics. MetaPusher updates a meta-learned pushing model from task transitions and repairs its planning tree as predictions change [368]. Sandwich-Residuals instead fits small corrections around a frozen predictor using latent prediction error [367]. Both adapt neural parameters at deployment. Their comparison with contextual prediction therefore turns on what changes: model parameters or the evidence supplied to an unchanged model.

From goal-reaching to demonstration-conditioned futures. Latent planning supplied an early way to connect an intended state to control. Universal Planning Networks learned a differentiable planner that optimizes a candidate action sequence [179]. The optimized variables are actions, so this search can run with unchanged model weights. Demonstration-conditioned prediction adds a preceding inference problem: determine which state evolution the new example calls for, rather than assume that the goal is already supplied.

This inference can remain in a latent space. OSVI-WM predicts demonstration-conditioned latent states and decodes them into attributed waypoints [21]. Demo-JEPA instead passes a predicted latent goal to a cross-entropy-method planner, retaining that goal until an achievement threshold permits progression [22]. Both combine world-model-based control with an explicit waypoint or goal executor: the predicted future supplies targets for a separate execution stage, rather than only features inside an action head. In Demo-JEPA’s real-robot comparison, the diffusion action head has higher mean success on familiar behavior grounding, while planning has higher mean success on unseen tasks. The familiar-task advantage does not persist in its simulation comparison. The relative value of search and learned decoding depends on both task novelty and the execution setting, even when the two share a future representation.

Prediction can also supply a route or a discrete transition model. TADreamer reconstructs navigation waypoints from generated video and registers them to measured geometry before planning [335]; metric calibration connects imagined motion to feasible execution. GAVEL instead rolls plans through an explicit graph of action preconditions, effects, and location beliefs, repairing model-detectable violations before invoking semantic replanning [336]. Their shared predictive interface does not require a common visual representation.

Human-video world-action systems broaden the intermediate toward robot-domain visual evolution. Two dependencies become central: current robot progress must determine which part of the demonstration remains relevant [23], and the demonstrated procedure must change the predicted future even when the current scene admits a familiar continuation [24]. These dependencies connect the older goal-reaching formulation to one-shot task transfer. The predicted evolution must preserve the procedure selected by the example.

The two dependencies operate over different timescales. A demonstration supplies a procedure that should persist, whereas recent robot observations determine the next physically relevant transition. WorldScape Policy 2.0 combines persistent demonstration context, event-level memory, and short-term visual dynamics within a joint video-action model [25]. This moves the design problem beyond choosing RGB or latent futures: the model must keep the taught procedure stable while updating its account of the scene. Memory preserves the procedure while prediction tracks its current physical realization.

The planning and control clocks can be separated explicitly. MaP-WAM uses sparse observed segment histories to generate a language subgoal and visual plan; its executor reuses the cached plan until progress




Figure 11 Two roles of visual prediction: training (left) and deployment (right). Flames mark trainable modules; snowflakes mark fixed modules. (a) SynthICL trains the context layers and both heads, then removes the subgoal head; its frozen DINOv3 backbone is omitted [77]. (b) Zero-WAM trains the video and action branches, then freezes both and replaces recorded future video with generated video [24].

estimates trigger another planning step [401]. Matching the current observation to nearby plan frames calibrates that progress. At transition, actual execution frames replace the imagined plan in episodic memory. This separation matters: a forecast can guide action, but recording it as an achieved event would let prediction errors become historical evidence. The executor’s auxiliary future-video head can be omitted at inference, while the upstream visual plan remains part of deployed control.

Retrieval can constrain prediction before detailed motion is generated. The retrieval-augmented policy ReCAP supplies trajectories to a frozen residual predictor [152]; its Cosmos Policy branch jointly generates future images and actions. The retrieved motion narrows the continuation, while prediction relates it to the evolving scene. Other ReCAP backbones route retrieved evidence through different action interfaces, so the predictive classification applies to its Cosmos Policy branch.

The gain in inspectability comes with a new failure mode. A predicted scene can express the correct intention yet be unreachable, miss a necessary contact, or accumulate error during execution. Pixel prediction retains appearance detail at a computational cost; latent prediction offers compression but makes the retained task detail less visible. Action-conditioned simulation, goal-conditioned generation, and joint action–world policies consequently use related representations for different purposes [180]. The value of the intermediate depends on how reliably its task information survives conversion into executable motion.

Learning from futures without executing through them. Future supervision can improve an action representation without requiring a future to be generated online. Let $\mathcal{L}_a$ be an action-training loss and $\mathcal{L}_v$ an auxiliary future-prediction loss. The training and deployment interfaces can then be written as

$$\theta^* \in \arg \min_{\theta} [\mathcal{L}_a(\theta) + \lambda \mathcal{L}_v(\theta)], \quad A_t \sim \pi_{\theta^*}(\cdot \mid h_t, C_t). \quad (12)$$

Here $\lambda \geq 0$ weights the auxiliary loss. Optimization acts on the designated trainable components, with any frozen backbone held fixed. The retained components of θ^* initialize the deployed parameters θ_0 ; the auxiliary prediction head can be removed when action generation does not depend on its output. Equation (10), by comparison, retains a future-dependent action computation at deployment.

A geometric intermediate can improve training without becoming an online execution dependency. Gripper traces provide motion-specific supervision, while reasoning dropout or removal of the future-trace head

permits execution without generating those traces [181, 182]. The learned action representation then carries the benefit, avoiding a separate prediction stage at each decision. Internal semantic reasoning can play a similar role: PonderPounce conditions action inference on an asynchronous cognition token while keeping subgoal text inside the model [183]. The intermediate shapes the representation used for action inference.

The distinction is therefore between two computational roles of prediction (Figure 11). In auxiliary predictive supervision, future targets shape the policy during training, but actions do not depend on an explicit future prediction at deployment. In prediction-conditioned control, the predicted future remains an input to action generation. In panel (a), context layers denote the trainable context and state–context transformers after the frozen visual backbone. In panel (b), the video and action blocks denote the trainable prediction branches, not the input frames. Auxiliary prediction strengthens the representation used by a context-conditioned policy. Deployment-time predictions guide execution directly, coupling control quality to the generated future.

Predictive control separates task interpretation from its physical realization through an anticipated consequence. The anticipated consequence exposes an intermediate for inspection or planning. Its contribution to control depends on both forecast accuracy and the action interface that realizes it, making prediction and decoding distinct sources of error. Skill- and agent-based execution makes a different abstraction: it specifies operations and delegates their consequences to available skills or tools.

3.4 Skill- and agent-based execution

Skill- and agent-based execution assumes that existing skills or robot tools can realize a task once context has specified their selection, order, and arguments. A video showing a red block placed in the left bin before a blue block enters the right bin can become an ordered procedure. The executor supplies the available motor competence; contextual inference determines how to use it. Two branches make this separation concrete: learned skill composition and VLM and LLM agents that generate programs or tool requests. After defining their common specification–executor interface, we compare these branches and use feedback-driven revision to show when a changed specification suffices and when execution itself must be learned.

The specification can take the form of code, spatial targets, or an ordered sequence of learned skills. Its defining role is to govern an execution stage that can track or advance it, as described in Section 2. At that action interface, the deterministic form is

$$g = f_\theta(C_t, h_t), \quad A_t = \kappa(g, o_t), \quad (13)$$

where f_θ extracts the execution specification g from context C_t and robot history h_t , and κ realizes it as the next action block A_t using observation o_t . For stochastic specification generation and execution, the same interface becomes

$$\pi_\theta(A_t | h_t, C_t) = \int k_\theta(A_t | g, h_t) p_\theta^g(g | C_t, h_t) dg. \quad (14)$$

Here p_θ^g generates execution specifications and k_θ is the executor’s action distribution. Discrete programs or skills replace the integral with a sum; deterministic generation and execution recover Equation (13). In this family, g specifies operations, their arguments, or their composition for a separate executor. A transferred geometric reference belongs to the geometric family when reference tracking is the primary dependency; hybrids can combine both interfaces. A predicted future can also be decoded into g , composing this interface with world-model-based control. Neural components of the executor belong to θ ; numerical arguments and generated code belong to g . Revising these arguments or code can change behavior without training the executor’s neural weights.

Skill composition and program execution instantiate this interface as

$$\begin{aligned} \text{Skill composition: } & g = (z_1, \dots, z_{N_g}), \quad A_t = \kappa(z_{\ell_t}, o_t), \\ \text{Program synthesis: } & g = f_\theta(C_t, h_t), \quad A_t = \text{Run}(g, h_t). \end{aligned} \quad (15)$$

In the first row, ℓ indexes skills, z_ℓ is a skill specification, N_g is the number of skills, and $\ell_t \in \{1, \dots, N_g\}$ identifies the active skill. In the second, f_θ generates program g , and Run executes it until it produces the next action block, including permitted perception and control calls. Tool-using agents can generate one request at a time: the tool name and arguments form g , the robot adapter implements κ , and its returned observation updates the next decision. A generated program may call a geometric tracker, thereby composing this family with geometric demonstration transfer.

Composing reusable skills. An early insight was that rearranging existing motor capabilities can express a new task. Neural Task Programming [26] made this explicit in 2017: a demonstration instantiated a hierarchy whose leaves invoked robot routines. The transferable knowledge was the decomposition and its arguments, while execution reused the available primitives. This separated generalization in task length, ordering, and object assignments from learning the physical skills themselves.

Human-video teaching requires recognizing behavior across bodies, scenes, and timing. XSkill learns shared human-robot skill prototypes [29]; UniSkill scales skill representation learning to unaligned cross-embodiment video [30]. Both use human examples to guide a robot-side skill-conditioned policy, extending the interface from reordering known routines to recognizing transferable behavior in another body’s demonstration.

Skill abstraction must preserve the taught distinction: attributed waypoints retain motion targets [164], object-relation graphs retain dependencies [165], and interaction-aware procedures retain stage changes [93]. The executor’s interface determines which of these constraints reaches motion. Semantic subtasks can alternatively condition a continuous action predictor internally, as in StellaVLA [163]. An abstraction’s computational role therefore depends on whether it supplies a separate executor or the policy’s own action inference.

A composed procedure also needs evidence for when to advance. StageGuard distills demonstration-grounded completion reasoning into a lightweight monitor that decides whether to continue, advance, or skip a skill [337]. Navi-Agent verifies local navigation subgoals against a persistent topology of visual places and executed transitions [338]. Observed progress governs advancement through the composed procedure.

Navigation illustrates two levels of example-conditioned execution. Instruction-landmark examples in LM-Nav [359] and instruction-subtask examples in A²Nav [360] guide a language parser that specifies what pretrained navigators should execute. Select2Plan [352] retrieves visual decision examples for selecting marked movements or waypoints. The reasoning output is an executable choice, while motion realization belongs to the available navigator. Terrain feedback can further change the costs attached to those choices, as in VLM-GroNav [353]. Section 6.2 compares the information supplied by these contexts.

VLM and LLM agents: programs, tools, and feedback. Language-model pretraining expanded adaptation to program generation in 2022. Code as Policies [27] uses command-program examples to recombine perception, control, arithmetic, and feedback logic; program examples also specify available objects and interfaces [28]. Multimodal reasoning extends this route to demonstrated procedures. Show-Harness converts video frames into a textual outline for an unchanged planner and an embodiment-specific action interpreter [31]. The outline carries procedural structure into execution; spatial and contact constraints must also appear in this representation to guide the action interpreter.

Preserving contact structure requires a richer executable abstraction. RAPID derives local trajectory-optimization primitives, relational composition constraints, and a reconstructed verification environment from a human demonstration [390]. The inferred predicates and simulated scene variants support iterative program revision before the resulting primitives and strategy are frozen for transfer. Compared with an outline over supplied skills, teaching here helps construct both the available operations and their checks. Transfer depends on whether the inferred relations preserve the demonstrated interaction and whether reconstruction makes verification physically informative. Its nonprehensile experiments support this demonstration-conditioned pathway; the separate LIBERO-Pro experiment instead constructs programs without demonstrations. Reconstruction and program search contribute to the time and interaction needed to acquire the new task.

The command interface determines which decisions contextual inference must resolve. GPT-Policy uses goal images or demonstrations, available state and action records, and interaction feedback to request Cartesian

targets or gripper commands [94]. RoboDawn instead grounds discrete translation, rotation, and gripper commands through in-context examples [365]. In both cases, fixed models request actions from existing motor competence and use returned feedback to revise subsequent requests.

Tools expand the evidence and behavior available to this inference. Robo-Harness K1 exposes calibrated depth, spatial measurements, persistent visual anchors, and grasp hypotheses to ground generic motion requests [379]. The agentic framework RACaP exposes typed Policy APIs: a ReAct agent combines working memory, retrieved experience, and visual feedback to choose calls, arguments, and recovery steps [380]. Both retain fixed neural weights during deployment; RACaP also freezes its generated source code.

The memory-executor interface sets another limit. In 2AM, the agent retains task history but sends only subtask language and optional 2D grasp, place, and move hints to an episodically stateless action model [403]. These cues bind remembered intent to current-image instances without asking the agent to produce robot trajectories. Their validity differs: a grasp cue expires after acquisition, whereas a destination can persist during transport. The action model is trained to interpret noisy or incomplete hints. External task memory works through a learned steering interface. Its expressiveness determines which temporal and contact-dependent requirements can reach the action model.

Previewable actions add a further use of context: a candidate can be revised before it changes the scene. World Action Agent (WAA) selects interaction-centered views, renders proposed poses with kinematic planning feedback, and converts in-view corrections into robot motion [388]. Its main agent consults a separate Skill Agent for demonstration-derived procedures and visual checks. Tool requests form the execution specification, while geometric planning checks their feasibility before control.

After execution, the interface determines which feedback reaches the next decision. ETA, implemented in OpenETA [349], requires a fresh observation after each world-changing tool call. RoboHarness addresses transitions between pretrained policies with different operating conditions: execution memory selects a policy and retrieves states that guide entry into its operating region [350]. Feedback thus supports both revising an action and selecting the competence needed to execute it.

Generated programs can also construct spatial affordance and constraint maps for geometric motion planning [184]. The interface allocates competence: a named pick-and-place routine supplies perception and grasp synthesis, whereas pose-level control exposes those decisions to the agent. Improved success may consequently reflect a stronger supplied executor as well as better contextual inference.

A plausible procedure can still request an action that the current scene cannot support. Skill-affordance estimates address this gap before execution by coupling linguistic relevance to physical feasibility [185]. Scene and execution feedback address it after a transition by revising what should happen next [186]. When feasibility depends on an unknown property, targeted measurements such as proprioceptive estimates of mass or stiffness can change the manipulation strategy [187]. These are complementary uses of evidence: predict applicability, observe the outcome, and acquire information missing from either. They connect program-level reasoning to the actual conditions under which a skill works.

Feedback-driven revision changes either the procedure or the learned capability used to execute it. A frozen language model can adjust a primitive’s direction, distance, or duration from outcome errors [9]; these values remain arguments of g . CaP-X contrasts feedback-driven code generation with reinforcement learning of the coding policy [166]. ENPIRE exposes both choices within physical experimentation: revise perception and control programs, or collect experience to train a neural policy [89]. VoxPoser’s online dynamics-learning extension similarly changes a learned component [184]. The practical distinction is what the correction must supply. Recombining available operations changes their use; acquiring a missing predictive or motor capability requires a learning process for that component.

Across these forms, explicit structure exposes task order, constraints, and intervention points. Its central trade-off is *portability versus physical detail*: a compact procedure can cross scenes and bodies while omitting a required grasp or compliance condition. The executor can recompute only details left open by the teaching. DualManip separates the corresponding update rates: live geometric correspondence revises grasp contacts




Figure 12 Four technical paradigms convert context into actions through (a) direct action inference, (b) geometric motion-reference transfer, (c) prediction-conditioned control, and (d) skill or tool execution. Representative implementations are Vid2Robot [14], SemAnCorr [38], Zero-WAM [24], and UniSkill [30], respectively.

while semantic reasoning is invoked when geometric validation fails [410]. This hybrid preserves task intent across motion without requiring full replanning for every observation.

Table 5 separates context operations from their carriers. *Encode* commonly supports learned future prediction or skill inference; *Parse* exposes procedures and calls; *Retrieve* and *Verify* can augment either route. A video can become a future, a motion reference, a skill, or a program, each preserving different task constraints for its executor. The execution column exposes a second difference: a predicted future can guide planning or condition action decoding, while a motion reference can support tracking or replay. The carrier and execution operation jointly specify the transfer assumption.

Choosing an interface by the constraint it preserves. A useful interface preserves the task distinction most likely to be lost before execution. Direct action inference can exploit learned correspondence when the teaching resembles its training relationships. Geometric transfer exposes a binding contact or relative motion for inspection and tracking. Predictive control makes consequences available for selection, while skill and program interfaces expose ordering and preconditions. These choices relocate the generalization burden: to the learned interpreter, the matched interaction, the future model, or the supplied executor. Hybrids compose these assumptions. A predicted future can yield waypoints, and a program can invoke a geometric tracker. Their placement follows the intermediate that carries the demonstrated requirement into execution, allowing a hybrid to combine several families.

3.5 Shared mechanisms: correspondence and memory

All four interfaces face three linked requirements. A taught task must first be bound to the objects in the current scene, including replacements for those in the demonstration. Its evidence must remain available at the decision that needs it, and its interpretation must express a change the controller can execute. The discussion therefore moves from transfer across objects to access across time and then to the command interface that realizes a correction.

Figure 12 shows the intermediates passed to control: attended demonstration features, transferred contact frames, a predicted robot future, and a skill sequence. Evidence selection and interpretation (Figure 10) precede these computations. Correspondence binds the evidence to the current scene; memory preserves it until the relevant decision.

Table 5 Predictive and skill/agent interfaces (Sections 3.3–3.4). Carriers identify intermediate representations; families follow the full control path. Listed deployment branches keep neural weights fixed. Codes: D examples, H history, I instruction/goal, F feedback, E scene/map; V visual, S state/action, G geometry, L symbolic. Slashes combine entries.

Method	Source	Form	Context operation	Carrier	Execute
<i>I. World-model-based control</i>					
OSVI-WM [21]	D	V	Encode	Future	Plan
Demo-JEPA [22]	D	V	Encode	Future	Plan
HOST [23]	D	V	Encode	Future	Decode
Zero-WAM [24]	D	V	Encode	Future	Decode
WorldScape Policy 2.0 [25]	D,H	V,L	Encode	Future / Actions	Decode
ReCAP (WAM) [152]	D	V,S	Retrieve / Encode	Future / Actions	Decode
MaP-WAM [401]	H	V	Align / Encode	Future	Decode
FARE [402]	H	V,S	Verify	Future / Actions	Decode
<i>II. Skill- and agent-based execution: learned skills</i>					
XSkill [29]	D	V	Encode	Skill	Decode
AWDA [21, 164]	D	V	Encode	Motion	Track
ORION [165]	D	V,G	Parse / Align	Program	Invoke
UniSkill [30]	D	V	Encode	Skill	Decode
<i>III. Skill- and agent-based execution: programs and tool calls</i>					
Code as Policies [27]	D,I	L	Parse	Program	Invoke
Show-Harness [31]	D	V	Parse	Program	Invoke
GPT-Policy [94]	D,H,I	V,S,L	Parse	Calls	Invoke
WAA [388]	D	V,G	Parse / Verify	Calls	Invoke
RAPID [390]	D	V,G	Parse / Verify	Program	Invoke
Robo-Harness K1 [379]	I,E	V,G,L	Parse / Align	Calls	Invoke
RACaP [380]	H,I,F	L	Retrieve / Parse	Calls	Invoke
2AM [403]	H,I	V,L	Parse	Calls	Decode
Mimir [394]	H,I	L	Parse / Verify	Calls	Invoke
SparseNav [396]	I,E	G,L	Parse / Align	Motion	Track
DROC [46]	H,F	L	Retrieve / Parse	Program	Invoke

Table 6 Object substitution illustrated by pouring. The demonstration uses a jug and a cup; all three settings preserve the goal of pouring into the receiver. Arrows indicate object replacement; the last column gives example motion adjustments.

What changes	Held vessel	Receiver	Motion adjustments
Held vessel	Jug → bottle	Same cup	Grasp; tilt angle
Receiver	Same jug	Cup → bowl	Pouring position; height
Both objects	Jug → bottle	Cup → bowl	Grasp; tilt; pouring position

Cross-object transfer of demonstrated tasks. A demonstration can specify a relation that remains meaningful after its objects are replaced. Consider a robot shown how to pour from a jug into a cup. It may then use a bottle with the same cup, keep the jug but pour into a bowl, or use both the bottle and the bowl. Table 6 contrasts these three settings using this illustrative task. The held vessel is the manipulated object; the cup or bowl is the receiving object, a role that can also be a support or contact surface in other tasks. Each role can change within a category or through a functionally compatible substitute. Novelty relative to the demonstration and novelty relative to training remain separate properties. We first examine which relations transfer between new instances, then consider functional substitution and the control needed to realize the transferred requirement.

In this example, the robot must still direct the liquid into the receiver. A bottle can require a different grasp and tilt angle from the demonstrated jug; a bowl can require a different pouring position and height from the cup. Replacing both couples these adjustments: the new vessel must be grasped and oriented so that its outlet directs the flow into the new opening. The transferred requirement is thus preserved while the motion changes. Each setting isolates a different set of coupled geometric adjustments.

From object identity to task-relevant relations. An unchanged object in a new pose often permits relocation of a reference. A new instance can change the grasp, contact surface, or clearance even when its role stays fixed. R-NDF transfers relations such as placing a bowl on a mug or a bottle in a container to previously unseen object pairs [346]. Its original setting uses five to ten demonstrations with pretrained category representations and a marked 3D keypoint. Descriptor alignment recomputes placement with those representations fixed; the paper also studies a learned energy-based refinement, which requires training an additional neural component. A part-based representation supports one-demonstration transfer across larger shape variation [347]. Its real-robot tests include new mug–rack and bowl–mug pairs, using learned part models and segmentation assistance. These results connect the relation shown in a demonstration to execution on new objects while exposing the prior geometry and human input that make the transfer possible.

From new instances to functional substitutes. Category changes make the preserved role more important than whole-object similarity. FUNCTO transfers tool use from an actionless human RGB-D video and task description, distinguishing spatial, instance, and category generalization [172]. Its functional keypoints locate how a replacement tool should be grasped and used; the experiments emphasize tool variation, with less emphasis on changes to the receiving object. SemAnCorr transfers a pouring motion from a kettle to a toy cup and a cutting motion from scissors to a nail clipper, using recorded contact and motion from kinesthetic teaching [38]. These different evidence interfaces support functional correspondence without requiring a common demonstration format. Cross-category matching can also be supplied: the part-decomposition study evaluates pre-pouring alignment for a watering-can substitution using provided part equivalencies [347]. This protocol isolates geometric alignment with the cross-object correspondence already specified.

Functional substitution must preserve the requirement actually taught. A general request to transfer contents can admit another vessel; a specified vessel or handle contact can remove that freedom. At a higher level of abstraction, preserving a purpose may require a different procedure rather than a transformed motion. The Imitator Game distinguishes category-preserving substitution from this affordance-adapted imitation [101]. The invariant is therefore a task-dependent choice: object relation, contact constraint, process, or intended effect. Section 7.1 tests whether the demonstration determines that choice.

Realizing a transferred requirement. The four control interfaces consume this requirement differently. A context-conditioned policy can bind it within action features; geometric transfer exposes a reference; predictive control can express a desired future; and a skill or program can bind its arguments to new objects. Each interface preserves the requirement in a different form for its executor. Geometry used to train a policy [189, 190] also differs from a reference supplied at deployment or a demonstration-fitted motion system [191]. Likewise, spatial affordance prediction and spatial action representations provide useful priors [192, 193]; a newly supplied demonstration must still determine which requirement the prior should realize.

Relational skill guidance provides an agent-based route to transfer. WAA’s frozen LIBERO-derived skills transfer to robosuite lifting, stacking, and restacking, where objects, scenes, and viewpoints change; grounding resolves poses from the current scene [388].

When interpretation produces a desired pose or feature configuration, a separate controller can track it. For a stationary target, a local visual-servo rule is

$$\mathbf{e}_t = \mathbf{s}_t - \mathbf{s}_t^*, \quad \dot{\mathbf{q}}_t = -k_{\text{vs}} \mathbf{J}_{\text{vis},t}^\dagger \mathbf{e}_t. \quad (16)$$

Here $\mathbf{s}_t$ and $\mathbf{s}_t^*$ are current and desired visual features, $\mathbf{e}_t$ their error, $\mathbf{q}_t$ the joint configuration, and $\mathbf{J}_{\text{vis},t}$ the feature Jacobian. Its pseudoinverse and positive gain k_{vs} map feature error to joint velocity [194]. Context changes the reference while feedback corrects its realization. New contact or dynamics can still invalidate that realization, even with correct correspondence. The final part of this subsection examines

this execution boundary. Transfer also requires the defining evidence to remain accessible when needed, motivating the memory mechanisms below.

Accessing and updating evidence across time. Alignment selects information within a demonstration, retrieval selects records from an archive, and memory preserves evidence across decisions. These operations can expose policy examples, motion references, predicted rollouts, or programs. The retrieved object’s use in control determines its method family; the access mechanism itself does not.

A useful memory must answer three questions: which fact matters, how it can be recovered, and when it stops being valid. Task requirements, physical-response estimates, and completed-step records need different update rules (Table 7). A moved part can invalidate its remembered pose while leaving the taught order intact. The read–update decomposition separates stored evidence from the policy input; subsequent comparisons relate selection, compression, and reset timing to the decision they support. Sections 5 and 8.4 develop validity during execution and reuse.

Carrying evidence between decisions. A recurrent state m_t and an external archive M_t provide complementary ways of retaining experience. Let $x_t = (o_t, a_t, o_{t+1})$ be an executed transition, augmented with feedback when available. A common read–update decomposition is

$$\begin{aligned} m_{t+1} &= u_\theta(m_t, x_t), & M_{t+1} &= w_\theta(M_t, x_t), \\ R_t &= \rho_\theta(h_t, M_t), & C_t &= b_\theta(C_t^{\text{task}}, m_t, R_t). \end{aligned} \tag{17}$$

The maps u_θ and w_θ update recurrent and external memory, ρ_θ retrieves records R_t , and b_θ constructs the policy input. The supplied task evidence C_t^{task} includes instructions, goals, examples, and newly received corrections in any supported modality; it reduces to D^s in a demonstration-only configuration. The retained state (m_t, M_t) and the selected input C_t are distinct. An implementation may omit either memory component, and its maps may combine analytic operations with fixed neural components. This decomposition generalizes the self-contained update in Equation (2) without restricting context to demonstrations.

Retaining the information that changes control. Early latent task inference compressed transitions to identify reward or dynamics [196, 197]. Selective retention matters because a demonstrator’s collection style can correlate with task identity in training logs without identifying the task reliably at deployment.

Physical adaptation depends on when evidence is acquired and how long it remains relevant. RMA uses recent task transitions; ICWM collects task-agnostic probes; LocoFormer carries interaction across trials [39–41]. RopeFormer extends cross-trial action–response conditioning to dynamic ropes, retaining policy context while resetting the physical trial [364]. Zeva and Zeva-Ego instead encode action effects at two timescales; the latter combines recent within-attempt traces with phase-matched memory across attempts on the same task instance [88, 363]. These fixed-weight interfaces trade the cost of collecting evidence against its relevance to subsequent control.

Event memory addresses a different representation need: an earlier observation may establish that a button was pressed, even when the current scene no longer reveals that event. Perceptual retrieval and semantic event records preserve such evidence with different detail and processing costs [42–44]. Recurrent designs compress visual history at frame or action-chunk timescales [198, 199]; event-triggered storage retains selected raw keyframes [200]. MemBodied combines a fixed-size associative interaction state with an initial-scene anchor, preserving early evidence alongside later updates [381]. ARMS addresses concurrent execution by asynchronously supplying perception, embodied state, and records of which arm acted and when [387]. These designs retain different temporal structure: a stable reference, accumulated events, or actor-specific action history. Their deployed memory updates change the evidence read by a trained policy. HIRE retains ordered wrench traces to distinguish visually aliased interaction states, coupling this history-conditioned action inference to a faster contact executor [404]. Its shuffled-history comparison tests whether temporal order, rather than force availability alone, supplies the missing state evidence.

Explicit memory can expose the binding that a compressed state leaves implicit. Mimir separates world facts from task progress, retrieving candidates for the active goal and marking completion only when feedback supports its postcondition [394]. OCC4M makes a complementary distinction between persistent object

identity and a remembered place [395]. A selected object’s target follows its updated track; a historical-location target remains fixed and is reprojected into the current view. Together, these designs show why memory must preserve the intended referent and its update rule. Object motion updates a tracked referent, while task progress advances after the required postcondition is observed. OCC4M evaluates moving views in simulation with calibrated support geometry and uses a fixed camera in its hardware study.

Past evidence and imagined futures answer different temporal questions. An observation window establishes what happened, as in SimpleMemVLA [45]; combining retrieved events with imagined states, as in MemoryVLA++, adds what could happen next [209]. Acquiring a new rule for what to remember poses a further generalization problem. In μ VLA, recurrence helps held-out tasks with familiar memory requirements but yields little improvement when the required memory structure changes [199]. The distinction is between recalling evidence under a learned rule and inferring which evidence an unfamiliar task makes relevant.

Where selection is performed changes both cost and flexibility. Workspace Models move VLM-based saliency selection into training: an encoder learns latent workspace tokens by reconstructing selected historical image patches, and a downstream policy learns to act from those tokens [397]. This amortizes online reasoning but also inherits the saliency supervision’s information choices. A comparison with runtime retrieval can vary decision latency and memory demand together, testing both familiar event recall and the retention of new teaching.

History can also reconstruct an unobserved part of the present, rather than a past event or future consequence. LIFD combines recurrent evidence with geometry-anchored scene-token completion before action prediction [339]. This distinction matters for classification: generating missing current state supports context-conditioned inference, whereas world-model-based control requires predicted consequences on the execution path.

Selecting, compressing, and using history. The benefit of memory depends on which event survives compression and whether the current decision needs it. Selection and ordering already matter in language ICL [124]; robotics adds short physical events whose removal changes the meaning of an otherwise similar trajectory. Robust Instant Policy’s downsampling study identifies the loss of brief events such as a grasp transition [201]; its full system additionally fits a policy at deployment.

The useful horizon follows the unknown. AMAGO’s MetaWorld ML-1 analysis finds that short histories can identify hidden goals, while longer contexts can reduce sample efficiency [202]. ReLIC instead reuses experience across episodes to locate objects in an unfamiliar home, with deployed neural parameters fixed [203]. Section 6.2 relates this longer horizon to acquisition of spatial knowledge.

Retrieval faces two decisions: which event answers the current uncertainty, and whether that uncertainty requires history at all. Question-answering supervision can teach relevance, as in HALO [204]; a calibrated history gate can give current observations precedence when they already suffice [105]. The decisions expose different errors: missing a relevant event or allowing an irrelevant record to override the current scene.

For an archive $M_t = \{\xi_j\}_{j=1}^{N_t}$, ξ_j is its j th record and N_t is the number of records available at decision t . Let J be a candidate set of record indices and J_t the selected set. A budgeted retrieval rule is

$$J_t \in \arg \max_{J \subseteq \{1, \dots, N_t\}, |J| \leq K} \sum_{j \in J} [s_\theta(h_t, \xi_j) - \eta], \quad R_t = (\xi_j)_{j \in J_t}. \quad (18)$$

Here $|J|$ counts selected records, K is the maximum record count, and η is the relevance threshold. The compatibility score s_θ is used as in Equation (6), with an archive record as its second argument. The retrieved sequence R_t contains the records indexed by J_t . A fixed ordering rule arranges selected records in R_t ; a fixed tie rule resolves equal scores. The empty set permits a read with no useful history. When retrieval supplies the action-reuse policy in Equation (5), a nonempty selection yields its normalized mixture; an empty read instead requires the system’s fallback policy or an information-acquisition step. This additive form describes thresholded selection; sequence-aware retrieval can score combinations when the value of one event depends on another.

Table 7 Four roles of contextual evidence and when to recheck their applicability. One episode can supply several roles, but each can remain valid for a different duration. Representative mechanisms illustrate each role; Table B2 compares storage and access.

Information role	Evidence and mechanisms	Recheck after
Task specification	Plate-first order; corrections and relational programs [46, 390]	A new task or correction
Correspondence	Matched target or phase; object tracks and trace retrieval [392, 395]	Object or execution-phase changes
Physical response	Action effects; trial context and hardware-shift tokens [88, 393]	Tool, material, or contact changes
Execution state	Verified completion; task agendas and observed segment records [394, 401]	New action, observation, or scene reset

Motion-sensitive features can retain executed change and temporal order [205]. The next computation determines whether this evidence supports progress estimation or physical-response inference. The memory mechanism need not change when its information serves a different control purpose.

Stage-dependent retrieval and action refinement. Relevant experience must match execution progress as well as object appearance. RTCF aligns the growing history to successful trajectories and transfers only bounded low-frequency components from the selected action chunk [195]. The frozen policy supplies the remaining motion components and gripper decisions. Temporal correspondence thus selects evidence, and selective residual transfer limits how it changes control. Its fixed evaluation bank isolates the use of retained trajectories from online archive growth.

Failure traces add evidence about continuations to avoid. TraceFlow builds a bounded guidance field from progress-aligned successful and failed action traces, steering a frozen flow-matching expert without changing its weights [392]. This extends the role of retrieval from selecting an example to shaping action generation through outcome contrast. Reported ordering gains coexist with unresolved counting and occlusion deficits. These results expose two linked requirements: retained event evidence identifies the relevant stage, and action-space guidance favors a better continuation within it.

Persistence and update timing. Persistence is determined by the state that survives a boundary between decisions or attempts. Recurrent state and verified subgoal records retain information without changing the action generator [206]; reference traces and task memory can also organize repeated calls to fixed primitives [207]. Fast-weight methods encode evidence in an optimized parameter subset, either during the sequence [10] or before rollout [11]. Learning a value function from deployment experience changes a further component while leaving the actor fixed [208]. These choices differ in reset behavior and acquisition cost. Clearing observations, resetting adaptive weights, and restoring a critic to its initial state isolate these different adaptation components even when the actor remains unchanged.

Table 7 makes the information needs concrete: recalling a completed placement preserves progress, whereas retaining an insertion response supports physical adaptation. The same storage mechanism can serve either purpose; Table B2 in Appendix B (p. 97) compares representation and update timing. Together the tables distinguish the required information, its storage and access, and the conditions for its reuse.

Memory capacity, access cost, and adaptation speed are separate constraints. Caching keys and values avoids repeatedly encoding the prefix, but attention over an expanding cache still accesses more history. Segment recurrence bounds direct access to earlier states; fast-weight memory summarizes the stream in adaptive parameters; external retrieval keeps a larger archive but must select the relevant record. Compression trades access cost against recoverable detail, while fast-weight adaptation before rollout adds preparation time to a subsequently fixed per-action cost. The useful comparison is how long a decisive event remains recoverable and how quickly new evidence can change a command. Section 7 connects these questions to context-scaling experiments and feedback-to-action delay.

Table 8 External artifacts used with fixed neural weights. Plans can guide the current task; retained guidance and callable routines can support later tasks. Table B3 compares their update scopes.

Artifact	Execution content	Examples
Plans and programs	Procedure; controller	Code as Policies [27]; ProgPrompt [28]; Show-Harness [31]
Reusable guidance	Corrections; constraints; patterns	DROC [46]; ASPIRE [47]; ARCHITECT [48]; WAA [388]
Executable repertoire	Callable skills; composition	LRL [50]; RATs [52]; RoboRSI [348]
Execution oversight	Context; verification; intervention	SHAPER [49]; Zetta [90]

Exposing changes the controller can execute. Memory makes evidence available; the command interface determines whether the inferred change can affect execution. ReSteer exposes the state dependence of this interface: policies that complete two instructions from their usual starting states can fail to switch between them midway [108]. Transition data and policy refinement improve switching. The implication for contextual control is that a new instruction or correction must be actionable from the state reached under the previous one.

Command granularity determines what the context interpreter can express. Steerable Policies lets an off-the-shelf VLM use history to choose subtask descriptions, motion instructions, or pixel targets [210]. With the upper-level model unchanged, the richer interface improves progression over subtask-only steering. A spatial cue can revise a grasp where repeating an object name cannot. This fixed-model branch connects contextual inference to a responsive executor; Section 5 examines which correction the failure evidence justifies.

The path from evidence to correction has three dependencies: the relevant event must remain available, the command interface must express the required change, and the controller must execute it from the current state. This applies to all four method families. Agentic systems expose an additional level of adaptation by revising the records, programs, or context-construction procedures that support later decisions.

3.6 External knowledge reuse and revision

External knowledge reuse retains a lesson or executable artifact beyond the attempt that produced it. Three forms have different consequences: guidance changes later inference, callable routines enlarge the available repertoire, and revised context or verification code changes how the system uses evidence. Table B3 in Appendix B records these update objects. The analysis first separates acquisition from reuse, then compares the three forms before considering consolidation into neural parameters. This progression connects immediate contextual adaptation to reuse and then to improvements in subsequent learning.

To separate within-attempt inference from persistent adaptation, let B_n collect the external artifacts retained before attempt n : memory records, skill code, controller code and its numerical settings, or context-construction code. Let $\mathcal{T}_n$ be the executed trace and feedback from that attempt. A proposal-and-validation loop has the form

$$\begin{aligned} \tilde{B}_{n+1} &\sim F_\theta(\cdot \mid B_n, \mathcal{T}_n), \\ B_{n+1} &= \begin{cases} \tilde{B}_{n+1}, & \chi_n = 1, \\ B_n, & \chi_n = 0, \end{cases} \quad \theta = \theta_0. \end{aligned} \tag{19}$$

The fixed model F_θ generates a proposed artifact collection $\tilde{B}_{n+1}$; the tilde marks a candidate awaiting validation. The acceptance indicator $\chi_n \in \{0, 1\}$ is one for acceptance and zero for rejection under the system’s validation procedure, potentially including additional rollouts. Writing the effective policy as $\pi_\theta(A_t \mid h_t, C_t; B_n)$ makes this external dependence explicit. The next attempt uses B_{n+1} to construct context or execute commands.

In skill- and agent-based execution, accepted artifacts can select an executor κ_{B_n} or determine the program g and its numerical arguments. These generated code and settings belong to the external artifacts; learned

neural parameters of any executor remain part of θ . Systems without a separate acceptance gate use $\chi_n = 1$. The attempt index n differs from the decision index t : one artifact revision can involve several model calls and robot transitions. Section 7.2 accounts for the trials needed to acquire and validate the artifact before measuring reuse.

Table 8 groups external artifacts by their role in execution; Table B3 (p. 98) further separates plans from executable programs. A *procedural plan* specifies an explicit step sequence or outline. *Reusable guidance* stores acquired rules, applicability conditions, or annotated lessons. An *executable repertoire* contains separately addressable skills, routines, or motion primitives; a *program or its settings* specifies the task controller. *Execution oversight* uses code that constructs model inputs or implements a separate verifier, monitor, or intervention mechanism. Checks embedded in the task program belong to that program. These objects can coexist: a stored lesson guides subsequent synthesis, while a callable function directly extends the execution vocabulary.

Two uses of the retained artifacts explain how this update reaches control:

$$\begin{aligned} \text{Guidance reuse: } C_t &= b_\theta(C_t^{\text{task}}, m_t, R_t; B_n), & A_t &\sim \pi_\theta(\cdot \mid h_t, C_t), \\ \text{Executable reuse: } g &= f_\theta(C_t, h_t; B_n), & A_t &= \kappa_{B_n}(g, o_t). \end{aligned} \quad (20)$$

The context builder b_θ and memory inputs are those of Equation (17); the semicolon makes dependence on the retained artifacts B_n explicit. Guidance reuse changes the evidence supplied to otherwise unchanged control. Executable reuse makes the retained library or controller implementation available to specification generation and execution; κ_{B_n} denotes that artifact-dependent executor. Both may occur in one system. Together, Equations (19) and (20) connect an artifact’s acquisition to its later use under fixed neural parameters θ .

Grounding an interface and learning through it. Show-Harness makes the distinction concrete. Its VLM selects semantic motion commands, while a body-specific interpreter supplies calibrated execution. Video can specify an otherwise missing task order. Yet the action-representation experiment shows that reading explicit command conventions is substantially easier than discovering arbitrary command meanings through probing [31]. The former supplies a correspondence; the latter requires inferring it from action effects. The interpreter supplies the low-level control in both conditions.

High-level commands leave grasp synthesis and feedback control to the executor; lower-level commands require the model to resolve more geometry and timing. Guava compares this allocation through workflows and action abstractions [211], while compositions of frozen motor capabilities, planning, and verification expose its system-level effect [212]. This allocation determines how existing competence is used; context interventions identify the additional task information supplied by examples or interaction.

SkipVLA makes this allocation explicit within a hybrid policy: a classical planner handles free-space transit, while the instruction-conditioned VLA handles contact-rich motion [340]. Its learned target predictor connects the two. The routing design improves execution efficiency by assigning transit and contact to interfaces suited to their respective demands.

The first extension is from open-loop synthesis to outcome-conditioned revision. Textual feedback can revise a plan [186]; intermediate observations can revise the plan and Python controller separately [213]; program checks can make some responses executable without another model call [28]. Constraint violations narrow revisions to code and sampling bounds [214], while primitive-level errors narrow them to control settings [9]. These interfaces determine the available edit space. Section 5 addresses the separate question of which edit the observed failure justifies.

Retaining guidance and executable routines. Reuse asks which part of a successful repair remains valid later. Constraints attached to plan or code revisions preserve its scope [46]; causal annotations explain its effect [215]; distilled rules and code patterns support later synthesis [48]. Attribution determines the appropriate change: failure to follow valid guidance calls for execution repair, while a defective skill may require revising guidance or code [216]. Retaining the cause helps avoid inappropriate reuse.

Visual evidence makes retained guidance more explicit. WAA turns expert demonstrations and human corrections into procedures with applicability conditions, reference images, and outcome checks [388]. In its one-demonstration stove study, review replaces knob motion as the success criterion with an observed activation signal. Fixed agents then interpret the acquired guidance.

Executable reuse makes recurring behavior callable. Function synthesis supplies missing operations [27]; program clustering and human feedback refine reusable candidates [50, 51]. Video-grounded retrieval supplies implementation details such as contacts and target poses [217], while repeated successful code can enter a library [166]. Transfer then depends on scene-specific arguments and valid execution preconditions.

Autonomous trials can supply the evidence previously provided by a teacher’s correction, but the retained object determines what later tasks gain. Rewriting a controller improves the current solution [218]; task-agnostic play can accumulate code skills and failure records for later use [52]. Distilling a validated repair into a failure signature, applicability condition, and strategy instead retains guidance for constructing a future solution [47]. These routes share outcome-driven acquisition while transferring different knowledge: a specific implementation, a callable operation, or a rule for choosing a repair.

A structured library also supports failure attribution. RoboRSI organizes diagnosis and revision through a task-skill tree, retaining successful workflows as callable code [348]. Its code-on/code-off comparison reports 174/600 versus 129/600 successes over 120 tasks and five layouts each. This tests retained routines; its separate policy-training case consolidates corrective trajectories into neural parameters. The retained object distinguishes these two forms of reuse.

Separating semantic effects from geometric realization supports reuse across scenes. Teach and Grow’s Skill Blocks specify effects, scope, executors, verification, and recovery [53]. Its simulation pilot evaluates local library additions with deterministic decomposition and a fixed executor; agent-based induction and Experience Memory updates remain proposed. The evaluated contribution expands a repertoire, while the proposed extension would learn how to build it.

Reuse can encompass a complete closed-loop policy. Programs refined through demonstrations and simulation feedback become examples for new policy generation [341]. Generalized TAMP experiments similarly separate simulator-assisted code synthesis from evaluation of frozen programs on unseen instances [407]. RIVET makes the interface explicit: generated perception and planning programs share object poses and a relation graph, then operate without runtime VLM regeneration [408]. HuGo instead refines high-level policy code from rollout trajectories and video while preserving a frozen whole-body controller [409]. Together, these designs locate adaptation in program construction and revision; subsequent execution tests whether the acquired logic transfers beyond its synthesis trials.

Adapting how evidence reaches control. Retained guidance improves how available routines are invoked. Parameterized call sequences preserve useful orderings, while global heuristics describe the operating conditions and failures of fixed primitives [207]. Hi-VLA’s tabletop study likewise finds greater benefits from prior-episode affordance summaries than longer within-episode history [219]. Such memory captures conditions of use, connecting agent reasoning to system identification of available capabilities.

Repeated failures in evidence selection or verification motivate adapting the surrounding inference procedure. Revising history-selection code changes what the planner can see; revising a critic or recovery tool changes how execution is checked and corrected. SHAPER targets procedural guidance and context construction with the planner and executor fixed [49]; Zetta targets critics, recovery tools, and governance around a frozen policy [90]. RACaP jointly evolves Policy APIs, the ReAct harness, and experience memory before freezing the resulting system for deployment [380]. This separates two timescales: external-artifact revision improves the available interface, whereas runtime context selects how the current task uses it.

Compiling a repair also changes when reasoning is paid for. HarnessPAI retains a task program during each rollout, including observation-conditioned checks, and revises it between rollouts using stage-indexed failure-to-repair memory [398]. Its program-level open loop therefore still permits sensor feedback during execution. This amortizes synthesis across repeated tasks, while a changed instruction or operating condition can require renewed validation. Its skill-library comparison shows faster program convergence on three

representative development tasks. Held-out acquisition tests would measure how this benefit transfers to subsequent learning. In its LIBERO evaluations, the reported 50 seeds include 15 used for program evolution, a distinction relevant to interpreting transfer.

Connecting contextual acquisition to neural consolidation. Neural consolidation turns acquired experience into future model competence. Distilling harness interaction traces [211, 388] or combining orchestration with autonomous data collection and policy refinement [220] transfers work initially performed through inference into trained parameters. EmbodiedSWE diversifies verified agent solutions into demonstrations for VLA training [382], supplying variations beyond the solved task instances. The training target determines what is consolidated: successful trajectories teach execution, whereas successful teaching exchanges, as in LMPC, teach responsiveness to later feedback [117]. Section 8.5 develops this distinction. The next chapter examines how training data connect task evidence to execution.

Takeaways

- ① Context determines an action distribution, a geometric reference, predicted consequences, or an execution specification. Equations (4), (7), (9), and (13) identify these four interfaces.
- ② Context interpretation and motor competence limit transfer separately. Retrieval and memory support transfer by supplying usable evidence to a controller that responds to later corrections.
- ③ With neural weights fixed, context can revise task outlines, executable routines, and programs. Retaining these artifacts supports reuse, and successor comparisons measure their contribution to learning subsequent tasks.

4 Data and Training: How Is Context Use Learned?

An input channel makes context accessible; training determines whether the policy uses it. When the current scene already identifies every target action, a model can fit demonstrations without learning to consult earlier teaching. The central data question is therefore which relationships make earlier evidence informative about a later decision. Acquisition sources determine what is measured, training units establish these relationships, and construction at scale supplies their missing variations. This connects the transfer assumptions of Section 3 to the experience needed to support them.

Source taxonomies, such as the Data Pyramid survey [221], describe where embodied data originate and how closely they match robot execution. ICL for robots additionally asks what earlier teaching or interaction reveals about a later decision. A record's source, channels, and role in learning are therefore separate dimensions. Explicit pairs, continuous streams, and linked attempts organize these dimensions differently; pretrained interpretation and geometric construction can also connect evidence to execution without jointly training the entire path.

4.1 Input sources and context modalities

An acquisition source determines what is measured directly and what must be inferred. Teleoperation records robot commands, UMI estimates motion through a robot-like device, egocentric video records human activity, and simulation provides modeled transitions. The four sources are compared on this common basis.

Language can accompany every acquisition source. Instructions specify constraints, command-program examples establish conventions, corrections revise decisions, and summaries retain earlier events. Code as Policies, DROC, and MEM illustrate the latter three roles [27, 44, 46]. Records should preserve each statement's source, time, and referent: a teacher's requested change and a model's interpretation of an event carry different evidence.

Figure 13 contrasts the directly observed interaction: robot execution in teleoperation, instrumented hand motion in UMI, human activity in egocentric video, and modeled execution in simulation. These sources leave different correspondence problems for the learner, as the following parallel comparisons explain.

Teleoperation: measurements and commanded interventions. A real robot provides synchronized scene observations, measured robot state, and issued commands. DROID and BridgeData V2 supply such robot demonstrations [56, 58]; RH20T additionally records force-torque measurements and, in one hardware




Figure 13 Four sources of contextual evidence for robots: (a) teleoperation, (b) UMI, (c) egocentric video, and (d) simulation. Representative resources include DROID [58], UMI [62], Ego4D [68], and MimicGen [74], respectively. Their recordings differ in visual, state, action, and contact channels; language can accompany each as instructions, examples, or feedback.

Table 9 Acquisition interfaces compared by state, actions, and contact evidence. UMI pose and width are estimated; extra sensing and simulated contact records depend on the setup. Table B4 lists resources.

Source	State	Actions	Contact
Teleoperation	Images; joints; gripper	Recorded commands	Force; tactile (optional)
UMI	Images; IMU; gripper pose	Mapped gripper targets	Visual; force (optional)
Egocentric video	Video; optional sensing	Inferred or paired actions	Visual interaction
Simulation	Rendered views; state	Simulator controls	Modeled contact

configuration, fingertip tactile signals [59]. These records can connect an attempted intervention to the body’s response. A gripper-close command specifies an action, while measured aperture and contact help establish whether the object was grasped. Joint torque or motor current also requires interpretation through the robot’s mechanics before it can describe external contact.

UMI: tracked motion without the robot arm. UMI reduces collection dependence on a robot arm by preserving a gripper-like interface in the demonstrator’s hand. Images and inertial measurements support device-motion estimation, while visual tracking estimates gripper width [62]. The resulting trajectory can supply end-effector targets without measuring arm-joint motion. Sharing the end effector between collection and execution reduces a further source of correspondence error [66]; synchronized tracking improves the fidelity of the recovered motion [223]. Extensions address acquisition, whole-body realization, and richer interaction sensing [224–226]. Their common trade-off is deliberate: an instrumented interface restricts how a human demonstrates, but reduces the inference needed to turn that demonstration into robot-compatible evidence.

Table 10 Context units, retained evidence, and adaptation roles. Support actions may enter context; future query actions are training targets.

Context unit	Evidence	Role
Support example	Demonstrations; paired examples	Task; procedure; contact
Scene preview	Exploration views; spatial records	Layout; locations; routes
Query history	Observations; actions; dialogue	Progress; recent response
Transition sequence	Observation–action–outcome	Physical response
Cross-attempt record	Episodes; summaries; corrections	Reuse across resets

Table 11 Three training organizations compared by evidence, supervision, and record relationship. Support can omit robot actions; linked attempts can use predictive, corrective, or reward-based supervision.

Organization	Earlier evidence	Learning signal	Link between records
Paired episodes	Support demonstration	Query actions	Shared task or environment
Continuous stream	Observation/action history	Subsequent actions	Temporal continuation
Linked attempts	Trials; outcomes; corrections	Selected actions or feedback	Consequences across resets

Egocentric video: visual procedure with optional human sensing. Egocentric recordings expose object interactions and their order without requiring a robot-compatible device. Ego4D and EPIC-KITCHENS-100 provide egocentric activity, and Ego-Exo4D adds synchronized external views [68–70]. Some collections also provide audio, gaze, inertial measurements, depth, or tracked hands. They measure human motion and the surrounding environment. Robot action targets must come from paired execution, reconstruction and retargeting, or a controller that grounds the observed procedure. An action-free video supplies visual observations without recorded robot commands.

Simulation: rendered observations and modeled responses. Simulation makes action consequences and privileged state available under a specified dynamics model. Demonstration environments such as RLBench, LIBERO, and ManiSkill2 support task execution with synchronized observations and controls [227–229]; trajectory or task-program construction produces additional executions [74, 230]. This control over the generating process is useful for contextual learning because a procedure can be held fixed while scene geometry changes, or varied while the scene remains comparable. Privileged object and contact states can verify those relationships during construction, even when the deployed policy receives cameras and proprioception. Section 4.5 examines how synthesis preserves the assigned teaching condition.

Channels and policy inputs. Table 9 separates measured state, action information, and contact evidence. A channel used during collection need not be a direct policy input: UMI uses IMU measurements for pose estimation, while its policy observes images, end-effector pose, and gripper width [62]. For ICL, the available fields must be specified independently for the support demonstration, the current query, and the supervision target. A human support may contain images alone while the executing robot queries its joint state and contact sensors. Conversely, recorded support commands can enter context, whereas future query commands serve as training targets. In the notation of Section 2, available sensory measurements enter o_t , issued commands are a_t , and their time-ordered records form h_t ; the selected evidence enters C_t . Textual instructions, examples, and feedback can accompany these records and enter C_t through the model’s context interface.

4.2 Training units and context dependence

Channels become useful training evidence through temporal and task relationships. Table 11 compares paired episodes, continuous streams, and linked attempts. We define the temporal units first, then use support–query supervision to distinguish input channels, action targets, and pairing criteria. Continuous experience and corrective records extend this dependency beyond explicit pairs.




Figure 14 Demonstrations serve three roles: (a) supervise motor actions, (b) train context use through paired episodes, and (c) specify a new task to a fixed policy. Panel (c) illustrates a can-uprighting demonstration specifying stapler reorientation; it is a conceptual example of cross-object transfer.

Temporal units of contextual evidence. A contextual training unit must identify what was observed, what was commanded, and what the resulting feedback established. Language can specify the intended task or annotate the outcome. The sequence must preserve when each item became available, so that a later consequence can supervise learning without appearing as evidence for an earlier decision.

Temporal alignment preserves the relation between an intervention and its consequence through sensor timestamps, command application times, coordinate conventions, and episode boundaries. Measured motor state must remain distinguishable from its command or setpoint. Aligned windows or timestamp-aware encoding accommodate different sampling rates while retaining brief contact events. Missing channels and reconstructed motion must also remain identifiable, because they supply different evidence from direct measurements.

Table 10 groups records by what they contribute to learning. Support can specify behavior or reveal an environment; query history locates execution; transitions reveal action consequences. These units overlap: history contains transitions, and retrieved segments can come from demonstrations or earlier attempts. Their role is independent of modality. Images or synchronized robot records can supply a demonstration, while state-action history can support physical adaptation without one, as in RMA [39].

Supervising a query from a support episode. Following the paired-execution principle of one-shot imitation [1], let support demonstration D^s and query observation-action trajectory τ^q share task specification ψ . The matched executions may differ in objects, timing, or embodiment. A common supervised objective predicts query actions from their history and support:

$$\mathcal{L}(\theta) = \mathbb{E}_{(D^s, \tau^q) \sim P_{\text{pair}}} \left[\frac{1}{|I^q|} \sum_{t \in I^q} \ell_{\theta}(A_t^q; h_t^q, D^s) \right]. \quad (21)$$

Here P_{pair} samples executions matched by specification ψ . Query history h_t^q ends at the current observation, while $A_t^q = a_{t:t+H-1}^q$ is the recorded target block. The nonempty set I^q contains positions with complete

targets; averaging over it normalizes each episode’s contribution. The loss ℓ_θ may use likelihood, regression, denoising, or flow matching. These objectives differ in action modeling but share the same supervisory dependency: the support must explain which query behavior is appropriate.

Equation (21) specializes the general context interface to $C_t = D^s$. In instruction-bearing episodes, the same query-action loss conditions on C_t containing both the demonstration and its associated text. Instructions specify goals and constraints; demonstrations can convey remaining choices of order or contact. A correction enters context when received, preserving what was available before each supervised action.

Figure 14 distinguishes motor supervision, training to use teaching, and deployment-time task specification. Uprighting a differently shaped object illustrates the third role: the operation must survive changed appearance. Paired supervision learns this dependency explicitly; continuous streams can expose it through the sequence of experience.

Support channels and action supervision. Equation (21) supervises query actions while leaving support channels open. Vid2Robot conditions on video without recorded support commands [14]; ICRT supplies support actions in a sensorimotor prefix [13]. Both supervise the query execution. Support content, the action loss, and the architecture that combines evidence are therefore separate choices.

Task identity and support–query pairing. The shared specification ψ determines the granularity of teaching. Pairing all executions labeled “put the object away” can preserve a destination while discarding grasp or route preferences. BPP makes this dependence explicit in its task-grouped demonstrations [54]. The pairing rule must retain the distinction the deployment example is expected to convey. For cross-object transfer (Section 3.5), support and query can change the manipulated object, the receiver, or both while preserving the same requirement. Conversely, retaining comparable objects while varying the required relation makes the support informative. Together, these pairings separate task meaning from object identity and provide examples of which substitutions preserve the intended function.

Pairing requires both semantic agreement and compatible action meaning. Additional demonstrations can distinguish task interpretations left ambiguous by one example, as Instant Policy illustrates with differently shaped boxes [16]. The resulting query target must still express the same intervention in consistent coordinates, timing, and gripper semantics. WALL-OSS-0.5 addresses this physical side through synchronized video and actions, camera repair, and end-effector calibration [233]. These requirements are complementary: semantic pairing determines which behavior is intended, while calibration determines what command realizes it. With shared action coordinates, each body retains its own reachability and contact mechanics.

A shared environment supplies a different pairing relation from a shared procedure. Route-following training pairs a demonstrated path with a perturbed execution, as in RPF [356]. Scene-conditioned training instead pairs a preview with navigation to goals within that environment, as in NOLO [351]; ReLIC [203] links episodes in one home while changing the target and start. The shared variable is therefore the route in the first case and the environment in the latter two. Scene identity and memory-reset boundaries determine which spatial knowledge remains available across examples or episodes.

From paired episodes to continuous experience. Explicit pairing makes the teaching–execution relationship part of the training unit. S1 supplies an in-context demonstration in its pretraining episodes [99]. GEN-1.5 instead reports next-action pretraining on randomly sampled continuous physical spans, without dedicated demonstration packing or an ICL-specific objective, and inserts sensorimotor examples as prompts at deployment [98]. A stream exposes earlier observations and actions without assigning them separate support and query roles. Both systems later accept task examples, acquired through different organizations of training experience. Action generation is a further choice: ContextFlow uses explicit robot support–query supervision with continuous flow matching [55]. The action loss and the organization of experience play different roles: one fits the action model, while the other links preceding evidence to subsequent behavior.

Behavioral diversity and contextual identifiability. Access to earlier evidence becomes useful when it resolves uncertainty about the continuation. If the current scene already determines every training action, the loss can be minimized without reading a demonstration. Comparable scenes with different valid procedures instead give teaching predictive value; repetitions of one procedure across scenes teach which




Figure 15 Task diversity and prompting affect transfer differently. (a) BPP drawing error with 10,000 demonstrations, averaged over three seeds [54]. (b,c) S1 compares demonstration and language prompting on seen and unseen tasks; its cumulative per-step success measure includes human recovery interventions [99].

variations leave it unchanged. The practical consequence is to count distinguishable teaching problems as well as trajectories. The comparisons below examine this dependency, while Section 8.1 asks how far the resulting inference transfers.

Context can also support shortcuts. Correlations with previous actions can improve imitation fit while impairing deployed control [234]; a prefix may identify collection habits rather than task requirements. Varying demonstrators, procedures, and physical conditions independently helps separate these explanations. Holding out task relations, alongside trajectories, tests transfer beyond familiar procedures. Multiple valid procedures in comparable scenes support both contextual training and tests in which the supplied example changes the correct action (Section 7).

The evidence distinguishes behavioral diversity from visual variety. ICRT’s selected DROID-only training condition fails its test tasks, whereas its multi-task scenes require different actions from similar initial observations [13]. BPP supplies a controlled allocation comparison: with 10,000 drawing demonstrations, mean Chamfer error is 9.5 pixels for 500 tasks with 20 demonstrations each, 5.1 for 1,000 tasks with ten each, and 3.4 for 2,000 tasks with five each (Figure 15a). Its less diverse folding setting can instead favor language conditioning [54]. Demo-JEPA separately varies task classes and trajectories per task [22]. Across these studies, useful diversity adds decisions that an example must resolve, supported by enough executions to learn the relationship.

S1 separates scaling gains on seen and unseen tasks under matched data, compute, and architecture apart from the prompt embedding [99]. Across 1k–100k training hours, both prompting methods improve on seen tasks. On unseen tasks, demonstration prompting rises from 1% to 66%, while language prompting rises from 0% to 9% (Figure 15b,c). Thus similar seen-task performance can conceal different returns to new-task teaching. The cumulative per-step metric includes human recovery, mainly for the language baseline. The S1 comparison concerns combined data scaling and prompting, while BPP directly isolates task diversity at a fixed demonstration budget. Matching the procedural information conveyed by text and video would further isolate modality from teaching detail (Section 7.1).

Linked attempts and corrective supervision. Learning from a failed action or a diagnostic probe extends context dependence across attempts. The training unit must link what was tried to what happened next, preserving reset boundaries and any change in the task or physical condition. Algorithm Distillation models such learning histories [6]; LocoFormer retains experience over varied morphologies and dynamics [40]. Reward-based meta-training can value an informative action through its benefit to later attempts [2]. A short probe gains learning value when its outcome informs the next decision. Linked attempts expose this dependency during training.

In recovery-conditioned supervision, the failed attempt and its feedback enter context while a later continuation supplies corrective targets. RoboTTT’s DAGger Distillation uses the full rollout as context but applies action loss only to human corrections, also improving its GDN comparator [10]. The failed actions

record what was tried; the selected corrections specify what to imitate. This training organization can support different memory mechanisms.

Equation (21) extends to this case by replacing support-only conditioning with C_t . For block-level losses, I^q includes only complete blocks composed entirely of actions selected for imitation. An objective supporting action-wise masking can instead suppress unselected positions within a block. Other executed actions remain in history. The context mask controls available evidence, whereas the target mask controls desired behavior. Runtime context contains only evidence available before the decision; future outcomes may supply supervision. Attempt boundaries must separately record physical resets, memory resets, and changes of tool or dynamics.

Long-context control requires representations that preserve decisive events. Past-Token Prediction adds reconstruction of earlier actions to future-action prediction, first learning short-context visual features and then caching them for longer-context training [235]. Caching reduces repeated computation; the cached visual features determine which task details remain available to the controller.

The supervisory unit determines what evidence can explain a later action: pairs link teaching to realization, streams link earlier events to current control, and linked attempts connect outcomes to revised decisions. Their common requirement is a decision that earlier evidence helps disambiguate. The remaining analysis examines acquisition sources that preserve this dependency with usable action correspondence.

4.3 Data coverage and action correspondence

Data coverage has two distinct meanings: the range of behavior recorded and the range of teaching–execution relationships recoverable from it. Robot records, motion-bearing human demonstrations, and ordinary video trade direct action correspondence against collection breadth. The comparison below follows that trade-off, asking which annotations make each source usable for contextual training. Simulation supplies controlled counterparts in Section 4.5; Table B4 lists resource-level details.

From reusable experience to annotated procedures. Motor coverage supplies reusable behavior before teaching can specialize it. Autonomous multi-task experience and shared demonstration collections addressed this need [236, 237]; Bridge Data included 7,200 demonstrations across 71 tasks and ten environments. Later collections broadened scenes and embodiments [56–58]. Contextual training additionally needs relationships between these executions that preserve what a new example should teach.

Procedure annotations identify why a continuation changes. AgiBot World’s March 2025 release spans 1,001,552 trajectories, 2,976.4 hours, 217 tasks, and 106 scenes, including subtask and recovery records [60]. Task intervals, step instructions, and later corrections connect demonstrations, autonomous attempts, and human interventions [238, 239]. Linking an attempt to its correction preserves the evidence that makes a different continuation appropriate.

Broader embodiment coverage makes shared task meaning and shared action meaning increasingly separate problems. Galaxea’s common platform and subtask timing provide a comparatively consistent collection interface [61]; RoboMIND 2.0 and RoboCOIN broaden the range of platforms [240–242]. Procedural metadata identifies which executions teach the same requirement, while kinematic metadata describes how their realization differs. A cross-body pair needs both. This explains why adding embodiments can expand motor coverage while increasing the work needed to construct reliable contextual pairs.

Robot records supply executable supervision. Robot records become contextual supervision when a query execution is associated with an example of its intended procedure [56–58, 60]. Cross-scene retrieval can locate candidate executions [243], but matching must distinguish alternatives that begin in similar scenes.

Human–robot pairing exposes an asymmetry between teaching and supervision. A human video can specify the intended interaction, while robot measurements reveal the commands and contact response needed to realize it. RH20T provides corresponding human videos together with robot images, joints, gripper state, force–torque signals, audio, and tactile sensing in one configuration [59]. These richer query records can supervise a policy whose eventual human support contains fewer channels. Task-level human–robot organization helps define the transfer setting [101], but a shared label such as *fold* can still merge different

procedures. Matching ordered events and terminal outcomes, as in WorldScape’s robot and UMI pairs, preserves a finer teaching relation [25]. Pairing granularity therefore determines whether the model learns an activity category, a procedure, or an interaction detail.

Mixed capture can reduce pairing cost. XR-2 combines 531.7 hours of robot teleoperation with roughly 1,000 hours of dual-UMI demonstrations [244]. Robot views, state, and actions supply executable targets; egocentric views and gripper motion supply candidate support. Shared laundry tasks narrow the search, while procedure-level matching identifies compatible executions.

Motion-bearing demonstrations reduce correspondence uncertainty. Instrumented teaching trades restrictions on collection for less uncertainty in transfer. Handheld grippers make the demonstrator’s observations and motion closer to a robot-compatible interface [62]. Preserving the end effector across collection and execution further reduces what must be inferred: YUBI reports 8,434 hours over 119 tasks and transfers the same gripper between arms through per-robot inverse kinematics [66]. Sharing the gripper reduces hand-to-robot correspondence uncertainty, leaving each arm’s reachability and control limits to be handled during execution.

Bimanual teaching makes relative motion fidelity more consequential than the visual plausibility of either hand alone. Stereo-inertial tracking, synchronized views, and inter-gripper reconstruction preserve this coordination [223]. HiFi-UMI’s reported 2K release also includes subtask boundaries, and its policy studies use UMI post-training without target-task teleoperation. Calibration preserves the physical relation between hands; boundaries identify which procedure that relation instantiates. Modular tracking can simplify acquisition [224], while a simulation-trained whole-body controller can supply realization beyond the demonstrated gripper motion [225]. More accessible capture and more accurate coordination are consequently distinct improvements to the teaching interface.

Contact sensing extends what the teaching record can express. OmniUMI synchronizes vision and motion with tactile signals, internal grasping force, and external interaction wrench [226]. Such measurements can distinguish a visually similar motion performed with different contact conditions. Their transfer still depends on compatible sensing, calibration, and execution interfaces; a measured wrench and a commanded gripper setting carry different information.

Collection can be conditioned on the current policy’s weaknesses. HIL-UMI compares handheld human trajectories with policy predictions on the same observation stream and targets additional data to disagreement and low-advantage segments [342]. The resulting demonstrations support further policy training. Thus policy-aware collection improves which experience is acquired, while remaining distinct from using that experience as fixed-weight deployment context.

Removing the collection device shifts this burden from measurement to inference. Tracked hands and segmented point clouds can be converted into gripper poses [16]; motion inferred from human play can supply pseudo-actions, retrieved context–query relations, and retargeted humanoid supervision [245]. These approaches broaden the usable recording pool, but the action evidence inherits reconstruction and retargeting error. The trade-off is most consequential near contact, where a small positional error can change whether a grasp or insertion is feasible. A common action format enables joint use of these records while their reconstruction reliability determines their supervisory value.

Mixed-source training can preserve this distinction instead of forcing all trajectories to carry the same supervisory weight. Aligning camera-space actions, embodiment, and chunk duration provides a common temporal interface, while reliability-weighted human pseudo-actions supplement a robot-supervised primary objective [246]. In ACE-Ego-0, the auxiliary objective emphasizes more dependable position channels. This weighting lets human records expand coverage while measured robot actions anchor execution. Carrying reconstruction confidence into contextual pairs likewise avoids treating uncertain hand tracking as an equally reliable robot target.

Zeva-Ego separates this acquisition of a physical prior from deployment adaptation [363]. Its Action-Centric Encoder converts egocentric visual transitions into continuous action-token supervision for mid-training. A later interaction-conditioned stage trains the action model to use observed action effects; deployment freezes



Figure 16 Temporal anchors for contextual teaching: opening the door enables loading, and closing it completes the procedure. These transitions align a teaching example with progress in a later execution.

the weights and updates memory. Human experience expands prior competence, while robot interaction supplies evidence about the current physical conditions.

Motion annotation reduces the uncertainty left by ordinary video while retaining a human-centered collection process. Hand-object geometry [63, 64] and contemporaneous hand and finger tracking, as in EgoDex’s 829 hours [65], expose motion that otherwise must be reconstructed from RGB. Robot learning still requires conversion into feasible contacts and controls. EgoScale’s 20,854 hours of action-labeled human data are connected to robot behavior through aligned mid-training and downstream post-training [67]; its one-demonstration experiment uses policy updates with one robot demonstration per task and additional human trajectories. These stages show why improved action evidence and fixed-weight task acquisition are separate contributions: the former can support either contextual inference or further policy fitting.

The transition from human pretraining to robot deployment also changes the meaning of language supervision. Xiaomi-Robotics-1 combines over 100K hours of UMI and egocentric records with cross-embodiment post-training [247]. State-transition descriptions during pretraining and imperative instructions during post-training serve different prediction targets. For demonstration-conditioned learning, the corresponding design choice is which task details remain unspecified by the instruction and must be recovered from the example. This connects source breadth to the pairing criterion in Section 4.2.

Ordinary video broadens procedural coverage. Ordinary video reverses the trade-off of instrumented capture: activity coverage is broad, while robot-compatible action information is sparse. Egocentric activity labels help locate candidate procedures [68, 69]; paired viewpoints and short human-object changes offer cues about correspondence and outcomes [70, 71]. These sources can identify what happened without fully specifying how a different body should reproduce it. For contextual teaching, selection must preserve the choice that distinguishes this example from another clip of the same activity.

Annotation is useful when it removes uncertainty about the relation teaching must preserve. Common capture protocols reduce incompatible recording conventions [248]; atomic segments and reconstructed hand and camera motion expose where an interaction begins and ends [249]. Object-state and contact annotations can resolve physical events hidden in RGB [250]. These additions are not interchangeable: reliable boundaries help preserve order, whereas contact-sensitive transfer needs relative motion and object response. The appropriate annotation follows the task distinction the support example must convey.

Procedural and conversational datasets contribute another missing dimension. Assembly101 records natural variations, mistakes, and corrections during assembly and disassembly; HoloAssist pairs a performer’s egocentric activity with an instructor’s spoken guidance and intervention annotations [72, 73]. Their contribution is the relationship between a partial procedure and information that changes its continuation. Transferring that relationship to robot ICL requires robot-compatible continuation targets or an executor that can realize the corrected instruction. Preserving mistakes, instructions, and corrections as linked events retains this distinctive supervision.

Human-video priors and demonstration following use this evidence differently. Learning latent plans from human play can reduce the robot experience needed for control [251], even when human and robot executions are not paired. Following a newly supplied video additionally requires preserving the demonstrated choice in a compatible robot execution. WAM-TTT learns that connection through paired human-robot meta-training, while its deployment-time fast-weight update needs only action-free human video [11]. This is parameter




Figure 17 Reported source duration across embodied collections and general-model pretraining [58, 60, 61, 65, 67, 247, 252]. Action-aligned recordings support grounded teaching episodes; broader audiovisual corpora supply knowledge across activities and domains. The contrast motivates comparing explicit context-use training with broader pretraining. Arrowheads denote lower bounds.

adaptation under the distinction in Section 2. Robot supervision establishes the connection during training; the later teaching input can omit recorded robot commands.

Preserving a teaching relation requires more structure than a shared task name. A mobile-manipulation episode contains identifiable substeps and transitions that can anchor retrieval or a corrected continuation (Figure 16). Human support and robot execution can traverse those stages at different speeds, so task correspondence matters more than frame equality. Constructing this usable teaching signal incurs a cost beyond recording the raw activity, motivating the acquisition comparison below.

4.4 Acquisition cost and contextual capability

Data value depends on the cost of producing usable teaching, not recording alone. We first examine correspondence costs, then compare explicit contextual training with broad priors and robot grounding.

Collection cost and correspondence. Filtering, synchronization, reconstruction, pairing, and validation contribute to acquisition cost. Egocentric video exposes hand-object interaction; UMI also measures device motion through a robot-like interface. The required conversion depends on the task: order-sensitive packing needs event boundaries and contrasting sequences, whereas tight insertion needs calibrated relative motion and contact evidence. Additional views or longer recordings improve coverage only if they preserve the distinctions needed for the task.

Figure 17 compares source duration and evidence type. Xiaomi-Robotics-1 combines UMI and egocentric recordings; general video spans broader activity with less robot-action correspondence. Hours measure exposure, while annotation effort and usable support-query pairs characterize the work needed to turn it into teaching.

GPT-Policy makes the contribution of correspondence concrete [94]. It compares robot video with the same video augmented by aligned end-effector poses, gripper states, and commands. Plug reinsertion succeeds in zero of three video-only trials and two of three augmented-reference trials. This is a combined state-action comparison. Goal-image and occluded-goal settings provide endpoint or self-interaction evidence (Section 6), showing how the content of the acquired evidence shapes its contribution.

Learning targets and acquisition routes. Pretraining can expand both executable behavior and the requirements inferred from teaching. Support-query episodes explicitly supervise the example-continuation relationship [1]; GPT-3 showed that broader pretraining can also support example-based specification at

inference [121]. Both routes can yield fixed-parameter adaptation. Their transfer to unfamiliar rules is an empirical question distinct from the choice of training recipe.

S1 and GEN-1.5 illustrate explicit contextual pairing and continuous physical experience (Section 4.2). Their shared demonstration interface motivates comparing which training relationships produced context use, with tasks, information, motor interfaces, and training budgets matched.

Mechanisms of fixed-parameter context use. Fixed-parameter computations offer candidate explanations for context-dependent behavior. Bayesian accounts infer a latent rule shared within a sequence [318]; distributional studies identify repeated structure that favors contextual inference [111, 319]. Controlled models can bind items to labels and retrieve their association [232], or implement gradient-descent-like updates through linear self-attention [112]. Comparisons with LLaMA show that similarities between contextual and parameter learning depend on the setting [322]. These accounts distinguish what training makes inferable from the computation used at deployment. For robots, tracing a taught distinction into the representation that guides action connects the training relationship to its physical consequence.

General-model priors and robot grounding. Astra adds a different source of prior competence to this comparison. Its public training summary describes broad text, image, and audiovisual sources, including video sources above one million hours, represented partly through transcripts, captions, or metadata [252]. In GPT-Policy, a fixed general VLM interprets task references and interaction feedback, proposes robot-tool requests, and receives outcomes from a constrained controller [94]. This places contextual interpretation in a general model and physical realization in an explicit robot interface. The disclosure characterizes source exposure; the causal contributions of individual training components to the observed robot behavior remain unresolved.

These acquisition routes supply complementary resources. Broad priors can interpret procedures and familiar interactions, while grounded records connect them to the receiving robot’s commands. GPT-Policy’s action-reference comparison illustrates this complementarity at contact-sensitive execution. The useful comparison therefore concerns the task distinctions and action correspondence obtained per unit of acquisition effort, with prior competence and the execution interface specified.

From exposure to usable teaching. The reported comparisons support different conclusions: BPP varies task diversity at fixed demonstration count, S1 varies scale and prompting, and general-model disclosures describe cross-domain exposure. Together they motivate measuring the marginal value of an additional recording: does it add a missing motion, distinguish another valid procedure, or reveal how to correct a failure? Section 8.1 connects these contributions to selection, composition, and rule inference. Retrieval and synthesis can then target the missing relationship rather than increase volume indiscriminately.

4.5 Scaling data construction

The coverage and cost analysis motivates three complementary construction routes. Retrieval reuses existing evidence; spatial supervision and task-preserving synthesis construct missing counterparts; failure-correction collection adds continuations absent from successful demonstrations. A demonstration can supply an execution reference or seed new training examples. The construction route determines which teaching relationships enter the learning units defined in Section 4.2.

Autonomous collection can increase experience while changing which situations the learner encounters. Task orchestration supplies new executions [76]; learned reward and success signals can select experience for subsequent policy training [254]. A contextual learner additionally needs the evidence that explains each selected continuation. Retaining its instruction, demonstration, or correction preserves a candidate teaching-execution pair. Contrasting pairs are needed to establish which behavior changes with that teaching. This links acquisition policy to the learning objective: selecting only successful motion favors motor competence, while preserving the preceding uncertainty and its resolution can train responsiveness to new evidence.

Retrieval for training and contextual pairing. Training retrieval asks which external behavior can improve a robot’s learned competence. Motion embeddings can locate human clips whose reconstructed trajectories provide reward guidance [255]; semantic, hand-pose, and motion similarity can select human

samples compatible with robot demonstrations [256]. In both cases, relevance is defined by the supervision entering the policy’s weights. This differs from contextual pairing, where the selected clip must explain a particular query procedure even if both behaviors are already within the motor repertoire. PHASE selects variable-length contact phases using tactile evidence, preserving insertion-stage compatibility in the retrieved training set [411]. This supplies a concrete counterpart to runtime retrieval: the selected evidence changes the policy learned offline.

Pair construction has a different criterion: the retrieved behavior must agree with the intended procedure. RHyME composes human clips for a recorded robot trajectory through sequence-level optimal transport [80]. Existing video replaces newly captured support, but constituent behaviors must occur in both collections, and clip boundaries can introduce pairing errors. During deployment, retrieval must additionally match current progress and motion compatibility. Retrieval-based ReCAP supplies state–action chunks to a frozen residual predictor [152]; expanding that archive changes runtime context through the mechanism in Section 3. Training selection, pair construction, and runtime retrieval therefore require different judgments of relevance.

Constructing spatial supervision. Spatial supervision reduces ambiguity that action loss alone must resolve. Synthetic cross-scene point correspondences associate the same task geometry across executions [161]; reconstructed contact heatmaps, hand configurations, and motion associate human interactions with affordances [176]. The first teaches where an analogous relation occurs, while the second teaches which interaction geometry is plausible. Contextual training connects either spatial prior to the supplied example, so that teaching selects which relation the robot preserves. Spatial accuracy and responsiveness to teaching then become complementary training objectives.

One demonstration as a reference or a data seed. A single demonstration can be reused as an executable reference or as a seed for training data. Retrieval and alignment support the first route in MT3 [36]. The second route generates additional trajectories by preserving bimanual coordination across object-dependent changes [257, 258], or by varying reconstructed scenes and embodiments [259]. These routes place generalization in different objects: a runtime reference or a policy fitted to its transformed executions. For contextual training, each generated trajectory remains associated with the demonstration whose task relation it realizes. This association turns motion augmentation into supervision for interpreting new teaching.

Task-preserving synthesis. Task-preserving synthesis needs both controlled variation and a criterion for what remains unchanged. Randomized human-arm renderings vary embodiment appearance [260]; geometric pseudo-demonstrations and waypoint renderings construct motion-correspondent pairs [16, 77]. Executing transformed segments in a simulator adds physical consequences and a success filter [74]. Environment generation and compilation of task semantics into costs and transition conditions broaden the available variations [261, 262]. These operations preserve different evidence. Rendering supplies correspondence, while executed validation tests realization under modeled dynamics; contextual supervision requires both to remain tied to the governing task specification.

KnowDemo makes the preserved requirement explicit [369]. It extracts task conditions from human video, grounds them in the target scene, and screens alternative contacts and subtask orders before planning and simulation. The verified trajectories then train a robot policy. This connects behavioral diversity to task constraints: varying an incidental grasp can expand training coverage, while varying a required order changes what the example teaches.

Simulation provides controlled task and scene variation alongside executable action records. RoboCasa expanded household environments and compositional tasks in 2024; RoboCasa365 adds 365 tasks and 2,500 kitchen scenes, with 612 hours of human-operated demonstrations in simulation and 1,615 hours generated with MimicGen [75, 263]. These are simulated robot trajectories, including those collected by human operators. For ICL, the opportunity is to hold a procedure fixed across environments or vary the procedure within a matched environment, then verify that each generated execution preserves its assigned teaching condition.

Cross-embodiment synthesis must distinguish a changed appearance from a newly validated interaction. Generating human-video counterparts for robot tasks starts from a robot execution [24]; retargeting human motion and rendering robot arms into its scene starts from a human execution [222]. In the latter, visible object responses remain those of the human recording. Physical-consistency and forward–inverse checks can

test whether a converted transition admits a compatible action [264, 265]. The teaching relation adds another requirement: the conversion must preserve order, goal, and binding contacts. A physically plausible transition can still be the wrong continuation for the supplied demonstration.

Agent collection can obtain additional executions through shared interfaces. GUMI records pre-action observations and semantic commands for humans and agents [31]. An ICL collection could use this interface to execute contrasting valid orders in comparable scenes, retaining each task specification, recording, and outcome. Pairing each successful execution with its governing specification would preserve the distinction needed for contextual supervision.

Processing must preserve intentional variation. An AgiBot-based study finds that demonstrator-specific speed variation can assign different chunks to similar observations, and that reducing this ambiguity improves downstream learning [266]. Uniform speed conversion is unsuitable when insertion requires alignment phases or pouring requires pauses. The corresponding ICL decision is selective: normalize incidental collection style while retaining a pause, route, or contact order when it is part of what the example teaches.

Failure-correction collection. Successful demonstrations leave recovery states underrepresented. Corrective collection changes the state distribution on which supervision is obtained: DAGger queries an expert on learner-visited states [267], DART uses perturbations to elicit nearby corrections [268], and selective intervention concentrates human effort [269]. XR-2’s separation of policy actions from human takeovers improves folding from 58% to 93% across three retraining rounds [244]. This expands the repairs represented in weights. Contextual recovery requires retaining the linked failure and feedback as well, so the learner can infer which repair the current evidence calls for.

The source of a correction determines what can be learned about its applicability. Synthetic perturbations, intervention in predicted rollouts, and planner-guided exploration expose different failure continuations [270–272]. Learning from intervention and autonomous outcomes can use these records without imitating every action [273]. For contextual learning, the record must retain the relation among the failed action, the corrective change, and its consequence. Action-level arguments and return values make that relation more specific than a task-level failure flag, as ASPIRE illustrates [47]. Richer failure records are useful when they help distinguish the causes a later correction must address.

Corrective records also need to expose mistaken task-state estimates. After a failed grasp, the gripper may be empty while a stored plan marks pickup complete. Linking the repair to observations that establish this discrepancy supplies supervision for revising context as well as motion. Retaining the repair as guidance changes the external artifacts B_n in Equation (19); using its trajectories as policy targets changes neural parameters. Section 7 tests whether either acquisition improves later tasks.

4.6 Connecting teaching data to a pretrained policy

Data construction must be matched by an interface through which teaching can change behavior. Three dependencies organize this integration: access to the evidence, preservation of its task distinction in the control intermediate, and an action encoding that can realize it. Access comes first. An observation-conditioned policy such as ACT needs an added context channel to distinguish procedures compatible with the same scene [274]. Language-conditioned policies expose goals and constraints [118, 119]; goal images specify endpoints [275]; sequential demonstrations can supply order or contact preferences [13, 55]. Adding an encoder provides access, while teaching-dependent targets establish use.

The missing dependency determines which training records a method needs. Foundation policies supply perception and motor interfaces [118, 275–279], but the four families consume different task-specific intermediates: actions, geometric references, futures, or executable specifications. Subtask annotations and skill representations expose higher-level structure [30, 60, 61, 241]; phase-aligned examples connect that structure to a current decision [18]. Direct supervision of an intermediate is one way to establish this connection. Joint training, pretrained interpretation, and geometric construction can also supply it. The requirement shared by the four interfaces is that the execution intermediate preserve the distinction selected by the evidence.

Human pretraining and deployment teaching act at different stages of this connection. Richer robot pretraining can increase the benefit of joint fine-tuning on human and robot demonstrations with reconstructed motion and subtask labels [280]. For ICL, the resulting policy must additionally respond to a new example with its parameters fixed. Conditioning on language, quality, speed, control modality, and visual subgoals, as in the $\pi_{0.7}$ recipe [281], illustrates how inputs can distinguish otherwise mixed executions. Demonstration conditioning can extend these inputs when a video conveys the distinction that the action generator needs.

Encoding is the final bottleneck between an informative training pair and an executable command. Regularized ICRT action tokens and hierarchical tokenization preserve different aspects of motion [282, 283]. LipVQ-VAE’s improvement from a low robot-success baseline shows that relieving a representation bottleneck can help while leaving execution difficult. The teaching signal passes through successive constraints: pairing identifies the intended relation, context encoding retains it, and action representation makes its realization expressible. These three stages must preserve the same taught distinction. The next chapter follows that requirement into deployment, where the current physical state determines how it can be realized and revised after failure.

Takeaways

- ① Data sources determine the available measurements; teaching–execution relationships determine what context can teach. Robot data supports motor competence; alternative taught behaviors within the same scene expose the additional dependence on demonstrations.
- ② Retrieval can select training data, construct a demonstration–execution pair, or supply examples during control. Only the last directly changes a frozen policy’s input at deployment; the first two change the experience from which it learns.
- ③ Useful scaling adds missing motor competence, task-defining alternatives, or evidence that changes a later attempt. The value of added experience follows the new distinctions it makes available to the learner.

5 Grounding: How Does Context Guide Recovery?

Objects can move, grasps can fail, and resistance can change after a task has been correctly interpreted. Recovery aims to preserve the taught requirement while revising the state estimate and continuation. It establishes the current state, diagnoses the deviation, and selects a feasible repair. Retaining the repair with its applicable conditions makes that feedback useful beyond the current attempt.

5.1 Grounding the task in the current physical state


Grounding requires two complementary estimates: the task requirement that remains binding and the physical state that may have changed. An insertion can preserve its destination and prerequisite order despite a moved part. We establish these estimates before examining motion feasibility and the entry conditions of transferred segments.

Preserving the task reference during execution. A transferable task reference must distinguish progress toward a goal from adherence to the taught procedure. Cross-view progress representations can align executions despite appearance or embodiment changes [78, 79]; demonstration-conditioned skill inference selects which procedure that progress belongs to [29, 80]. Current observations must then establish whether its prerequisites occurred and its required contacts remain valid. Terminal appearance can conceal an incorrect order, so verification must use the relations supplied by teaching [284]. This connects correspondence in Section 3.5 to recovery: a deviation is meaningful only relative to the procedure the robot was actually asked to follow.


Figure 18 distinguishes three failures in video-conditioned manipulation. An occluded interaction calls for a view that restores the missing geometry; a lost grasp calls for checking contact before resuming transport; a wrong target calls for revisiting the instruction and demonstration correspondence. Each repair addresses a different missing fact. Contact-intensive tasks add another ambiguity: a stalled insertion can reflect misalignment, unfamiliar resistance, or excessive force. A bounded probe can distinguish physical responses, whereas an irreversible operation may require inspection or clarification before another attempt.

(a) Visual occlusion




Human demonstration



Robot execution



(b) Grasp loss



(c) Target-selection error





Figure 18 Three failures call for different recovery evidence: (a) visual occlusion, (b) grasp loss, and (c) incorrect target selection. Human demonstrations on the left establish the intended behavior; robot executions on the right reveal where realization departs from it.

Maintaining a current, task-relevant scene. A repair must distinguish an intended change from its observed outcome. A placement command or predicted future specifies what should happen; completion evidence supports recording the step as finished. If the part subsequently moves, its remembered pose needs revision while the earlier placement instruction can remain binding. Memory and geometry-indexed retrieval recover unavailable observations [285, 286]; dynamic scene records replace spatial facts contradicted by new measurements [287–289]. When neither establishes applicability, another view or a reversible probe can distinguish the remaining explanations [290, 291]. A navigation failure can similarly invalidate a remembered route or its terrain assessment. VLM-GroNav [353] aligns visual terrain examples with sinkage or slip indicators before revising traversal estimates. This distinguishes physical feedback about a surface from geometric evidence that a passage is open.

Applicability can be checked before a remembered event influences control. CoreSense examines episode scope, provenance, temporal validity, and consistency [343]. Its recorded-data and simulated-signal evaluation measures evidence selection. Physical recovery evaluates the subsequent step: whether the selected evidence supports a repair that succeeds on the robot.

Connecting intended change to feasible motion. A corrected interpretation needs an executable continuation. Forward-inverse consistency checks agreement between proposed actions and modeled transitions; force-conditioned control adds measured contact feedback [265, 292]. These mechanisms make a context-selected correction executable. When the required motion lies outside the controller’s repertoire, recovery also requires a new skill or a change in the physical setup.

Checking whether a transferred segment still applies. Geometric one-shot methods distinguish initial alignment from recovery during execution. DOME and ODIL first establish a relative starting configuration, whereas staged alignment can re-establish it at later subtask boundaries [33, 37, 177]. A transport segment also depends on the object remaining grasped: after a slip, geometric alignment can persist while the segment’s entry condition fails. Its entry condition and intended completion therefore carry different information: the former establishes whether the recorded motion can begin, and the latter whether the next segment should follow. Learned subgoal transitions similarly govern progression within trained long-horizon policies [293].

Failure can invalidate a reference or expose missing control competence. Robustness via Retrying recomputes a continuation through image registration and predictive control [294]; ScrewMimic uses a demonstrated

screw-motion representation to guide further bimanual learning [295]. The former changes the continuation through existing components; the latter improves the component that realizes it. This distinction directs recovery toward revising the current reference or expanding the motor capability needed to realize it.

5.2 From failure evidence to corrective execution

Recovery proceeds from assessment to diagnosis and correction. Assessment establishes a deviation from the taught procedure; diagnostic evidence distinguishes its plausible causes; the controller’s correction interface determines which repair can be executed. The resulting continuation may be absent from the successful demonstration.

Task-conditioned progress assessment. Failure assessment requires a reference for what should have happened. Action consistency, visual progress, and calibrated rollout scores can detect departures from expected execution [81, 82, 296]. For ICL, the reference follows the current teaching, allowing an unfamiliar but valid order to define expected progress. GVL illustrates example-conditioned assessment by using observation–progress examples to improve a frozen VLM’s estimates from shuffled frames, including human examples for robot trajectories [83]. The taught reference makes progress assessment responsive to the current task.

Progress and physical acceptability can disagree. A connector can reach its apparent destination with an unsuitable contact force, just as a kit can reach the right inventory in the wrong order. Verification must therefore test both the context-defined requirement and the measurements needed to establish acceptable execution.

Diagnostic evidence acquisition. Diagnosis turns failure detection into a repair choice by identifying the likely cause. Multisensory diagnosis and learned failure knowledge organize clues to a stalled action [84, 85], but misalignment and an unseen obstruction can still require different interventions. Uncertainty about the intended task can be reduced by clarification and reasoning over alternative plans [86, 297, 298]. Uncertainty about resistance instead requires sensing or interaction, including tactile evidence absent from RGB [299]. The unresolved cause determines the information source: teaching clarifies intent, while sensed action effects reveal the resistance encountered by the robot.

Task-context exploration and question-based teaching formalize the choice to acquire evidence [300, 301]. The useful check is the one that discriminates causes relevant to the next action: a view for occlusion, contact feedback for grasp loss, or clarification for ambiguous intent. Affordance Agent Harness routes detection, segmentation, and other skills, verifies evidence, and retries before committing an actionable region [376]. Its episodic memory supplies priors for recurring objects. This illustrates how verification can govern evidence acquisition before a motion controller uses the result.

Correction granularity and feasibility. Correction granularity should match the cause of failure. A wrong target requires changing task selection; a displaced approach point requires a spatial update; insufficient displacement may require only a magnitude adjustment. Corrective-language training makes supervisory instructions actionable [87], while sparse adjustments around a frozen VLA can address magnitude errors when the target is already correct [142]. Subtask, motion, and pixel-target interfaces expose these changes at different resolutions [210]. The controller must also accept the correction from the state already reached, connecting diagnostic precision to the steerability analysis in Section 3.

Iterative in-context policy improvement links an observed error directly to a motion primitive’s settings. A frozen language model uses primitive parameters, execution errors, and correction examples to propose an adjustment; its comparison favors relative errors over raw states and targets [9]. The error is useful because it identifies a change the primitive can make. Neural parameters stay fixed, while the generated control settings change within the primitive’s available range.

When the frozen action policy cannot continue from the failure state, CoRe uses history and counterfactual continuations to choose a rejoin state, then realigns displaced objects before returning control [302]. The teaching still specifies the task; restoration makes an existing skill usable again.

Recovery can instead revise the history that conditions prediction. FARE compares a full causal cache, a recovered prefix, and a reset cache under the same latest observation, selecting their predicted continuations

Table 12 How teaching and feedback support assessment, correction, and reuse. Each row links supplied evidence to the decision or control element it changes. ENPIRE denotes its program-refinement branch.

Work	Available evidence	Affected output
Assessment and assistance		
GVL [83]	Observation–progress examples	Progress estimate
IntroPlan [298]	Retrieved reasoning examples	Plan; clarification
Corrective execution		
ICPI [9]	Execution errors; correction examples	Motion-primitive parameters
RACER [87]	Current images; corrective language	Action continuation
CorrectVLA [142]	Human feedback	Sparse action bias
Experience reuse		
DROC [46]	Retrieved language corrections	Plan; skill calls
Zeva [88]	Action–effect history	Subsequent actions
ENPIRE [89]	Execution feedback	Executable program

with a frozen world-action model [402]. This changes the model’s retained evidence rather than physically restoring the scene. A compatible historical prefix must therefore be re-grounded in the current state. Its ablations find that always invoking recovery search or committing an unverified prefix can underperform the base policy, supporting selective intervention with verification. A continuation must also preserve completed steps and restore failed prerequisites, as captured by task graphs and verified subgoal memory [206, 303].

Feedback is useful when it identifies a decision the robot can revise: its estimate of progress, its next action, or a rule retained for later use (Table 12). A contextual repair must change that decision while preserving the task requirements outside the correction’s scope. This connects immediate recovery to the longer-term validity of the lesson it produces.

5.3 Applicability and reuse of corrections

Reuse retains a conditional lesson: which intervention corrected which failure, and under what conditions. Guidance or executable artifacts preserve it through Equation (19); corrective training can incorporate such lessons into policy weights (Section 4).

Different lessons require different reset rules. Language-correction memory retains task constraints, whereas action–effect history retains physical response [46, 88]. A workcell reset clears completed-step memory; a changed connector or tool can invalidate a response estimate while leaving the assembly order intact. For a jammed connector, an adjustment justified by measured resistance should not be reused when misalignment is the cause. Retaining the observation, intervention, and outcome allows later attempts to reassess that condition instead of replaying the repair unconditionally.

Program-based reuse makes repairs executable. ENPIRE’s program-revision branch uses fixed neural models; Zetta retains critics and recovery routines around a frozen policy [89, 90]. Interruptible execution exposes intervention points [304]. Later trials test the repair’s benefit under applicable conditions. Task structure and irreversibility determine what can be repeated, linking recovery to the application demands below.

Takeaways

- ① Context specifies the recovery reference and helps interpret failure. Physical measurements establish quality; a responsive control interface makes the correction executable.
- ② Corrective training expands available repairs; contextual recovery selects or revises one from fresh evidence. Reuse requires retaining the correction with its outcome and applicable physical conditions.

6 Applications: What Must Context Resolve?

The value of context depends on the uncertainty an application makes consequential. Manipulation can leave the intended interaction unspecified; navigation can hide the route or local spatial structure; a long procedure




Figure 19 Cross-embodiment skill transfer in UniSkill [30]: (a) tissue pulling, (b) lid closing, and (c) cloth placement. Human demonstrations above and robot executions below preserve the interaction despite changes in embodiment and object appearance.

can hide a completed prerequisite; an unfamiliar body can hide its response to action. These demands determine what teaching, observation, or interaction must supply. Repeated task changeover combines them with acquisition and reuse costs. Any of the four method families can serve several demands, provided its intermediate preserves the information the application requires.

6.1 Task specification in manipulation

In rearrangement, an instruction or demonstration can specify which object belongs in which container, how components should be ordered, or which tool serves a purpose. A moved destination requires a new path that preserves the demonstrated object–destination relation. The Imitator Game makes differences between demonstration and execution explicit through scene, object, and functional changes [101]. These changes test preservation of intent across altered object correspondence. We first consider the distinctions an example must convey, then how they survive replacement of the objects involved, and finally how geometry or contact response changes their realization.

When an example conveys more than a task label. The value of a demonstration depends on what the user needs to specify. A final drawing leaves stroke order, direction, and tool approach unresolved; BPP’s DrawAnything study connects broader drawing-task diversity to better demonstration-conditioned generalization [54]. In its three-task laundry study, language is more reliable than behavior demonstrations, with some demonstration-conditioned failures executing the wrong fold. These contrasting settings identify an application condition: video is especially useful when it communicates a choice that a short task label or goal image leaves open. A goal image can specify the task when only a visible endpoint matters; a demonstration additionally conveys a required path or order.

Replacing objects makes the distinction between task meaning and execution concrete. A new workpiece can preserve an assembly relation while changing its grasp; a new fixture can preserve the goal while changing the insertion geometry. Replacing both requires correspondence at both ends of the interaction, as the object-pair transfers in Section 3.5 illustrate. The application must specify whether an object identity is binding or whether another object may serve the same role. Functional substitution is useful only when it preserves that requirement, including contact, order, and material constraints.




Figure 20 Task meaning and physical realization vary independently: (a) an unchanged baseline, (b) a new task relation, (c) a changed physical response, and (d) both changes together. Teaching specifies the intended relation; interaction reveals how the robot can realize it.

Personalized behavior adds persistent conventions to task specification. Language corrections can retain user preferences and guide subsequent plans [46, 305]. A request to place objects quietly should govern later placements, whereas a correction to one object’s location need not. Such applications require both an acceptable endpoint and adherence to the user’s continuing constraint.

Realizing these requirements across bodies adds a correspondence problem. In bimanual transfer, coordination must be preserved while each arm uses a feasible motion. Conditioning one predicted arm trajectory on the other addresses dependence within the action inference problem [92]. Sharing an end effector across arms reduces a different uncertainty, the physical mismatch between demonstration and execution [66]. Teaching specifies the intended cloth region or contact relation; action inference and embodiment design determine how coordinated motion realizes it.

Figure 19 shows the preserved interaction across human and robot executions: pulling tissue, closing a lid, or placing cloth. The relevant object part and destination remain identifiable while the robot changes its approach and contacts. A handle-specific instruction restricts that variation; an endpoint-only request leaves the grasp open.

Method selection follows the binding requirement. Geometric transfer exposes a contact or motion constraint; a skill interface exposes order and prerequisites; contextual policies and predictive control can encode either through their learned intermediates. The application supplies the criterion for judging these choices. Training pairs that equate executions solely by endpoint can erase a route or contact requirement that the user values (Section 4.2).

When intent and physical response change together. Deformable manipulation couples the procedure to a changing scene. A fold alters visible surfaces and accessible contacts, so the correct next action cannot be obtained by repeatedly matching the initial demonstration frame. Geometric and semantic information can jointly condition a folding policy [91]; preserving fold boundaries then provides a more useful correspondence than copying absolute hand trajectories. Instant-Fold’s simulated mode variation and real garment variation examine different parts of this transfer. Its RGB-D and interaction interface, including estimated and potentially corrected human interaction points, also supplies richer information than ordinary RGB demonstration video.

Material changes can leave the intended fold unchanged while altering the force needed to realize it. Visuo-tactile measurements distinguish completion from excessive compression [306]. Here the demonstration

supplies the procedure and action–effect evidence supplies the current material response, giving the two context sources different roles.

A new arrangement under familiar conditions primarily changes the requested relation. A familiar insertion under changed resistance primarily changes physical realization. A new fold on unfamiliar fabric couples both: the taught fold remains the reference, but each action alters the geometry and contact conditions for the next. This independent variation provides a basis for choosing teaching and feedback, even within one application.

Figure 20 crosses the two demands: a changed task relation, changed physical response, or both. This view identifies what evidence an application needs before comparing the method used to process it. Navigation extends this question from object interaction to spatial knowledge and traversability.

6.2 Navigation and spatial adaptation

Navigation requires more than a destination: a robot may need a demonstrated route, knowledge of an unfamiliar building, examples of how to interpret directions, or feedback about an unsuccessful passage. Table 13 organizes this evidence into four context types. The following paragraphs compare them in parallel by what they convey and how they change navigation. This application taxonomy concerns the supplied information; the method families in Section 3 concern the computation that turns it into behavior. Several context types can be used together.

Table 13 Navigation contexts grouped by the information they supply at deployment. Each type supports a distinct decision, and a system can combine several types. These information categories complement the four computational method families in Section 3.

Context type	Evidence and decision	Representative approaches
Route demonstrations	Ordered views or landmarks, sometimes with actions. Follow a demonstrated path or sequence of subgoals.	Zero-Shot Visual Imitation [355]; RPF [356]; sparse visual memory [357]
Environment observations	Scene previews or earlier visits. Locate objects and connecting passages for new goals.	SPTM [358]; NOLO [351]; ReLIC [203]; NavProbe [383]; SparseNav [396]
Instruction and decision examples	Instruction–landmark, instruction–subtask, or image–decision pairs. Interpret a new instruction or select a move.	LM-Nav [359]; A ² Nav [360]; Select2Plan [352]
Outcome feedback	Outcomes, failure notes, terrain response, or corrective dialogue. Revise guidance or traversal costs.	VLM-GroNav [353]; CMMR-VLN [361]; HAM-VLN [362]; Talk2Escape [384]

Route demonstrations specify where to go in sequence. A tour can teach a particular passage through an environment: pass the reception desk, traverse a corridor, and stop by the stairs. The ordered reference remains relevant even when execution deviates from the original camera poses. Zero-Shot Visual Imitation [355] treats demonstration images as successive goals for a pretrained goal-conditioned policy, including navigation with a TurtleBot in unseen offices. RPF [356] instead encodes views and recorded actions into a path memory that guides execution under motion noise and scene changes. Sparse visual memory [357] reduces the retained views while preserving the action sequence, examining how much visual evidence path following needs. These approaches differ in whether the reference specifies visual subgoals or an action-bearing route. Their shared mechanism aligns present progress with the demonstrated route.

Environment observations reveal where places and objects are. A preview or earlier visit can reveal the building without prescribing the next route. SPTM [358] converts an exploration recording into a graph of visual locations, then selects reachable waypoints for different goals using fixed retrieval and control networks. NOLO [351] conditions navigation actions on a scene video, while ReLIC [203] accumulates experience across episodes with changing starts and goals in one home. NavProbe uses a spatial index to find visual and geometric evidence relevant to the current subgoal [383]. Its executive consults these records when the available evidence is insufficient, then revises the subgoal before invoking navigation skills. The distinction is

how local knowledge becomes available: a graph, a supplied sequence, policy memory, or actively retrieved evidence. A new goal can reuse these facts; a new building requires acquiring them again. CLUE additionally lets exploration revise what the task requires: hypotheses about relevant objects and functions are tested against an online language-embedded map [406]. Context acquisition then resolves an underspecified objective as well as an unknown route.

The current instruction can also decide which facts are worth acquiring. SparseNav retains traversability geometry but grounds semantic landmarks only when the active sub-instruction makes their location useful [396]. Its sparse landmark memory and instruction-progress record then guide waypoint selection. This separates maintaining a navigable map from repeatedly labeling every visible object. Together with NavProbe’s targeted retrieval, it treats context construction as a decision about missing information. The resulting context is organized around the active instruction and the evidence needed for its next navigation decision.

Instruction and decision examples teach an interpretation. Example pairs can specify how a navigation request should be translated into executable choices. LM-Nav [359] supplies instruction–landmark examples to a language model; the extracted sequence is grounded in a topological graph and executed by a pretrained navigator. A²Nav [360] extends this interpretation to action demands: passing a doorway and entering it require different subtasks even when the landmark is identical. Its prompted parser selects among navigators trained for these demands. Select2Plan [352] moves example conditioning closer to the decision itself by retrieving image–answer pairs for movement or waypoint selection. Thus the example can teach a symbolic decomposition or a visual choice. Unlike a route recording tied to particular places, such a pairing can explain how to interpret a new query using an existing execution interface.

Outcome feedback revises what should be attempted next. Execution adds evidence that a successful demonstration may omit: a branch was unhelpful, a stopping decision was premature, or a visually open surface caused slip. CMMR-VLN [361] stores successful routes and notes about the first incorrect decision for retrieval in later navigation. HAM-VLN [362] attaches failure notes to visited places during an episode, making them available when those regions become relevant again. VLM-GroNav [353] grounds terrain judgments in physical interaction and updates waypoint selection and local traversal costs. Talk2Escape acquires additional guidance when retained evidence cannot resolve a failure: detected looping or trajectory divergence triggers a landmark-grounded question to a human or oracle [384]. The reply corrects navigation context. Feedback thus revises a decision through remembered outcomes, physical measurements, or newly requested clarification. These lessons remain useful while the associated instruction, place, and physical conditions still apply.

The four navigation contexts retain different things: a prescribed route, facts about an environment, an interpretation rule, or a conditional lesson. Their validity therefore changes at different boundaries. A new goal can reuse a map but replace the route; a new building invalidates local locations while leaving an interpretation convention useful; changed terrain can invalidate an earlier traversal estimate. Section 7.1 translates these distinctions into controls. Combining navigation with manipulation adds procedural state: after reaching a pickup site, the robot must establish whether the object was acquired. The next subsection examines these dependencies across a complete procedure.

6.3 Long procedures and irreversible operations

Long procedures change the teaching problem from selecting an action to maintaining a dependency structure. Interaction phases and sequential end-effector goals can expose it explicitly, including in teleoperated RGB-D demonstration interfaces [93]; a context-conditioned policy must retain the corresponding distinctions in its representation. An unfinished step, a completed step, and a failed prerequisite require different continuations even when the current image is similar. Long procedures therefore require correspondence between the teacher’s successful sequence and the robot’s interrupted or incomplete execution, beyond retaining a longer record. The analysis moves from dependency tracking to the additional constraints imposed by irreversible steps.

Irreversible procedures add a constraint on how progress and failure can be established. In laboratory workflows, intermediate products must remain valid while the robot decides how to continue. Benchmarks spanning atomic skills, compositions, and extended workflows expose the interaction of procedural inference




Figure 21 Three decision roles of context: (a) task specification selects the object and destination; (b) interaction guidance conveys how the cloth should move onto the plate; (c) execution-state tracking uses the result of removing the towel to advance from uncovering the target to placing the lemon. Representative implementations are VIMA [100], UniSkill [30], and GPT-Policy [94], respectively.

and dexterity [95]; household planning with low-level embodiment abstracted away isolates dependencies among steps [307]. Varying the supplied procedure within these settings would test whether contextual teaching preserves dependencies through the available executor.

Long procedures also couple memories with different validity periods. Progress memory records which steps have been completed; physical memory estimates how to perform the remaining operation. Zeva’s two-timescale analysis tests reuse of action effects across reset attempts [88]. In a laboratory workflow, however, a physical reset may be impossible: repeating an insertion and repeating a dose have different consequences. The informative next interaction must therefore respect both the taught dependency structure and the state of the intermediate product.

Reversibility determines how the recovery mechanisms in Section 5 can be used. A placement may permit local correction; dispensing or fastening can require inspection before a retry. The application therefore needs evidence of completed operations, alongside the demonstrated dependency structure, to decide whether a failed step can be resumed.

The distinctions in Figure 21 concern what information is missing from the next decision. In (a), *task specification* identifies the required object–destination relation among alternatives already present in the scene. In (b), *interaction guidance* conveys the motion relation that should survive the change from a human hand to a robot gripper. In (c), *execution-state tracking* establishes that the plate has been uncovered, making lemon placement the next feasible step [94]. These examples instantiate the information roles in Section 2: panel (b) combines a specified interaction with cross-embodiment correspondence, while panel (c) updates execution state under an unchanged goal. They are complementary information requirements within a task, and can be supplied together through language, demonstrations, and interaction history.

A new procedure can require either a new arrangement of available operations or a motion the current policy has not acquired. Assembly adaptation through fast weights [10] and GEN-1.5’s brush-to-dustpan example after five minutes of fine-tuning [98] place new evidence into parameters. They provide a useful comparison with fixed-weight teaching: context can select, order, and parameterize available operations, including new motions within the executor’s capabilities; further policy learning can expand those capabilities. This separation localizes long-procedure failures to task interpretation or motor competence.

Long-procedure evaluation must separate adherence to a new ordering from adaptation to changed physical response. Body and dynamics adaptation isolates the second demand: how a familiar requirement is realized when the physical system changes.

6.4 Adaptation to body and dynamics

Body adaptation isolates physical response while keeping the requested task specified. Locomotion makes this separation explicit; manipulation introduces analogous uncertainty through payloads and contact. Their combination in mobile manipulation then requires procedure and progress memory alongside the physical estimate. A direction or speed command can already specify the goal; recent action consequences reveal how the current body and terrain respond. Learned adaptation modules address this physical uncertainty [39], while longer-context experience broadens the setting toward different morphologies, actuator changes, and evidence retained across attempts [40]. A broader range of bodies expands the response variations that the policy must infer from its interaction history.

Locomotion also separates a change in the requested behavior from a change in the body that realizes it. Example-conditioned foot-contact generation can specify a gait-like procedure for a pretrained controller [96]. Active tracking across wheeled, quadrupedal, and aerial platforms instead uses a shared velocity interface, with body-specific realization delegated to lower-level control and visual history identifying the moving target [97]. The shared interface supports cross-platform execution by assigning target tracking to visual history and body-specific motion to lower-level control. The contextual contribution depends on whether the evidence specifies a gait, tracks a target, or identifies an unfamiliar response.

Manipulation introduces analogous uncertainty through a new gripper, payload, or deformable object. RopeFormer retains action–response history across trials on unseen ropes, adapting rotation and whipping with fixed policy weights [364]. MetaPusher instead updates a dynamics model during pushing and revises the plan for the unfamiliar object [368]. The goal can remain fixed in both cases: interaction reveals how to realize it under changed physical conditions. Combining this evidence with demonstrations requires distinguishing a changed task requirement from a changed response.

Histories of a stable command under different body or terrain responses provide the training relationship for this adaptation. Mobile manipulation combines this requirement with procedural teaching and the spatial knowledge discussed in Section 6.2. In a delivery procedure, the example specifies the destination and handling order, lifting an object changes the payload, and execution memory records whether pickup has completed. The contextual learner must retain the procedure and progress while updating the physical estimate as the load changes. Evaluating the combination requires both adherence to the demonstrated handling and reliable transport under the changed response.

Procedural teaching and response adaptation become especially valuable when the same capabilities serve recurring tasks. Their practical benefit then depends on which acquired knowledge can be reused and which conditions must be checked again.

6.5 Task changeover, teaching, and reuse

Frequent procedural change over a reusable motor repertoire gives contextual teaching a concrete application. Two kitting orders can use the same parts and grasps but require different placement sequences or handling conventions. A demonstration is useful when it communicates that difference beyond the bill of materials; a shorter instruction can suffice for a fully specified count or destination. Across batches, the chosen sequence persists while completed-step memory resets. A changed tool requires a new physical estimate while the order may remain valid. This separates three parts of changeover: specifying what changes, acquiring the evidence needed to validate it, and retaining the result for subsequent batches.

The recurring change determines the economical teaching channel. Counts or destinations may be stated directly; a fold can require a sequential example; a changed material may require a test interaction. Precision and irreversibility determine the verification needed before accepting the new procedure (Section 5). Changeover therefore measures the complete acquisition process, including the evidence needed to validate it.

Table 14 Teaching and reuse interfaces relevant to production. GEN-1.5, S1, and GLOW condition execution with fixed weights; physical feedback can instead revise programs or train policies.

System	Task evidence	Adaptation	Context or artifact
Demonstrations condition robot execution			
GEN-1.5 [98]	Sensors; trajectories	Physical prompting	Demonstration; history
S1 [99]	Video demonstration	Video conditioning	Demonstration
GLOW [389]	Single human video	Contextual action generation	Encoded skill context
HOST [23]	Human video; robot progress	Predict; decode	Demonstration
Physical feedback revises procedures and learning experiments			
ASPIRE [47]	Task; execution traces	Diagnose; repair	Repair guidance
ENPIRE [89]	Task; physical trial outcomes	Revise code; train policy	Programs; policies

The relevant endpoint is repeatable execution at the required quality. Reusing a task specification can reduce preparation for later batches, while changed tools or materials create a new validation cost. This distinguishes the value of contextual teaching at changeover from the value of a retained procedure during production. Section 7.4 measures both through teaching, interaction, and operating effort.

Teaching and reuse interfaces. Teaching can specify a task for existing competence or revise the resources used in later execution (Table 14).

Sensorimotor and video prompts specify tasks without retraining [98, 99]; HOST instead predicts task evolution before decoding actions [23]. GLOW reports watering and box-storage examples that reuse encoded skill context across changed objects or positions [389]. The qualitative cases illustrate reuse of supplied teaching, while aggregate benchmark scores summarize broader capability. Context interventions isolate the demonstration’s contribution (Section 7.1).

Outcome-driven reuse changes what is retained. ASPIRE preserves repair guidance [47]; ENPIRE can revise a control program or train a policy [89]. Prompt preparation, physical trials, and training incur different changeover costs. Their value depends on whether the change requires another procedure, correction of a recurring failure, or a motion outside existing competence. The next chapter evaluates these contributions to task acquisition and reuse.

Takeaways

- ① Applications determine the missing information: an object relation, spatial knowledge, procedural state, or physical response. Each requires teaching or feedback that resolves the corresponding uncertainty.
- ② Unclear requirements call for teaching; uncertain execution conditions call for observation or interaction; missing motor competence calls for an expanded executor or additional skill learning. These deficits require different remedies.
- ③ Task changes create value for new teaching; repeated operations create value for retained procedures. Acquisition and verification costs determine that value, especially when actions are irreversible or time-critical.

7 Evaluation: When Does Context Help?

Evaluation traces the contribution of context from changed evidence to physical success. Familiar scenes, executor capabilities, and retry opportunities provide the comparison conditions needed to isolate that contribution. Does changed evidence produce the appropriate change in behavior? Does that change survive physical execution and transfer? Does retaining the lesson improve a later attempt enough to justify its acquisition cost? Context interventions, intermediate diagnostics, and retention tests separate these questions. Reported comparisons ground the analysis; proposed controls test the learning mechanisms behind them.

(a) Teaching change**(b) Object change**

Figure 22 Two tests of contextual transfer. (a) Changing the demonstration while holding the robot scene fixed changes UniSkill’s predicted continuation [30]. (b) LIBERO-PRO changes objects to test target selection [106]. Together, these axes motivate tests that vary teaching and object identity independently, measuring both the inferred requirement and its physical realization.

7.1 Context use and transfer across tasks and environments

Context interventions test sensitivity to teaching; changes in tasks, objects, environments, or bodies locate its transfer boundary. Varying evidence quantity, retention distance, and distractors tests access to the required information.

Sensitivity to task-defining context. Demonstrations specifying different feasible behaviors in the same robot scene isolate the influence of teaching. In a kitting task, changing the demonstrated part count should change the selected count; changing the demonstrator while preserving that count should preserve the requirement. Together, meaning-changing and meaning-preserving interventions test what information the learner uses. Success without context provides the prior-only comparison. A fully specified target tests whether the required behavior is executable.

Instruction benchmarks illustrate why these controls matter. LIBERO varies spatial, object, goal, and procedural factors [228], yet task-index embeddings can achieve high performance when the scene reveals task identity [309]. LIBERO-CF supplies alternative feasible instructions within familiar layouts and finds that tested VLAs often continue the scene-associated task [107]. For demonstration ICL, replacing a video must therefore change the intended order or relation, rather than merely perturb its appearance. Robustness tests supply the complementary invariant: preserve the required procedure across viewpoint, object appearance, or instruction rephrasing [106, 310]. RoboStressBench separates material, viewpoint, lighting, and geometry stresses to diagnose upstream visual recognition, reasoning, and planning [377]. VidPair-Halluc instead pairs similar backgrounds with different foreground semantics [378]. Applied to demonstration evaluation, this design would test whether inferred task content follows the demonstrated event rather than scene similarity; a subsequent rollout tests physical realization.

Show-Harness’s video-order comparison provides task information left unspecified by the no-video instruction [31]. An equally informative textual order is the relevant control for a claim about video’s particular advantage. Comparing language-only, demonstration-only, and joint context with matched task information separates modality from information content. Holding the instruction fixed while changing the demonstrated convention tests learning from examples; holding the demonstration fixed while changing a correction tests use of feedback. Discovering an unfamiliar action convention proved harder: without supplied conventions, the system succeeds in one of 20 episodes and infers 23.3% of mappings correctly. Interpreting a demonstrated procedure and discovering an unfamiliar action interface thus require separate evidence.

Table 15 Proposed evaluation controls and resources that illustrate the corresponding capabilities. Comparisons hold other task conditions fixed; cross-attempt tests separate physical resets from memory resets.

Capability	Intervention	Readout	Example resources
Context use	Change teaching	Task-consistent actions	LIBERO-CF [107]; LMAct [102]
Relation transfer	Hold out relations	Relation adherence	VIMA [100]; MOSAIC [12]; Imitator Game [101]
Evidence retention	Retain or clear history	Recall; memory updates	RoboMME [103]; MEMOBench [104]
Cross-attempt reuse	Separate reset types	Gain; interference	RopeFormer [364]; Zeva-Ego [363]
Physical transfer	Change dynamics	Adaptation; goal fidelity	MemMimic [105]; PACE-Bench [308]
Human-data transfer	Vary human data; fix robot data	Trained-policy success	H2RBench [366]
Scene robustness	Change scene	Target; outcome	LIBERO [228]; LIBERO-PRO [106]
Execution steering	Change command	Redirection; success	ReSteer [108]
Physical execution	Supply target or prerequisites	Target selection; success	RoboSemanticBench [109]; Behavior-Skill [110]
Route teaching	Change route; retain goal	Path; landmark-order adherence	RPF [356]; sparse visual memory [357]
Spatial adaptation	Change preview; reset memory	Success; route efficiency	NOLO [351]; ReLIC [203]
Example interpretation	Change examples; retain query	Parsed subtasks; selected moves	A ² Nav [360]; Select2Plan [352]
Outcome reuse	Remove notes; retain retries	Repeated errors; path cost	CMMR-VLN [361]; HAM-VLN [362]
Traversability	Withhold terrain feedback	Success; physical disturbance	VLM-GroNav [353]

Sensitivity to teaching and robustness of execution require different observations. A changed predicted continuation can expose a changed task interpretation; a rollout establishes whether the robot realizes it. UniSkill’s decoded futures and LIBERO-PRO’s object substitutions illustrate these complementary views (Figure 22). The former diagnoses an inferred skill within skill- and agent-based execution; the latter tests target selection and physical behavior. Evaluation must relate both to a controlled change in the supplied evidence.

Table 15 pairs evaluation questions with resources that make the relevant variation possible. The rows concern different dependencies, so their metrics are interpreted within the stated task setting. In particular, the human-data row measures transfer into trained policy weights, whereas contextual tests intervene on teaching or retained evidence at deployment. This distinction connects evaluation to the data roles in Figure 14.

Object substitution and relation transfer. Task novelty has to be defined at the level of the relation being learned. Task suites make adaptation to held-out tasks measurable [227, 311], but withheld names can combine several kinds of change. New object combinations and task templates probe different generalization in VIMA [100]; the Imitator Game separates changes in demonstrated motion from changes in functional intent [101]. A familiar procedure on a new body tests realization, whereas a new ordering of available actions tests contextual composition. A useful transfer claim identifies both the unseen relation and the motor competence available to realize it.

Object substitution supplies a complementary test to changing the demonstrated requirement. Following Table 6, retain the requirement while replacing only the manipulated object, only the receiving object, and then both. Within each condition, compare the correct demonstration, no demonstration, and a demonstration specifying another feasible behavior in the same scene. Comparing object conditions measures substitution

scope; comparing demonstration conditions measures dependence on teaching. To isolate the demonstration's contribution, keep the instruction fixed or hold its task information constant. Report task interpretation and physical completion separately, verifying the latter through robot execution.

The choice of substitutions determines which transfer capability is tested. FUNCTO distinguishes instance and category changes [172], whereas object-pair transfer exposes changes at both ends of a relation [347]. Functionally compatible substitutes with different appearances, paired with visually similar but incompatible distractors, can test whether correspondence follows the taught interaction. Training exposure, demonstration count, human correspondence or segmentation assistance, neural updates, and retry opportunities affect the interpretation of this comparison. Undisclosed pretraining leaves training novelty uncertain even when context dependence is established.

Withholding a relation while preserving familiar objects and executable actions isolates relation novelty; substituting objects while retaining the relation tests its physical transfer. Cross-robot imitation and paired human-video interfaces then extend the change to embodiment and observation [12, 312]. Whether the demonstration remains available is a further axis: executing with a supplied reference and retaining its lesson after that reference is removed are evaluated separately in Section 7.2.

Capture conditions also affect the transfer axis. RoboReel links 2,000 human recordings from four views to simulated trajectories in ten environments, separating clean and distracted recordings from evaluation-only distraction and pose-defined tasks [312]. These variations test whether the same demonstrated requirement survives a changed observation interface. Reporting the neural update rule alongside these variations identifies whether transfer uses context or task-specific policy fitting.

The real-to-sim benchmark H2RBench tests how human demonstrations support policy training [366]. It reconstructs four real manipulation tasks in simulation, varies human-video quantity at a fixed robot-demonstration budget, and compares relative simulation and real-robot performance. Its shared training protocol tests human-data transfer; frozen-policy context interventions test whether a newly supplied example changes execution. Both matter when a system combines human-video pretraining with contextual adaptation.

Navigation across goals and environments. The four context types in Table 13 require different interventions. Changing a route demonstration while keeping its destination fixed tests adherence to the taught path or landmark order. Replacing a scene preview and clearing memory between goals test spatial evidence and within-environment reuse, respectively; NOLO's context-removal experiment [351] and ReLIC's cross-episode attention comparison [203] examine these contributions. Changing instruction-decision examples while retaining the current query tests the demonstrated interpretation. Removing failure notes while retaining backtracking tests whether a remembered outcome changes later decisions, as in HAM-VLN's reflection ablation [362]. A new building then tests acquisition after environment-specific locations cease to apply.

Success weighted by path length (SPL) combines reaching the goal with route efficiency [354]; a prescribed route additionally requires an adherence measure, since a shortcut can reach the same destination. Success across successive goals measures environment reuse, while repeated wrong turns and stopping errors expose the contribution of feedback. Goal visibility, stopping distance, action discretization, and access to maps or localization define the comparison. Terrain adaptation adds physical cost: VLM-GroNav [353] reports trajectory length and IMU energy density alongside success. Exploration and retries spent acquiring context must be included when relating these gains to total interaction.

Simulation, physical navigation, and recorded-scene decisions provide different evidence about execution. NOLO's unseen-layout transfer is evaluated in simulation; its real maze uses separate training and context recordings from that maze. Select2Plan [352] combines simulated object navigation with waypoint assessment on recorded rover scenes, whereas LM-Nav [359] executes routes on a physical mobile robot. HAM-VLN's reported comparisons use 100-episode subsets of unseen simulation splits. Together, these settings span simulated transfer, recorded-scene decisions, and physical route execution.

Instruction adherence needs alternatives that the scene alone cannot resolve. RoboFollow pairs shared scenes with several feasible behaviors and separates stage-level intent from execution scores [399]. Its layout and semantic recombination tests distinguish choosing the requested object or procedure from physically completing

it. This provides an operational counterpart to the context interventions above: if only one behavior is plausible in a scene, high success can conceal weak use of teaching. RoboFollow tests instruction grounding. Testing unfamiliar rule learning requires a held-out rule family and examples that distinguish its candidate rules.

Evidence quantity, retention, and distractors. The amount of teaching evidence, retention distance, and distractor load test different properties. LMAct varies the number of expert episodes and includes replay controls, showing that more examples often provide little benefit to the tested models [102]. A flat curve can reflect early saturation, weak context use, or inadequate control. Comparing an ambiguous task with a fully specified version helps distinguish these causes. Memory benchmarks make the missing evidence more explicit through object, spatial, sequential, or imitation demands [103, 313]. The necessary history depends on the hidden information: AMAGO reports a MetaWorld setting whose task can be identified from short history [202].

Checkpoint-based evaluation localizes memory failures. MEMOBench distinguishes Storage, Update, and Compression demands across 30 history-dependent tasks [104]. These test whether decisive events survive retention, revision, and compression. Supplying a missing decisive record directly at the policy input, in its expected format, can help distinguish storage or retrieval failure from failure to use available evidence. Intermediate interventions then test physical realization (Section 7.3). WorldLines extends state validity to household traces: Memory QA and a downstream planning probe test whether remembered state changes support current answers and action preconditions [375]. EvoEmpirBench isolates dynamic spatial reasoning in locally observable mazes and match-2 games; Agent-ExpVer distills and revises experience across tasks [374]. These settings separate state recall and revision from physical execution, complementing robot rollouts with more targeted memory tests.

7.2 Sources of improvement and comparable evidence

Fourteen reported comparisons (Table 16) connect context use, physical transfer, and later learning to observable outcomes. Their findings motivate controls that attribute gains to supplied evidence, parameter updates, or retained artifacts.

What reported comparisons establish. The first group shows why evidence quantity and usefulness must be separated. ICRT’s training-data comparison links context use to task alternatives within a scene; LMAct shows that adding expert episodes alone often yields little benefit [13, 102]. BPP’s drawing comparison favors demonstrations that specify a trajectory, whereas Show-Harness’s action-convention comparison exposes information needed to ground unfamiliar commands [31, 54]. NOLO’s context-removal comparison extends this analysis to scene previews [351]. Together, these results connect improvement to the information supplied by context.

The second group connects transfer to the execution interface. MT3’s low-demonstration advantage relies on alignment followed by interaction replay [36]; part-wise warping improves the tested transfers with supplied part correspondence [347]. Demo-JEPA’s comparison with an action head shows that the preferred realization can change with task familiarity [22]. RAPID’s scene-variant ablation adds a complementary question: whether the acquired program is verified beyond the demonstrated configuration [390]. Transfer therefore depends jointly on the represented relation, correspondence assistance, motor competence, and conditions used to validate the result.

The third group distinguishes using retained experience from acquiring a better learner. Zeva tests access to earlier interaction; RTCF and TraceFlow test action correction from successful or failed traces; FARE tests selective revision of the history conditioning prediction [88, 195, 392, 402]. Their scope differs: TraceFlow’s ordering gains coexist with unresolved simulation deficits, while FARE’s always-search variant underperforms its base policy. LMPC instead trains a successor between sessions and measures its response to new teaching [117]. Cumulative success over retries measures yet another quantity, as in Zeva’s evolution milestones. These comparisons locate gains in a specific use of experience; the interventions below test whether that use explains the outcome.

Controlled attribution. Table 17 separates interventions on the training checkpoint, deployment evidence, adaptive parameters, and retained artifacts. Evidence-acquisition controls additionally test how informative observations are obtained.

The first control compares checkpoints under the same context interventions; the remaining controls change what is available during execution. Repeated-task and transfer trials then establish whether the identified dependence improves subsequent behavior.

Acquiring the ability to use context. Contextual training is useful when it improves the model’s response to task-defining evidence, beyond improving execution generally. Explicitly paired examples offer one route [1, 12]; multimodal prompts, sensorimotor streams, and corrective histories expose other dependencies reviewed in Section 4.2 [10, 13, 100]. Each checkpoint can be evaluated on the same tasks with informative context, absent context, and context specifying a different feasible behavior. The within-checkpoint comparison measures the effect of evidence; comparing that effect across checkpoints tests whether training improved its use. Performance with a fully specified task separately checks the available execution competence.

Table 16 Reported comparisons of context use and transfer. $\rightarrow$ denotes baseline-to-variant change; $\uparrow/\downarrow$ denote increase/decrease; $>$ denotes ordering. SR: success rate (%); pp: percentage points; fractions: successes/total; qual.: qualitative; BC: behavior cloning. Comparisons are within study; weights stay fixed during deployment (LMPC trains successors between sessions). Full protocols and source locations: Supplement S1.

Study	Target	Contrast	Readout	Setting
I. Training and task evidence				
ICRT [13]	Training data	DROID-only / multi-task	DROID-only: no progress	Shared test tasks
BPP [54]	Prompt	Goal image $\rightarrow$ demo	Fidelity $\uparrow$ (qual.)	Unseen drawings; matched start
LMAct [102]	Context size	0 $\rightarrow$ 512 episodes	Often little gain	Games; simulated control
Show-Harness [31]	API semantics	Arbitrary names; no conventions	1/20 successes	Fixed harness
NOLO [351]	Scene context	No video $\rightarrow$ video	SR: 33.58 $\rightarrow$ 43.65	Habitat; 3 seeds
II. Transfer and physical realization				
MT3 [36]	Action reuse	BC $\rightarrow$ retrieval; whole $\rightarrow$ phases	SR $\uparrow$; demos $\downarrow$	New instances; align + interact
Part-based transfer [347]	Warp granularity	Whole object $\rightarrow$ parts	11/27 $\rightarrow$ 23/27	Mug-rack placement; assisted parts
Demo-JEPA [22]	Action interface	Action head / planning	Known, real: head $>$ plan; unseen: plan $>$ head	Sim. + real; task-group means
RAPID [390]	Verification	Scene variants: off $\rightarrow$ on	SR: 53.2 $\rightarrow$ 75.9	8 tasks; 50 scenes/task; 3 runs
III. Retained experience and subsequent learning				
Zeva [88]	Persistent memory	Off $\rightarrow$ on	SR: +10–20 pp	5 real tasks; brief trace retained
RTCF [195]	Memory correction	Off $\rightarrow$ on	SR: 86.4 $\rightarrow$ 88.4	Fixed bank; 2,000 episodes/variant
TraceFlow [392]	Trace guidance	Base $\rightarrow$ guided	Ordered: 21/50 $\rightarrow$ 39/50; sim. aggregate: unchanged	1 real packing task; simulation suite
FARE [402]	History revision	Base $\rightarrow$ selective; always-search ablation	SR: 91.5 $\rightarrow$ 93.2; always-search: 91.0	RoboTwin Easy; fixed checkpoint
LMPC [117]	Successor training	Base $\rightarrow$ successor	SR: 39.4 $\rightarrow$ 66.3	Held-out teaching tasks

Table 17 Attribution controls expressed as target–intervention–readout. $\leftrightarrow$ denotes paired test conditions; $\times$ crosses checkpoints with context conditions. Match tasks, physical resets, base execution resources, and attempt budgets. Measure both behavioral change and task completion.

Target	Intervention	Readout
Training checkpoint	Checkpoint $\times$ context	Learned context use
Task teaching	Original $\leftrightarrow$ replacement	Requirement adherence
Retained history	Retain $\leftrightarrow$ clear	Use of prior evidence
Neural parameters	Adapt $\leftrightarrow$ restore	Adaptation gain
Reusable guidance	Retain $\leftrightarrow$ withhold	Lesson transfer
Skill library	Expanded $\leftrightarrow$ base	Use of new skills
Active observation	Enabled $\leftrightarrow$ disabled	Acquisition benefit

Using new teaching and retained observations. Supplied teaching and retained history require different interventions. Changing an example with the checkpoint, scene, and executor fixed tests whether its task requirement affects behavior; video-conditioned policies and predictive interfaces provide concrete settings for this comparison [14, 24]. Clearing history at a matched state instead tests access to earlier evidence. Dynamics-revealing transitions [39, 364] and remembered task events [42, 44] resolve different unknowns. Memory-focused settings distinguish these requirements [103, 104]; evaluating the action selected under each history tests whether the recalled evidence is used appropriately. Counterfactual Memory Audit compares alternative histories that yield identical current inputs [385]. It controls policy-sampling randomness and tests the resulting actions under both histories. This distinguishes an action change caused by memory from a justified choice and its physical benefit.

Adapting neural parameters. Parameter adaptation requires a matched-input comparison and a distinct reset control. Meta-imitation changes an initialization through a demonstration loss [7, 8]; fast-weight approaches restrict the updated subset and its timing [10, 11]. With demonstrations and test instances fixed, restoring pre-adaptation weights separately from removing context identifies the contribution of weight changes. Optimization and interaction costs characterize the resources required. This comparison isolates an optimization-based contribution; fixed-parameter ICL holds deployed neural weights unchanged in both conditions.

Reusing guidance and executable skills. Artifact ablations distinguish a retained lesson from an expanded execution repertoire. Withholding guidance tests whether acquired knowledge improves fresh inference, as motivated by correction memory and repair guidance [46, 47]. Restoring the pre-acquisition library instead tests newly acquired routines while preserving the original skills [50, 52]. A geometric archive can be treated similarly [36]. Fresh task instances, the base executor, and attempt budgets are held fixed so that the contrast concerns reuse rather than additional opportunities to solve the task. Table B3 identifies the artifact associated with each branch.

Acquiring new evidence. Active sensing changes the evidence available to the learner, complementing interventions that hold the evidence fixed. VA-Bench gives MLLM agents RGB demonstrations, camera control, metric target commands, and execution feedback through a fixed controller [344]. Its active-versus-passive viewpoint comparison isolates evidence acquisition, while held-out geometry and long-horizon composition test whether that evidence supports transferable execution. These interventions distinguish obtaining a better observation from interpreting the same observation more effectively.

Matched retention and reuse tests. Figure 23 separates acquisition, retention, and transfer through a drawer-placement example. In the recorded episode, the robot initially reaches for scissors; human feedback establishes that it must open the drawer first [46]. Panel (a) identifies the acquired ordering constraint. Panels (b,c) propose comparing access to that constraint with its removal: repeating the task tests retention, while replacing the object and drawer tests transfer. The illustrated continuations make the predicted effect observable before completion—opening first versus premature grasping. Repeated trials measuring success and further human corrections would establish whether that difference improves execution.




Figure 23 Testing the use of a learned ordering constraint. (a) Feedback establishes that the drawer must be opened before grasping the object. (b,c) Conceptual illustrations contrast following that order with premature grasping on repeated and new instances. Panel (a) follows the DROC example [46]; panels (b,c) propose candidate outcomes for retention tests. Success and further human corrections provide the readouts.

The retention comparison holds task instructions, physical initialization, base competence, and resource budgets fixed; access to the selected lesson is the changed factor. Independent scene resets remove physical progress, and matched retry budgets control additional opportunities for success. Selective resets identify dependence conditional on the components left intact; redundant records can preserve the same information after one carrier is removed. Training exposure remains a separate axis when pretraining data are undisclosed. Figure 6 identifies the possible carriers and timescales, while Table 17 specifies interventions on each. This links a particular acquired fact to its availability and subsequent behavioral effect.

Acquiring an artifact consumes experience before it can save experience. Separating play from downstream testing and comparing against additional retries, as in RATs, exposes this trade-off [52]. Separating synthesis feedback from fresh validation serves the same purpose for program and oversight revision [90, 218]. A successful acquisition trial establishes that a candidate worked once; a retention benefit on later instances establishes useful reuse. Relating that gain to the trials and corrections spent acquiring it determines whether the retained knowledge reduces total learning effort.

Adherence to teaching, physical success, and acquisition cost characterize reuse within a shared task setting. Success at a fixed attempt budget or attempts to a common criterion provide comparable outcomes, including unsuccessful trials; repeated matched trials quantify variation. The next section traces these gains through the stages of physical execution.

7.3 Physical realization and learning across attempts

Channel and intermediate interventions trace contextual information into physical outcomes. Independent physical and memory resets then test whether retained evidence improves later attempts and recovery.

Channel interventions complement memory resets. With the task, checkpoint, and recorded history matched, removing language, visual, proprioceptive, or contact inputs tests which evidence supports adaptation. The comparison must retain the policy’s valid input format and distinguish a missing measurement from a measured zero. Goal completion and demonstration adherence should also be reported separately when contact is specified: a handle-grasp requirement calls for both successful transport and preservation of handle contact.

Intermediate interventions localize failure along the four control interfaces. For a contextual policy, supply the relevant demonstration segment; for geometric transfer, supply the correct correspondence and reference; for

predictive control, supply a task-consistent future; for an agent, supply the correct executable specification. Each intervention must preserve the controller’s expected format. Recovery after the substitution locates a bottleneck in the replaced stage or its input. Failure in both conditions calls for tests of downstream execution and interface compatibility. The contrast is most informative at the earliest decision where inferred and required behavior diverge. Target-selection assessment and restored-precondition tests provide related diagnostics [109, 110]. Both conditions require measuring prescribed paths, order, or contacts alongside the terminal state.

Physical adaptation requires evidence about the changed condition. MemMimic distinguishes event recall from adaptation to initially unknown physical response, while PACE-Bench varies physical conditions with a stable goal and interface [105, 308]. For synthesized teaching videos, the video-generation benchmark H2R-Bench evaluates goal, event, contact, and embodiment consistency before robot execution [314]. These tests are useful when tied to the specific information that the contextual policy must recover. Probe-then-commit audits further separate whether a belief covers the action-relevant truth from whether the executor predicts failure accurately [412]. Their offline ground-truth checks show why sharper beliefs alone need not justify physical commitment.

Independent physical and memory resets separate the sources of improvement across attempts. Resetting the scene removes task progress, whereas resetting memory removes acquired evidence. Retaining each separately distinguishes physical continuation, contextual adaptation, and additional chances from repeated execution.

For trial-memory adaptation, the reuse gain at attempt n is $\Delta_n^{\text{reuse}} = S_n^{\text{keep}} - S_n^{\text{reset}}$. Here S_n is the same predefined success or adherence rate at attempt n ; the superscripts keep and reset specify retained and cleared memory. Their difference Δ_n^{reuse} measures the benefit of retaining evidence from earlier attempts. The comparison holds task distribution, physical initialization, executor, other retained artifacts, earlier attempt count, and execution budget constant; only access to evidence from those attempts changes. Per-attempt rates isolate the benefit of retained evidence, whereas cumulative success also benefits from independent retries. A subsequent tool or dynamics change tests whether the learner revises a lesson whose physical conditions have become obsolete.

Recovery evaluation connects the detected failure to the acquired evidence, the revised decision, and the outcome of that revision. End-to-end success includes all attempted tasks; conditional repair success concerns the subset of detected, eligible failures. Distinguishing autonomous repair, human correction, and external reset makes their respective contributions visible. Later attempts establish whether the correction remains useful, while newly taught valid procedures test whether the failure detector follows the current task reference. Section 5 describes the corresponding mechanisms.

RoboRecover makes the starting deviation reproducible by replaying action prefixes, then evaluating continuation under the original goal [400]. Its 2,000 scenarios distinguish the stage to resume from the deviation to correct, across RoboTwin and LIBERO. Initializing candidate policy history at the recovery point isolates continuation from the reconstructed off-nominal state. A contextual-recovery extension should compare retained and cleared prefixes at that same reconstructed state, preserving the continuation budget. This connects reproducible physical deviations to the selective memory resets above.

7.4 Teaching effort, response delay, and reuse cost

Contextual adaptation has three operational timescales: acquiring the task, responding to fresh evidence, and reusing the result. Task acquisition includes instruction preparation, demonstration recording, exploration, and validation. Extra search can improve execution with unchanged task information; varying relevant context independently of search budget identifies the contribution of learning [315, 316]. SAIL’s reconstruction and simulated trials illustrate preparation costs beyond model inference [317].

Acquisition is measured end to end by *changeover time*: the interval from receiving a specification to satisfying a stated quality and operating criterion. A common criterion permits comparison of contextual teaching, fine-tuning, trajectory reuse, and programming. Elapsed time, human labor, robot interaction, and resets capture different costs: ENPIRE’s program refinement requires physical trials and verification even with

fixed neural weights [89]. Repeatable task performance therefore provides a more informative endpoint for changeover than a first successful attempt.

During execution, the relevant delay runs from informative feedback to the affected command. Model inference time, command rate, and servo frequency describe different parts of this interval. Cached demonstration features can be reused, while a new contact observation must reach control before the corresponding motion is committed. SimpleMemVLA reports latency falling from 1.02 to 0.68 seconds with identical outputs through scheduling and shared-prefix processing [45]. This scheduling gain reduces response delay while preserving the decision itself. State-dependent steering tests such as ReSteer make the arrival point of feedback explicit [108].

Repeated batches test whether acquisition costs are repaid. Accepted output and intervention time across repeated tasks quantify this return; changing a tool or requirement then reveals the cost of obsolete knowledge. These outcomes motivate the next chapter: broader contextual learning must expand what an episode can teach while preserving the ability to revise what no longer applies.

Takeaways

- ① Context interventions establish dependence on task-defining evidence beyond familiar-scene recognition. Physical execution and held-out relation tests then establish how far that requirement transfers.
- ② Common task outcomes expose the contributions of training, supplied context, adaptive neural state, and retained knowledge. Acquisition and reuse carry different resource costs.
- ③ Teaching effort, feedback-to-action delay, and cost per accepted execution connect contextual learning to deployment value across repeated attempts.

8 Future Directions: What Limits Transfer and Improvement?

Broader contextual learning has four connected objectives: infer unfamiliar rules, preserve binding physical constraints through abstraction, revise retained corrections when conditions change, and improve how subsequent tasks are learned. The directions below connect these objectives to training, representation, and experience acquisition, with behavioral tests that distinguish progress on each.

8.1 From context dependence to a reusable learning rule

The learning distinctions introduced in Section 2.1 separate familiar-task selection, composition, and inference of an unfamiliar rule. Building on the learning mechanisms in Section 4.4, we ask which training relationships would extend this behavior to unfamiliar requirements.

Table 18 Proposed tests of behavioral transfer in a placement task. Keep available objects and motor actions matched; vary what teaching must supply. Novelty is relative to the controlled training exposure.

Learning claim	What is already familiar	What the example must establish
Familiar-task selection	Object-to-location assignments	Choose a known left- or right-bin routine
Compositional transfer	Placement relations; order constraints	Combine familiar color and order rules
Rule inference	Objects; executable placements	Infer an unfamiliar size-to-slot rule

Table 18 distinguishes selecting a familiar task, composing known relations, and inferring a new rule. An example can identify a routine without teaching a new relation: demonstration benefits under incorrect labels in some NLP settings expose the role of task-format recognition [320]. Conversely, function-learning studies show inference of unseen functions, and controlled regression connects task diversity to more general learning behavior [253, 321]. The illustrative robotic tests keep available objects and motor skills fixed while varying what teaching must establish: a familiar bin choice, a new composition of known constraints, or a size-to-slot mapping from a family excluded during training. Rule inference requires enough examples to distinguish the candidate mappings. BPP’s continuous drawing tasks and the Imitator Game’s functional substitutions provide concrete task variations from which to design these tests [54, 101].




Figure 24 Four research directions: (a) infer taught relations, (b) compose task steps, (c) preserve function across physical changes, and (d) use corrective feedback. Together they extend contextual learning toward interpreting, transferring, and retaining task requirements.

The acquisition routes in Section 4.4 motivate a testable hypothesis: coverage of task relations and coverage of executable motions make distinct contributions to contextual transfer. Broad pretraining may support unfamiliar requirements, while grounded pairs establish precise action correspondence. Varying pretraining scale and explicit pairing independently, with the executor and task information fixed, would test their contributions. Motion coverage, relation diversity, interaction experience, and model size are distinct factors; agentic systems also expose library size and acquisition cost [53]. Relation-level outcomes and continuous errors are needed alongside all-steps success, because thresholded metrics can make smooth component gains appear abrupt [323].

Cross-modal transfer supplies a complementary test. GPT-Policy provides a concrete interface in which a general VLM receives teaching and acts through fixed robot tools [94]. Expressing the same unfamiliar rule through text, demonstrations, or both would test whether its inferred meaning transfers across formats. Matching task information and motor capabilities isolates rule inference from access to additional details, such as a procedural step visible in one format but absent from another.

The task examples in Figure 24 connect these research questions to observable behavior. Panel (a) asks whether a taught relation determines actions in a new scene; (b) extends it to a sequence of dependent steps; (c) changes the contact geometry while preserving the interaction’s function; and (d) uses human feedback to revise the next action. The following sections develop the corresponding teaching, transfer, and retention requirements. Section 8.5 then considers whether experience improves how subsequent requirements are learned.

8.2 Instruction tuning for more complex contextual learning

MetaICL and Flan provide language-model precedents for example-based learning and instruction tuning [113, 324]. For ICL for robots, the training problem is to make new teaching select among feasible continuations within a controlled motor repertoire. The three cases below progress from combining requirements to revising them and resolving missing information.

Figure 25 illustrates these cases through separate kitting episodes with comparable parts and motor skills. Paired, continuous, and corrective records (Section 4.2) provide possible training formats; the intermediate tray states reveal which requirements survive execution.

Requirement binding and composition. Panel (a) combines a demonstrated plate-first order with a request for two plates and three bushings. Reversing the demonstrated order should reverse intermediate tray states while preserving the final inventory. The example separates *compositional complexity*, which varies skill order and dependencies, from *constraint complexity*, which adds jointly binding counts, contacts, or dwell times. Measuring each requirement and their joint satisfaction tests whether a new composition was followed through the method’s intermediate representation.




Figure 25 Illustrative kitting tasks for three contextual learning problems: (a) combine demonstrated order with requested quantities, (b) change the destination while retaining the order, and (c) resolve an ambiguous destination before acting. Intermediate tray states reveal adherence to teaching beyond the final inventory.

Selective revision of a procedure. Panel (b) changes the destination to tray B before placement, retaining the order and quantities. DROC’s language guidance, RACER’s corrective supervision, and RoboTTT’s separation of failed context from corrective targets offer components for learning such revisions [10, 46, 87]. Training should reward the targeted change together with unaffected requirements. Varying correction timing, scope, and precedence then tests *revision complexity*: reaching tray B while losing the taught count is an overbroad revision. Intervening tasks test whether the revised procedure remains available later.

Decision-directed information acquisition. Panel (c) leaves both trays feasible without specifying which to use. Asking for the destination resolves this ambiguity before placement, building on question-based teaching and uncertainty-aware assistance [86, 301]. The acquired answer must preserve the already taught order and quantity. Other uncertainties call for different interventions: a view or probe can reveal execution conditions (Section 5), whereas a missing grasp skill requires an additional controller or motor learning. Behavioral Exploration offers a related action-based route: a policy uses interaction history to select expert-like behaviors that explore beyond previous attempts [333]. Clarification and exploration target different unknowns; training should connect each request or action to the information it supplies.

These cases motivate comparing action pretraining, demonstration-conditioned training, and mixed instruction–demonstration–correction training under a common backbone and experience budget. Withheld compositions and corrections test requirement adherence, physical completion, and later reuse. The next section examines which requirements should remain invariant when objects and execution conditions change.

8.3 Preserving taught relations across physical change

Cross-object transfer in Section 3.5 establishes useful relations when the corresponding parts and admissible substitutions can be identified. The remaining challenge is to infer which demonstrated details are binding and when no valid substitute exists. A handleless vessel may serve a pouring goal but cannot preserve an

explicitly required handle grasp. Similarly, a pause that lets adhesive set is part of the process, whereas an incidental pause may be shortened. These cases require selective invariance: retain the taught requirement while changing only details that its meaning permits. The part-decomposition study supplies cross-category part equivalencies [347]; learning to infer and validate such equivalencies would extend its transfer scope.

The crossed comparisons in Section 7.1 locate these limits by separating changed requirements from changed objects. Beyond successful substitution, a learner must recognize when the available objects cannot satisfy the demonstrated relation and acquire the missing information. Contrasting demonstrations or a clarification can establish whether a pause or contact is required; interaction can reveal whether a geometrically compatible replacement has unfamiliar resistance or dynamics [40, 88]. Correspondence and physical feedback therefore answer different questions: where an interaction should occur, and whether that interaction realizes the intended effect.

The available executor sets the boundary of this transfer. If the revised requirement can be expressed through existing actions or skills, contextual inference can select a new realization. If the necessary motion lies outside that repertoire, further motor learning contributes a separate stage, as in human-video-initialized WHIRL [150]. Such a hybrid process can preserve a context-inferred task relation while learning its physical realization.

Navigation introduces a related separation between reusable strategy and local knowledge. An exploration strategy can transfer across buildings, while a remembered object location is tied to one layout. ReLIC [203] illustrates the value of retaining such local experience across goals. Broader transfer requires identifying which spatial facts remain valid when the environment changes, alongside preserving the task relations considered above.

8.4 Retaining teaching while revising obsolete experience




Figure 26 Illustrative reuse of corrective context: (a) an order error, (b) plate-first feedback, and (c) execution in a new layout with fixed policy weights. The intermediate tray makes retention of the taught order observable before the kit is complete.

Physical transfer preserves a requirement across changes in its realization; temporal reuse must also accommodate new teaching. Video outlines and lower-cost planning latents retain different details [31, 325]. A later correction can make an initially incidental contact or pause decisive, after a task-specific summary has discarded it. Delayed corrections can test what new teaching compact summaries and retrievable source segments support at matched storage and access costs. Preserving current task meaning and enabling future reinterpretation impose different requirements on retained evidence.

Consolidation into weights adds training after contextual acquisition. The $\pi_{0.7}$ coaching example and the RECAP training framework for $\pi_{0.6}^*$ connect interaction or correction to subsequent training [273, 281]. Immediate context-driven behavior and persistence after consolidation or reset characterize different stages of acquisition. Retained programs, records, or policy weights can help later tasks, but obsolete corrections can also interfere with new teaching.

A taught requirement can remain useful after the scene that elicited it has changed. In the kitting example, retained plate-first feedback supplies information absent from the new layout, illustrating how a fixed policy could combine an earlier correction with current geometry (Figure 26). Keeping, omitting, or replacing that

correction tests whether the later decision depends on the acquired lesson. Retention must also survive the representation changes used to make such lessons cheaper to store and apply.

Across successive deployments, the criterion also applies to model upgrades and changing environments [127]. Replacing the interpreter or action interface raises two questions: whether retained teaching still governs behavior, and whether revised requirements can override it. An intervening task can leave the earlier procedure available for later reuse; a new instruction for that procedure can supersede it. Section 7.4 connects this selective retention to reduced teaching and validation effort.

8.5 Physical recursive self-improvement

Physical feedback can improve the current solution or the process used to acquire future solutions. Retained context, program revision, and neural updates provide different means of doing so, with their learning benefit determined by what the successor acquires from new experience. Improved teachability connects these updates to the choice of experience and the cost of acquiring the next task. In the survey’s six-horizon framework, S5 names this subsequent-learning objective and S6 extends it to knowledge exchanged between robots.

Table 19 Feedback-driven improvement through update target, operation, and reusable output. Targets can overlap; successor comparisons measure their contribution to subsequent learning. ENPIRE and REVOLVE denote their program and policy branches, respectively.

Update target	Operation	Reusable output	Examples
Interaction context	Retain	Action–effect memory	RopeFormer [364]; Zeva-Ego [363]
Executable artifacts	Revise	Skills; programs	RoboRSI [348]; ENPIRE [89]; HarnessPAI [398]
Neural components	Fit; retrain	Dynamics; policy	MetaPusher [368]; REVOLVE [345]
Acquisition process	Refine	Teaching; training procedure	LMPC [117]; Eureka [115]; DrEureka [116]

Table 19 separates where experience is stored from how it changes subsequent learning. RopeFormer and Zeva-Ego use retained action–response history to adapt a fixed policy to encountered conditions [363, 364]. Programs retain executable solutions, and neural updates change predictive or motor competence. Any of these can contribute to physical recursive self-improvement when experience improves acquisition or verification on subsequent tasks [326]. The decisive comparison is the later learning process, not the storage format of the change.

Program improvement can supply data for subsequent model improvement. HarnessPAI uses validated program rollouts to fine-tune the action model; its published sequence ends at policy refinement, before a further round of harness evolution around the improved checkpoint [398]. TraceFlow instead grows an external trace bank with weights fixed, and its reported stacking gains peak before the final round [392]. These examples motivate measuring the successor’s acquisition rate on new tasks and the persistence of gains across rounds, beyond demonstrating a stronger current solver.

Learning to become more teachable. LMPC links accumulated teaching to improved subsequent acquisition [117]. Within a session, a fixed language model converts instructions and corrections into reward code for model-predictive control. Across sessions, successful interactions train it to predict teaching exchanges; the rollout variant searches those exchanges for a shorter route to success. Test-task success rises from 39.4% to 66.3%, while mean turns in successful sessions fall from 2.4 to 1.9. This connects experience to a successor’s response to new teaching.

For action policies and VLM agents, the corresponding comparison holds new tasks, motor interfaces, and interaction budgets fixed. Evaluating each successor with and without informative context separates improved use of teaching from stronger prior execution. Familiar tasks and fully specified targets measure execution competence; withheld requirements and matched teaching measure acquisition. This extends the controls in Section 7.2 across successive versions of a learner.

Improving execution, supervision, and acquisition. Retained guidance and context-building code can improve selection or interpretation of future evidence [47, 49]. Autonomous practice can improve resetting, execution, or a successor policy [114, 141]. RoboRSI’s successive skill revisions exemplify repertoire growth; its release-level cumulative coverage records tasks solved at least once [348]. REVOLVE updates two distinct components: failure and correction records retrain the manipulation policy, while mismatch memory refines the supervising agent’s judgments [345]. Selective resets and successor comparisons attribute gains to the retained solution, control, or supervision.

Predictive components offer further update locations. Q-Planning trains a value function on deployment experience while retaining the action policy [208]; SILVR trains a visual planner on its executions [327]; joint refinement changes both policies and the models that assess or train them [328–330]. These loops improve the control and training signals available to contextual learning. Their contribution to teachability is measured by how the resulting successor responds to new evidence, alongside its execution success.

Changing the acquisition procedure makes the recursive objective explicit. Eureka revises reward code from policy-training feedback; DrEureka also configures domain randomization for transfer [115, 116]. ENPIRE extends agent-directed experimentation to physical environments, revising programs, data use, and learning algorithms through established interfaces [89]. Retaining effective learning choices together with their applicable conditions offers a route to improving future acquisition. Successor comparisons measure this benefit alongside repertoire growth. Evidence from simulated practice and physical experiments characterizes the operating conditions of this improvement.

Making the next experience informative. An improving learner can also change what it asks, observes, or tests. EnvGen, OMNI-EPIC, and SimWorld Studio use learner performance to shape subsequent simulated experience [261, 331, 332]. For ICL, teaching–execution pairs near interpretation failures could target conflicting examples, unfamiliar compositions, or corrections with a limited scope. Physical trials then test whether the inferred requirement survives contact and embodiment changes.

A failed insertion makes the distinction concrete. An ambiguous target calls for clarification; an uncertain contact response calls for a diagnostic motion. Retaining when each intervention resolved uncertainty could reduce unnecessary demonstrations and trials on later assemblies. The improvement loop then links requirement inference, evidence acquisition, physical verification, and reuse of an acquisition strategy.

Collective knowledge evolution. S6 extends the learning cycle to exchanges between robots (Figure 1). An arm can demonstrate a placement relation, a humanoid realize it under different kinematic constraints, and the receiver return a verified execution or correction. The exchanged lesson preserves task order and object relations while allowing grasps and balance control to change.

Shared progress and skill representations preserve task content across appearances [29, 30, 79]; multi-robot training supplies compatible competence [57, 132]; morphology interpolation and kinematic constraints support physical realization [131, 133]. Collective knowledge evolution adds a reciprocal loop: each execution helps revise knowledge another robot can use. Measuring transfer in both directions, with the teaching and practice required by each body, tests whether the exchange improves the group’s subsequent learning.

Across individual and collective learning cycles, requirement adherence, correction efficiency, and transfer measure the benefit to subsequent acquisition. Interaction, intervention, and training costs measure the investment required to obtain it.

Takeaways

- ① Diversity in what examples teach, motion coverage, and model scale address different limits. Learning unfamiliar combinations and relation families requires more than recognizing new instances.
- ② Composition, selective correction, and information acquisition expand what an episode can teach over an available motor repertoire.
- ③ Retained corrections should change later decisions through context and remain applicable when layouts change. Across longer learning cycles, the central outcome is a successor that learns more effectively from new teaching.

9 Conclusion

ICL for robots addresses a gap between possessing reusable competence and knowing how to apply it in the present task. Teaching and interaction resolve that gap with deployed neural parameters fixed. Four interfaces carry this information into action: contextual action inference, geometric reference transfer, prediction of consequences, and skill or program execution. Their benefits depend on compatible representations, applicable interaction geometry, realizable predictions, and sufficient executor capabilities, respectively. Correspondence and memory connect the evidence across situations and time; grounding establishes whether it remains usable.

The central synthesis is that contextual learning depends on preserving task-relevant distinctions from evidence to execution. Loss during selection, interpretation, realization, or reuse calls for different data, representations, or control capabilities. Across manipulation and navigation, the corresponding tests ask whether changed teaching appropriately redirects behavior, its requirement survives physical change, and retained experience helps later attempts. Acquisition costs and motor competence determine the value and scope of these gains.

The next challenge is to make these dependencies transferable: infer unfamiliar combinations of requirements, revise only the constraints affected by feedback, and discard experience whose conditions no longer hold. Physical recursive self-improvement extends this agenda to a successor that acquires new tasks more effectively; collective knowledge evolution asks whether verified exchanges improve learning across embodiments. Both are assessed by what later learners acquire and the total teaching, interaction, and training effort required.

Image attribution. Figures combine original vectors, AI-assisted illustrations, and cited task imagery. AI-assisted scenes appear in Figures 1, 10, 20, 23, 25, and 26; the S6 robot illustrations in Figures 1 and 3 are original vector drawings. Source motion and task annotations are preserved.

Method imagery comes from Vid2Robot, SemAnCorr, Zero-WAM, and UniSkill [14, 24, 30, 38]. Additional examples come from ICRT, MemoryVLA, WAM-TTT, DROC, MT3, and CorrectVLA [11, 13, 36, 42, 46, 142]; research examples use VIMA, Instant Policy, and HoloAssist [16, 73, 100]. Acquisition examples use DROID, UMI, Ego2Robot, and MimicGen [58, 62, 74, 222]. Other recordings come from Galaxea Open-World [61], LIBERO-PRO [106], GPT-Policy [94], and HumanAndRobot.

Vid2Robot imagery is adapted under CC BY-SA 4.0; SemAnCorr, Zero-WAM, UniSkill, LIBERO-PRO, VIMA, Instant Policy, HoloAssist, and DROID imagery under CC BY 4.0; and Galaxea imagery under CC BY-NC-SA 4.0. Editable sources include image provenance and institution-mark sources.

References

- [1] Y. Duan, M. Andrychowicz, B. C. Stadie, J. Ho, J. Schneider, I. Sutskever, et al. One-Shot Imitation Learning. *arXiv preprint arXiv:1703.07326*, 2017. URL <https://arxiv.org/abs/1703.07326>.
- [2] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel. RL²: Fast Reinforcement Learning via Slow Reinforcement Learning. *arXiv preprint arXiv:1611.02779*, 2016. URL <https://arxiv.org/abs/1611.02779>.
- [3] N. Mishra, M. Rohaninejad, X. Chen, and P. Abbeel. A Simple Neural Attentive Meta-Learner. *arXiv preprint arXiv:1707.03141*, 2017. URL <https://arxiv.org/abs/1707.03141>.
- [4] K. Rakelly, A. Zhou, D. Quillen, C. Finn, and S. Levine. Efficient Off-Policy Meta-Reinforcement Learning via Probabilistic Context Variables. *arXiv preprint arXiv:1903.08254*, 2019. URL <https://arxiv.org/abs/1903.08254>.
- [5] M. Xu, Y. Shen, S. Zhang, Y. Lu, D. Zhao, J. B. Tenenbaum, et al. Prompting Decision Transformer for Few-Shot Policy Generalization. *arXiv preprint arXiv:2206.13499*, 2022. URL <https://arxiv.org/abs/2206.13499>.
- [6] M. Laskin, L. Wang, J. Oh, E. Parisotto, S. Spencer, R. Steigerwald, et al. In-context Reinforcement Learning with Algorithm Distillation. *arXiv preprint arXiv:2210.14215*, 2022. URL <https://arxiv.org/abs/2210.14215>.
- [7] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine. One-Shot Visual Imitation Learning via Meta-Learning. *arXiv preprint arXiv:1709.04905*, 2017. URL <https://arxiv.org/abs/1709.04905>.
- [8] T. Yu, C. Finn, A. Xie, S. Dasari, T. Zhang, P. Abbeel, et al. One-Shot Imitation from Observing Humans via Domain-Adaptive Meta-Learning. *arXiv preprint arXiv:1802.01557*, 2018. URL <https://arxiv.org/abs/1802.01557>.

- [9] M. Van der Merwe and D. K. Jha. In-Context Iterative Policy Improvement for Dynamic Manipulation. In *Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pages 3888–3901, 2025. URL <https://proceedings.mlr.press/v305/merwe25a.html>.
- [10] Y. Jiang, Y. Chebotar, R. Zheng, F. Hu, Y. Ge, J. Wu, et al. RoboTTT: Context Scaling for Robot Policies. *arXiv preprint arXiv:2607.15275*, 2026. URL <https://arxiv.org/abs/2607.15275>.
- [11] Y. Feng, B. Han, J. Lyu, K. Liu, Y. Zheng, Y. Wan, W. Liu, S. Han, R. Li, Y. Zhang, F. Liu, X. Shi, L. Liu, Y. Wang, Z. Zhang, and H. Wang. WAM-TTT: Steering World-Action Models by Watching Human Play at Test Time, 2026. URL <https://arxiv.org/abs/2607.06988>.
- [12] Z. Mandi, F. Liu, K. Lee, and P. Abbeel. Towards More Generalizable One-shot Visual Imitation Learning. *arXiv preprint arXiv:2110.13423*, 2021. URL <https://arxiv.org/abs/2110.13423>.
- [13] L. Fu, H. Huang, G. Datta, L. Y. Chen, W. C.-H. Panitch, F. Liu, et al. In-Context Imitation Learning via Next-Token Prediction. *arXiv preprint arXiv:2408.15980*, 2024. URL <https://arxiv.org/abs/2408.15980>.
- [14] V. Jain, M. Attarian, N. J. Joshi, A. Wahid, D. Driess, Q. Vuong, et al. Vid2Robot: End-to-end Video-conditioned Policy Learning with Cross-Attention Transformers. *arXiv preprint arXiv:2403.12943*, 2024. URL <https://arxiv.org/abs/2403.12943>.
- [15] N. Di Palo and E. Johns. Keypoint Action Tokens Enable In-Context Imitation Learning in Robotics. *arXiv preprint arXiv:2403.19578*, 2024. URL <https://arxiv.org/abs/2403.19578>.
- [16] V. Vosylius and E. Johns. Instant Policy: In-Context Imitation Learning via Graph Diffusion. *arXiv preprint arXiv:2411.12633*, 2024. URL <https://arxiv.org/abs/2411.12633>.
- [17] K. Sridhar, S. Dutta, D. Jayaraman, and I. Lee. RICL: Adding In-Context Adaptability to Pre-Trained Vision-Language-Action Models. *arXiv preprint arXiv:2508.02062*, 2025. URL <https://arxiv.org/abs/2508.02062>.
- [18] S. Yang, Z. Liu, X. Li, R. Xiao, K. Zhu, and M. Li. ICI-VLA: In-Context Imitation with Spatiotemporally Aligned Demonstrations for Vision-Language-Action Models, 2026. URL <https://arxiv.org/abs/2609.07581>. Version 1, accessed 11 September 2026.
- [19] Y. Mu, Y. Zhuang, F. Ni, B. Wang, J. Chen, J. Hao, and P. Luo. Decomposed mutual information optimization for generalized context in meta-reinforcement learning. In *Advances in Neural Information Processing Systems*, 2022. URL <https://arxiv.org/abs/2210.04209>.
- [20] F. Ni, J. Hao, Y. Mu, Y. Yuan, Y. Zheng, B. Wang, and Z. Liang. Metadiffuser: Diffusion model as conditional planner for offline meta-rl. In *International Conference on Machine Learning*, 2023. URL <https://proceedings.mlr.press/v202/ni23a.html>.
- [21] R. G. Goswami, P. Krishnamurthy, Y. LeCun, and F. Khorrami. OSVI-WM: One-Shot Visual Imitation for Unseen Tasks using World-Model-Guided Trajectory Generation. *arXiv preprint arXiv:2505.20425*, 2025. URL <https://arxiv.org/abs/2505.20425>.
- [22] J. He, G. Li, J. Zhang, C. Hou, Z. Che, and S. Zhang. Demo-JEPA: Joint-Embedding Predictive Architecture for One-shot Cross-Embodiment Imitation, 2026. URL <https://arxiv.org/abs/2605.20811>.
- [23] G. Chen, M. Wang, T. Cui, Z. Zhou, Q. Shao, S. Li, et al. HOST: Robots Acquire Manipulation Skills in Seconds from a Single Human Video. *arXiv preprint arXiv:2607.20033*, 2026. URL <https://arxiv.org/abs/2607.20033>.
- [24] J. Zhou, Q. Zhang, G. Xu, C. Fan, Y. Zhao, R. Wang, et al. Zero-WAM: In-Context World-Action Modeling from Human Videos for Open-Ended Task Generalization. *arXiv preprint arXiv:2608.26103*, 2026. URL <https://arxiv.org/abs/2608.26103>.
- [25] H. Su, Z. Liu, X. Jin, H. Dou, C. Hu, B. Li, Z. Liu, R. Xu, J. Fang, X. Zhang, Z. Yang, X. Yang, C. Gao, J. Yan, Y. Li, and W. Wu. WorldScape Policy 2.0: Empowering Steerable World Action Modeling with Reasoning-Augmented Memory, 2026. URL <https://arxiv.org/abs/2607.18840>.
- [26] D. Xu, S. Nair, Y. Zhu, J. Gao, A. Garg, L. Fei-Fei, and S. Savarese. Neural Task Programming: Learning to Generalize Across Hierarchical Tasks. *arXiv preprint arXiv:1710.01813*, 2017. URL <https://arxiv.org/abs/1710.01813>.
- [27] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, et al. Code as Policies: Language Model Programs for Embodied Control. *arXiv preprint arXiv:2209.07753*, 2022. URL <https://arxiv.org/abs/2209.07753>.
- [28] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. *arXiv preprint arXiv:2209.11302*, 2022. URL <https://arxiv.org/abs/2209.11302>.
- [29] M. Xu, Z. Xu, C. Chi, M. Veloso, and S. Song. XSkill: Cross Embodiment Skill Discovery. In *Conference on Robot Learning*, 2023. URL <https://proceedings.mlr.press/v229/xu23a.html>.
- [30] H. Kim, J. Kang, H. Kang, M. Cho, S. J. Kim, and Y. Lee. UniSkill: Imitating Human Videos via Cross-Embodiment Skill Representations, 2025. URL <https://arxiv.org/abs/2505.08787>.
- [31] Y. Chen, Z. Bai, Z. Cao, W. Zeng, K. Q. Lin, Y. Lin, G. Liang, K. Y. Ma, Q. Huang, and M. Z. Shou. Show-Harness: Just a VLM Agent Can Play Robots, 2026. URL <https://arxiv.org/abs/2609.10522v1>.

- [32] M. Argus, L. Hermann, J. Long, and T. Brox. FlowControl: Optical Flow Based Visual Servoing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020. URL <https://arxiv.org/abs/2007.00291>.
- [33] E. Valassakis, G. Papagiannis, N. Di Palo, and E. Johns. Demonstrate Once, Imitate Immediately (DOME): Learning Visual Servoing for One-Shot Imitation Learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022. URL <https://arxiv.org/abs/2204.02863>.
- [34] N. Di Palo and E. Johns. DINOBot: Robot Manipulation via Retrieval and Alignment with Vision Foundation Models. In *IEEE International Conference on Robotics and Automation*, 2024. URL <https://arxiv.org/abs/2402.13181>.
- [35] N. Heppert, M. Argus, T. Welschehold, T. Brox, and A. Valada. DITTO: Demonstration Imitation by Trajectory Transformation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024. URL <https://ditto.cs.uni-freiburg.de/>.
- [36] K. Dreczkowski, P. Vitiello, V. Vosylius, and E. Johns. Learning a Thousand Tasks in a Day. *Science Robotics*, 10:eadv7594, 2025. doi: 10.1126/scirobotics.adv7594. URL <https://arxiv.org/abs/2511.10110>.
- [37] Y. Wang and E. Johns. One-Shot Dual-Arm Imitation Learning. In *IEEE International Conference on Robotics and Automation*, 2025. URL <https://arxiv.org/abs/2503.06831>.
- [38] X. Dong, W. Baron, H. Chen, U. Yoo, J. Ichnowski, and W. Zhi. SemAnCorr: Semantic Anchored Correspondence for Zero-Shot Manipulation Skill Transfer, 2026. URL <https://arxiv.org/abs/2607.28382>.
- [39] A. Kumar, Z. Fu, D. Pathak, and J. Malik. RMA: Rapid Motor Adaptation for Legged Robots. *arXiv preprint arXiv:2107.04034*, 2021. URL <https://arxiv.org/abs/2107.04034>.
- [40] M. Liu, D. Pathak, and A. Agarwal. LocoFormer: Generalist Locomotion via Long-context Adaptation. *arXiv preprint arXiv:2509.23745*, 2025. URL <https://arxiv.org/abs/2509.23745>.
- [41] S. Wang, J. Shi, S. Fei, Z. Fu, L. Ji, J. Gong, and X. Qiu. In-Context World Modeling for Robotic Control. *arXiv preprint arXiv:2606.26025*, 2026. URL <https://arxiv.org/abs/2606.26025>.
- [42] H. Shi, B. Xie, Y. Liu, L. Sun, F. Liu, T. Wang, E. Zhou, H. Fan, X. Zhang, and G. Huang. MemoryVLA: Perceptual-Cognitive Memory in Vision-Language-Action Models for Robotic Manipulation, 2025. URL <https://arxiv.org/abs/2508.19236>.
- [43] A. Sridhar, J. Pan, S. Sharma, and C. Finn. MemER: Scaling Up Memory for Robot Control via Experience Retrieval. *arXiv preprint arXiv:2510.20328*, 2025. URL <https://arxiv.org/abs/2510.20328>.
- [44] M. Torne, K. Pertsch, H. Walke, K. Vedder, S. Nair, B. Ichter, A. Z. Ren, H. Wang, J. Tang, K. Stachowicz, K. Dhabalia, M. Equi, Q. Vuong, J. T. Springenberg, S. Levine, C. Finn, and D. Driess. MEM: Multi-Scale Embodied Memory for Vision Language Action Models. *arXiv preprint arXiv:2603.03596*, 2026. URL <https://arxiv.org/abs/2603.03596>.
- [45] C. Yin, W. Xu, J. Yang, S. Tam, H. Liu, Y. Yao, X. Zeng, J. Cui, Y. Wang, Z. Yin, and Y. Lin. SimpleMemVLA: A Simple but Effective Native-Video Memory for Vision-Language-Action Models, 2026. URL <https://arxiv.org/abs/2609.05533>. Version 1, accessed 11 September 2026.
- [46] L. Zha, Y. Cui, L.-H. Lin, M. Kwon, M. G. Arenas, A. Zeng, F. Xia, and D. Sadigh. Distilling and Retrieving Generalizable Knowledge for Robot Manipulation via Language Corrections. *arXiv preprint arXiv:2311.10678*, 2023. URL <https://arxiv.org/abs/2311.10678>.
- [47] R. Lu, Y. Wu, E. Kou, L. Fu, W. Xiao, A. Mandlekar, Y. Xu, G. Shi, K. Goldberg, A. Chen, M. Chowdhury, Y. Zhu, L. J. Fan, and G. Wang. ASPIRE: Agentic /Skills Discovery for Robotics, 2026. URL <https://arxiv.org/abs/2607.00272>.
- [48] D. Chen, A. R. Jain, E. Goossen, E. Romig, M. Murray, N. Walker, and M. Cakmak. A Few Words Go a Long Way: Language Guided Robot Policy Synthesis, 2026. URL <https://arxiv.org/abs/2607.23784>.
- [49] P. Wang, Z. Ma, Y. Chang, X. Luo, Y. Zhang, Z. Wang, X. Yang, S. Feng, Y. Yang, and D. Li. Self-Evolving Embodied Agents via Skill-Harness Evolution, 2026. URL <https://arxiv.org/abs/2608.11350>.
- [50] G. Tzifas and H. Kasaei. Lifelong Robot Library Learning: Bootstrapping Composible and Generalizable Skills for Embodied Control with Language Models, 2024. URL <https://arxiv.org/abs/2406.18746>.
- [51] Y. Meng, Z. Sun, M. Fest, X. Li, Z. Bing, and A. Knoll. Growing with Your Embodied Agent: A Human-in-the-Loop Lifelong Code Generation Framework for Long-Horizon Manipulation Skills. *arXiv preprint arXiv:2509.18597*, 2025. URL <https://arxiv.org/abs/2509.18597>.
- [52] J. Zhang, J. Ge, H. Yoo, L. Fu, Z. Yang, Y. Liu, et al. Playful Agentic Robot Learning. *arXiv preprint arXiv:2606.19419*, 2026. URL <https://arxiv.org/abs/2606.19419>.
- [53] C. Nie, Z. Liu, and H. Wang. Teach and Grow: An Agent-Centered Architecture for General Robot Learning. *arXiv preprint arXiv:2608.17209*, 2026. URL <https://arxiv.org/abs/2608.17209>.
- [54] A. Patel, B. Pekarek, J. E. C. Hernandez, and S. Song. Behavior Prompting Policy: Demonstrations as Prompts for Manipulation. *arXiv preprint arXiv:2606.30457*, 2026. URL <https://arxiv.org/abs/2606.30457>.
- [55] J. Ding, X. Dai, R. Herzig, N. Hroub, J. Mai, D. Dai, B. Ghanem, and M. Elhoseiny. ContextFlow: In-Context Flow Matching for Robot Manipulation, 2026. URL <https://arxiv.org/abs/2609.06852v1>.
- [56] H. Walke, K. Black, A. Lee, M. J. Kim, M. Du, C. Zheng, T. Zhao, P. Hansen-Estruch, Q. Vuong, A. He, V. Myers, K. Fang, C. Finn, and S. Levine. BridgeData V2: A Dataset for Robot Learning at Scale. *arXiv*

- preprint arXiv:2308.12952*, 2023. URL <https://arxiv.org/abs/2308.12952>.
- [57] Open X-Embodiment Collaboration. Open X-Embodiment: Robotic Learning Datasets and RT-X Models. *arXiv preprint arXiv:2310.08864*, 2023. URL <https://arxiv.org/abs/2310.08864>.
- [58] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, et al. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. *arXiv preprint arXiv:2403.12945*, 2024. URL <https://arxiv.org/abs/2403.12945>.
- [59] H.-S. Fang, H. Fang, Z. Tang, J. Liu, C. Wang, J. Wang, et al. RH20T: A Comprehensive Robotic Dataset for Learning Diverse Skills in One-Shot. *arXiv preprint arXiv:2307.00595*, 2023. URL <https://arxiv.org/abs/2307.00595>.
- [60] AgiBot-World-Contributors. AgiBot World Colosseo: A Large-scale Manipulation Platform for Scalable and Intelligent Embodied Systems. *arXiv preprint arXiv:2503.06669*, 2025. URL <https://arxiv.org/abs/2503.06669>.
- [61] T. Jiang, T. Yuan, Y. Liu, C. Lu, J. Cui, X. Liu, S. Cheng, J. Gao, H. Xu, and H. Zhao. Galaxea Open-World Dataset and G0 Dual-System VLA Model, 2025. URL <https://arxiv.org/abs/2509.00576>.
- [62] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, et al. Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots. *arXiv preprint arXiv:2402.10329*, 2024. URL <https://arxiv.org/abs/2402.10329>.
- [63] Y. Liu, Y. Liu, C. Jiang, K. Lyu, W. Wan, H. Shen, B. Liang, Z. Fu, H. Wang, and L. Yi. HOI4D: A 4D Egocentric Dataset for Category-Level Human-Object Interaction, 2022. URL <https://arxiv.org/abs/2203.01577>.
- [64] P. Banerjee, S. Shkodrani, P. Moulon, S. Hampali, S. Han, F. Zhang, L. Zhang, J. Fountain, E. Miller, S. Basol, R. Newcombe, R. Wang, J. J. Engel, and T. Hodan. HOT3D: Hand and Object Tracking in 3D from Egocentric Multi-View Videos, 2024. URL <https://arxiv.org/abs/2411.19167>.
- [65] R. Hoque, P. Huang, D. J. Yoon, M. Sivapurapu, and J. Zhang. EgoDex: Learning Dexterous Manipulation from Large-Scale Egocentric Video, 2025. URL <https://arxiv.org/abs/2505.11709>.
- [66] T. Ohkawa, J. Arima, Y. Noguchi, M. Tateno, M. Sugiura, T. Okubo, K. Ikeuchi, Y. Shin, H. Nishizawa, N. Kanazawa, Y. Wakayama, D. Fukunaga, K. Makihara, T. Motoda, F. Erich, Y. Domae, T. Matsushima, Y. Okumatsu, and K. Ota. YUBI: Yielding Universal Bidigital Interface for Bimanual Dexterous Manipulation at Scale, 2026. URL <https://arxiv.org/abs/2606.10244>.
- [67] R. Zheng, D. Niu, Y. Xie, J. Wang, M. Xu, Y. Jiang, F. Castañeda, F. Hu, Y. L. Tan, L. Fu, T. Darrell, F. Huang, Y. Zhu, D. Xu, and L. Fan. EgoScale: Scaling Dexterous Manipulation with Diverse Egocentric Human Data, 2026. URL <https://arxiv.org/abs/2602.16710>.
- [68] K. Grauman, A. Westbury, E. Byrne, Z. Chavis, A. Furnari, R. Girdhar, J. Hamburger, H. Jiang, M. Liu, X. Liu, M. Martin, T. Nagarajan, I. Radosavovic, S. K. Ramakrishnan, F. Ryan, J. Sharma, M. Wray, M. Xu, E. Z. Xu, C. Zhao, S. Bansal, D. Batra, V. Cartillier, S. Crane, T. Do, M. Doulaty, A. Erapalli, C. Feichtenhofer, A. Fragomeni, Q. Fu, A. Gebreselasie, C. Gonzalez, J. Hillis, X. Huang, Y. Huang, W. Jia, W. Khoo, J. Kolar, S. Kottur, A. Kumar, F. Landini, C. Li, Y. Li, Z. Li, K. Mangalam, R. Modhugu, J. Munro, T. Murrell, T. Nishiyasu, W. Price, P. R. Puentes, M. Ramazanova, L. Sari, K. Somasundaram, A. Southerland, Y. Sugano, R. Tao, M. Vo, Y. Wang, X. Wu, T. Yagi, Z. Zhao, Y. Zhu, P. Arbelaez, D. Crandall, D. Damen, G. M. Farinella, C. Fuegen, B. Ghanem, V. K. Ithapu, C. V. Jawahar, H. Joo, K. Kitani, H. Li, R. Newcombe, A. Oliva, H. S. Park, J. M. Rehg, Y. Sato, J. Shi, M. Z. Shou, A. Torralba, L. Torresani, M. Yan, and J. Malik. Ego4D: Around the World in 3,000 Hours of Egocentric Video. *arXiv preprint arXiv:2110.07058*, 2021. URL <https://arxiv.org/abs/2110.07058>.
- [69] D. Damen, H. Doughty, G. M. Farinella, A. Furnari, E. Kazakos, J. Ma, D. Moltisanti, J. Munro, T. Perrett, W. Price, and M. Wray. Rescaling Egocentric Vision. *arXiv preprint arXiv:2006.13256*, 2020. URL <https://arxiv.org/abs/2006.13256>.
- [70] K. Grauman, A. Westbury, L. Torresani, K. Kitani, J. Malik, T. Afouras, K. Ashutosh, V. Baiyya, S. Bansal, B. Boote, E. Byrne, Z. Chavis, J. Chen, F. Cheng, F.-J. Chu, S. Crane, A. Dasgupta, J. Dong, M. Escobar, C. Forigua, A. Gebreselasie, S. Haresh, J. Huang, M. M. Islam, S. Jain, R. Khirodkar, D. Kukreja, K. J. Liang, J.-W. Liu, S. Majumder, Y. Mao, M. Martin, E. Mavroudi, T. Nagarajan, F. Ragusa, S. K. Ramakrishnan, L. Seminara, A. Somayazulu, Y. Song, S. Su, Z. Xue, E. Zhang, J. Zhang, A. Castillo, C. Chen, X. Fu, R. Furuta, C. Gonzalez, P. Gupta, J. Hu, Y. Huang, W. Khoo, A. Kumar, R. Kuo, S. Lakhavani, M. Liu, M. Luo, Z. Luo, B. Meredith, A. Miller, O. Oguntola, X. Pan, P. Peng, S. Pramanick, M. Ramazanova, F. Ryan, W. Shan, K. Somasundaram, C. Song, A. Southerland, M. Tateno, H. Wang, Y. Wang, T. Yagi, M. Yan, X. Yang, Z. Yu, S. C. Zha, C. Zhao, Z. Zhao, Z. Zhu, J. Zhuo, P. Arbelaez, G. Bertasius, D. Crandall, D. Damen, J. Engel, G. M. Farinella, A. Furnari, B. Ghanem, J. Hoffman, C. V. Jawahar, R. Newcombe, H. S. Park, J. M. Rehg, Y. Sato, M. Savva, J. Shi, M. Z. Shou, and M. Wray. Ego-Exo4D: Understanding Skilled Human Activity from First- and Third-Person Perspectives. *arXiv preprint arXiv:2311.18259*, 2023. URL <https://arxiv.org/abs/2311.18259>.
- [71] R. Goyal, S. E. Kahou, V. Michalski, J. Materzyńska, S. Westphal, H. Kim, V. Haenel, I. Fruend, P. Yianilos, M. Mueller-Freitag, F. Hoppe, C. Thureau, I. Bax, and R. Memisevic. The "something something" video

- database for learning and evaluating visual common sense. *arXiv preprint arXiv:1706.04261*, 2017. URL <https://arxiv.org/abs/1706.04261>.
- [72] F. Sener, D. Chatterjee, D. Shelepov, K. He, D. Singhania, R. Wang, and A. Yao. Assembly101: A Large-Scale Multi-View Video Dataset for Understanding Procedural Activities, 2022. URL <https://arxiv.org/abs/2203.14712>.
 - [73] X. Wang, T. Kwon, M. Rad, B. Pan, I. Chakraborty, S. Andrist, D. Bohus, A. Feniello, B. Tekin, F. V. Frujeri, N. Joshi, and M. Pollefeys. HoloAssist: an Egocentric Human Interaction Dataset for Interactive AI Assistants in the Real World, 2023. URL <https://arxiv.org/abs/2309.17024>.
 - [74] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox. MimicGen: A Data Generation System for Scalable Robot Learning using Human Demonstrations. *arXiv preprint arXiv:2310.17596*, 2023. URL <https://arxiv.org/abs/2310.17596>.
 - [75] S. Nasiriany, A. Maddukuri, L. Zhang, A. Parikh, A. Lo, A. Joshi, A. Mandlekar, and Y. Zhu. RoboCasa: Large-Scale Simulation of Everyday Tasks for Generalist Robots, 2024. URL <https://arxiv.org/abs/2406.02523>.
 - [76] M. Ahn, D. Dwibedi, C. Finn, M. G. Arenas, K. Gopalakrishnan, K. Hausman, B. Ichter, A. Irpan, N. Joshi, R. Julian, S. Kirmani, I. Leal, E. Lee, S. Levine, Y. Lu, I. Leal, S. Maddineni, K. Rao, D. Sadigh, P. Sanketi, P. Sermanet, Q. Vuong, S. Welker, F. Xia, T. Xiao, P. Xu, S. Xu, and Z. Xu. AutoRT: Embodied Foundation Models for Large Scale Orchestration of Robotic Agents, 2024. URL <https://arxiv.org/abs/2401.12963>.
 - [77] C. Qian, R. Fan, Y. Ren, Y. Wang, and E. Johns. SynthICL: Scalable In-context Imitation Learning with Synthetic Data. *arXiv preprint arXiv:2606.08154*, 2026. URL <https://arxiv.org/abs/2606.08154>.
 - [78] P. Sermanet, C. Lynch, Y. Chebotar, J. Hsu, E. Jang, S. Schaal, and S. Levine. Time-Contrastive Networks: Self-Supervised Learning from Video. *arXiv preprint arXiv:1704.06888*, 2017. URL <https://arxiv.org/abs/1704.06888>.
 - [79] K. Zakka, A. Zeng, P. Florence, J. Tompson, J. Bohg, and D. Dwibedi. XIRL: Cross-embodiment Inverse Reinforcement Learning. *arXiv preprint arXiv:2106.03911*, 2021. URL <https://arxiv.org/abs/2106.03911>.
 - [80] K. Kedia, P. Dan, A. Chao, M. A. Pace, and S. Choudhury. One-Shot Imitation under Mismatched Execution. *arXiv preprint arXiv:2409.06615*, 2024. URL <https://arxiv.org/abs/2409.06615>.
 - [81] C. Agia, R. Sinha, J. Yang, Z.-a. Cao, R. Antonova, M. Pavone, and J. Bohg. Unpacking Failure Modes of Generative Policies: Runtime Monitoring of Consistency and Progress. *arXiv preprint arXiv:2410.04640*, 2024. URL <https://arxiv.org/abs/2410.04640>.
 - [82] C. Xu, T. K. Nguyen, E. Dixon, C. Rodriguez, P. Miller, R. Lee, P. Shah, R. Ambrus, H. Nishimura, and M. Itkina. Can We Detect Failures Without Failure Data? Uncertainty-Aware Runtime Failure Detection for Imitation Learning Policies. *arXiv preprint arXiv:2503.08558*, 2025. URL <https://arxiv.org/abs/2503.08558>.
 - [83] Y. J. Ma, J. Hejna, A. Wahid, C. Fu, D. Shah, J. Liang, Z. Xu, S. Kirmani, P. Xu, D. Driess, T. Xiao, J. Tompson, O. Bastani, D. Jayaraman, W. Yu, T. Zhang, D. Sadigh, and F. Xia. Vision Language Models are In-Context Value Learners, 2024. URL <https://arxiv.org/abs/2411.04549>.
 - [84] Z. Liu, A. Bahety, and S. Song. REFLECT: Summarizing Robot Experiences for Failure Explanation and Correction, 2023. URL <https://arxiv.org/abs/2306.15724>.
 - [85] J. Duan, W. Pumacay, N. Kumar, Y. R. Wang, S. Tian, W. Yuan, R. Krishna, D. Fox, A. Mandlekar, and Y. Guo. AHA: A Vision-Language-Model for Detecting and Reasoning Over Failures in Robotic Manipulation, 2024. URL <https://arxiv.org/abs/2410.00371>.
 - [86] A. Z. Ren, A. Dixit, A. Bodrova, S. Singh, S. Tu, N. Brown, P. Xu, L. Takayama, F. Xia, J. Varley, Z. Xu, D. Sadigh, A. Zeng, and A. Majumdar. Robots That Ask For Help: Uncertainty Alignment for Large Language Model Planners. *arXiv preprint arXiv:2307.01928*, 2023. URL <https://arxiv.org/abs/2307.01928>.
 - [87] Y. Dai, J. Lee, N. Fazeli, and J. Chai. RACER: Rich Language-Guided Failure Recovery Policies for Imitation Learning, 2024. URL <https://arxiv.org/abs/2409.14674>.
 - [88] F. Chen, X. Ding, B. Huang, X. Li, M. Wang, J. He, et al. Zeva: In-Context Causal Learning for Generalizable Embodied Manipulation. *arXiv preprint arXiv:2608.30880*, 2026. URL <https://arxiv.org/abs/2608.30880>.
 - [89] W. Xiao, J. Xie, T. Zhang, H. Lin, L. Fu, H. Xue, J. Lu, Y. Yang, C. Dai, Z. Wang, J. Wu, G. Wang, S. S. Sastry, K. Goldberg, L. Fan, Y. Zhu, and G. Shi. ENPIRE: Agentic Robot Policy Self-Improvement in the Real World, 2026. URL <https://arxiv.org/abs/2606.19980v1>.
 - [90] X. Ding, L. Mi, M. Huang, Z. Wang, C. Zhang, Z. Hao, et al. Zetta ζ : An Efficient Closed-Loop Embodied Harness for Self-Evolving Physical Intelligence. *arXiv preprint arXiv:2608.16590*, 2026. URL <https://arxiv.org/abs/2608.16590>.
 - [91] Y. Wang, C. Qian, and E. Johns. Instant-Fold: In-Context Imitation Learning for Deformable Object Manipulation. *arXiv preprint arXiv:2606.04269*, 2026. URL <https://arxiv.org/abs/2606.04269>.
 - [92] A. Palma, I. Spinelli, V. Prasad, L. Scofano, Y. Jin, G. Chalvatzaki, and F. Galasso. Bimanual Robot Manipulation via Multi-Agent In-Context Learning, 2026. URL <https://arxiv.org/abs/2604.20348>.
 - [93] Z. Chen, C. Gao, L. Shao, J. Shi, J. Huo, and Y. Gao. ManiLong-Shot: Interaction-Aware One-Shot Imitation Learning for Long-Horizon Manipulation. In *AAAI Conference on Artificial Intelligence*, 2026. URL <https://ojs.aaai.org/index.php/AAAI/article/view/38881>.

- [94] D. Cheng, T. Yi, Y. Fang, X. Zhang, F. Feng, Y. Li, G. Zhuang, R. Wang, S. Yang, W. Song, W. Xue, M. Wu, J. Gui, J. Wang, and T. Wu. In-context robot learning with vlm agents, 2026. URL <https://arxiv.org/abs/2609.19138>. GPT-Policy; first submitted 16 September 2026.
- [95] Z. Tang, S. Chen, S. Zhang, P. Yang, Y. Liu, W. Zhao, et al. LabDex: A Hierarchical Benchmark for Dexterous Manipulation in Laboratories. *arXiv preprint arXiv:2608.18618*, 2026. URL <https://arxiv.org/abs/2608.18618>.
- [96] Y. Tang, W. Yu, J. Tan, H. Zen, A. Faust, and T. Harada. SayTap: Language to Quadrupedal Locomotion. *arXiv preprint arXiv:2306.07580*, 2023. URL <https://arxiv.org/abs/2306.07580>.
- [97] K. Wu, H. Chen, J. Han, H. Liu, C. Wang, Y. Wang, Z. Li, S. Liu, and F. Zhong. AdaTracker: Learning Adaptive In-Context Policy for Cross-Embodiment Active Visual Tracking. *IEEE Robotics and Automation Letters*, 2026. doi: 10.1109/LRA.2026.3692075. URL <https://arxiv.org/abs/2604.20305>.
- [98] Generalist Team. GEN-1.5: Embodied Foundation Models are One-Shot Learners. Generalist AI Blog, August 19, 2026. URL <https://generalistai.com/blog/gen-1.5>. Accessed September 19, 2026. Industrial technical report.
- [99] Skild AI. Introducing S1: In-Context Learning for Robotics. Skild AI official blog, 2026. URL <https://www.skild.ai/blogs/s1>. Accessed September 19, 2026. Industrial technical report.
- [100] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, et al. VIMA: General Robot Manipulation with Multimodal Prompts. *arXiv preprint arXiv:2210.03094*, 2022. URL <https://arxiv.org/abs/2210.03094>.
- [101] X. Zhou, Y. Cai, F. Wang, R. Ju, H. Ren, R. Liu, et al. The Imitator Game: Benchmarking Robot Imitative Ability Beyond Action Prediction. *arXiv preprint arXiv:2608.22301*, 2026. URL <https://arxiv.org/abs/2608.22301>.
- [102] A. Ruoss, F. Pardo, H. Chan, B. Li, V. Mnih, and T. Genewein. LMAct: A benchmark for in-context imitation learning with long multimodal demonstrations. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 52303–52344, 2025. URL <https://proceedings.mlr.press/v267/ruoss25a.html>.
- [103] Y. Dai, H. Fu, J. Lee, Y. Liu, H. Zhang, J. Yang, C. Finn, N. Fazeli, and J. Chai. RoboMME: Benchmarking and understanding memory for robotic generalist policies. *arXiv preprint arXiv:2603.04639*, 2026. URL <https://arxiv.org/abs/2603.04639>.
- [104] H. Sun, H. Wang, J. Chen, W. Fang, Z. Su, J. Yi, W. Yi, H. Chen, and Z. Zhao. MEMOBench: A Process Level Memory Benchmark for Robotic Manipulation, 2026. URL <https://arxiv.org/abs/2609.07047>. Version 1, accessed 11 September 2026.
- [105] Y. Gao, J. J. Liu, S. Li, and S. Song. Gated Memory Policy: In-Context Memorization and Adaptation. *arXiv preprint arXiv:2604.18933*, 2026. URL <https://arxiv.org/abs/2604.18933>.
- [106] X. Zhou, Y. Xu, G. Tie, Y. Chen, G. Zhang, D. Chu, P. Zhou, and L. Sun. LIBERO-PRO: Towards robust and fair evaluation of vision-language-action models beyond memorization. *arXiv preprint arXiv:2510.03827*, 2025. URL <https://arxiv.org/abs/2510.03827>.
- [107] Y. Fang, Y. Feng, D. Jing, J. Liu, Y. Yang, Z. Wei, D. Szafir, and M. Ding. When Vision Overrides Language: Evaluating and Mitigating Counterfactual Failures in VLAs, 2026. URL <https://arxiv.org/abs/2602.17659>.
- [108] Z. Chen, A. Tian, L. Wang, B. Joffe, Y. C. Lin, Y. Chen, S. Karamcheti, and D. Xu. ReSteer: Quantifying and Refining the Steerability of Multitask Robot Policies, 2026. URL <https://arxiv.org/abs/2603.17300>.
- [109] B. Yu, Y. Zhang, H. Liu, S. Lian, Y. Wei, X. Lin, Z. Shen, C. Wu, R. Hu, B. Wang, C. Huang, and K. Chen. RoboSemanticBench: Diagnosing Semantic Grounding in Action Prediction for VLA Models, 2026. URL <https://arxiv.org/abs/2606.02277>.
- [110] C. Ma, L. Luo, X. Luo, X. Feng, H. Zhang, W. Liu, et al. Behavior-Skill: A Fine-Grained Benchmark for Evaluating Vision-Language-Action Policies in Long-Horizon Tasks. *arXiv preprint arXiv:2608.30536*, 2026. URL <https://arxiv.org/abs/2608.30536>.
- [111] S. C. Y. Chan, A. Santoro, A. K. Lampinen, J. X. Wang, A. Singh, P. H. Richemond, J. McClelland, and F. Hill. Data distributional properties drive emergent in-context learning in transformers. *arXiv preprint arXiv:2205.05055*, 2022. URL <https://arxiv.org/abs/2205.05055>.
- [112] J. Von Oswald, E. Niklasson, E. Randazzo, J. Sacramento, A. Mordvintsev, A. Zhmoginov, and M. Vladymyrov. Transformers Learn In-Context by Gradient Descent. In *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 35151–35174, 2023. URL <https://proceedings.mlr.press/v202/von-oswald23a.html>.
- [113] S. Min, M. Lewis, L. Zettlemoyer, and H. Hajishirzi. MetaICL: Learning to Learn In Context. In *Proceedings of NAACL-HLT*, 2022. URL <https://arxiv.org/abs/2110.15943v2>.
- [114] K. Bousmalis, G. Vezzani, D. Rao, C. Devin, A. X. Lee, M. Bauza, T. Davchev, Y. Zhou, A. Gupta, A. Raju, A. Laurens, C. Fantacci, V. Dalibard, M. Zambelli, M. Martins, R. Pevceviciute, M. Blokzijl, M. Denil, N. Batchelor, T. Lampe, E. Parisotto, K. Żołna, S. Reed, S. G. Colmenarejo, J. Scholz, A. Abdolmaleki, O. Groth, J.-B. Regli, O. Sushkov, T. Rothörl, J. E. Chen, Y. Aytar, D. Barker, J. Ortiz, M. Riedmiller, J. T. Springenberg, R. Hadsell, F. Nori, and N. Heess. RoboCat: A Self-Improving Generalist Agent for Robotic Manipulation, 2023. URL <https://arxiv.org/abs/2306.11706>.

- [115] Y. J. Ma, W. Liang, G. Wang, D.-A. Huang, O. Bastani, D. Jayaraman, Y. Zhu, L. Fan, and A. Anandkumar. Eureka: Human-Level Reward Design via Coding Large Language Models, 2023. URL <https://arxiv.org/abs/2310.12931>.
- [116] Y. J. Ma, W. Liang, H.-J. Wang, S. Wang, Y. Zhu, L. Fan, O. Bastani, and D. Jayaraman. DrEureka: Language Model Guided Sim-To-Real Transfer, 2024. URL <https://arxiv.org/abs/2406.01967>.
- [117] J. Liang, F. Xia, W. Yu, A. Zeng, M. G. Arenas, M. Attarian, M. Bauza, M. Bennice, A. Bewley, A. Dostmohamed, C. K. Fu, N. Gileadi, M. Giustina, K. Gopalakrishnan, L. Hasenclever, J. Humprik, J. Hsu, N. Joshi, B. Jyenis, C. Kew, S. Kirmani, T.-W. E. Lee, K.-H. Lee, A. H. Michaely, J. Moore, K. Oslund, D. Rao, A. Ren, B. Tabanpour, Q. Vuong, A. Wahid, T. Xiao, Y. Xu, V. Zhuang, P. Xu, E. Frey, K. Caluwaerts, T. Zhang, B. Ichter, J. Tompson, L. Takayama, V. Vanhoucke, I. Shafran, M. Mataric, D. Sadigh, N. Heess, K. Rao, N. Stewart, J. Tan, and C. Parada. Learning to Learn Faster from Human Feedback with Language Model Predictive Control, 2024. URL <https://arxiv.org/abs/2402.11450>.
- [118] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. OpenVLA: An Open-Source Vision-Language-Action Model. *arXiv preprint arXiv:2406.09246*, 2024. URL <https://arxiv.org/abs/2406.09246>.
- [119] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, S. Jakubczak, T. Jones, L. Ke, S. Levine, A. Li-Bell, M. Mothukuri, S. Nair, K. Pertsch, L. X. Shi, J. Tanner, Q. Vuong, A. Walling, H. Wang, and U. Zhilinsky. π_0 : A Vision-Language-Action Flow Model for General Robot Control. *arXiv preprint arXiv:2410.24164*, 2024. URL <https://arxiv.org/abs/2410.24164>.
- [120] NVIDIA, J. Bjorck, F. Castañeda, N. Cherniadev, X. Da, R. Ding, L. J. Fan, Y. Fang, D. Fox, F. Hu, S. Huang, J. Jang, Z. Jiang, K. Kautz, K. Kundalia, L. Lao, Z. Li, Z. Lin, K. Lin, G. Liu, E. Llontop, L. Magne, A. Mandlekar, A. Narayan, S. Nasiriany, S. Reed, Y. L. Tan, G. Wang, Z. Wang, J. Wang, Q. Wang, J. Xiang, Y. Xie, Y. Xu, Z. Xu, S. Ye, Z. Yu, A. Zhang, H. Zhang, Y. Zhao, R. Zheng, and Y. Zhu. GR00T N1: An Open Foundation Model for Generalist Humanoid Robots. *arXiv preprint arXiv:2503.14734*, 2025. URL <https://arxiv.org/abs/2503.14734>.
- [121] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. URL <https://arxiv.org/abs/2005.14165>.
- [122] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A Survey of Robot Learning from Demonstration. *Robotics and Autonomous Systems*, 2009. doi: 10.1016/j.robot.2008.10.024. URL <https://publications.ri.cmu.edu/a-survey-of-robot-learning-from-demonstration>.
- [123] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard. Recent Advances in Robot Learning from Demonstration. *Annual Review of Control, Robotics, and Autonomous Systems*, 2020. doi: 10.1146/annurev-control-100819-063206. URL <https://www.annualreviews.org/content/journals/10.1146/annurev-control-100819-063206>.
- [124] Q. Dong, L. Li, D. Dai, C. Zheng, J. Ma, R. Li, H. Xia, J. Xu, Z. Wu, B. Chang, X. Sun, L. Li, and Z. Sui. A Survey on In-context Learning. In *Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, 2024. doi: 10.18653/v1/2024.emnlp-main.64. URL <https://aclanthology.org/2024.emnlp-main.64/>.
- [125] A. Moeini, J. Wang, J. Beck, E. Blaser, S. Whiteson, R. Chandra, et al. A Survey of In-Context Reinforcement Learning. *arXiv preprint arXiv:2502.07978*, 2025. URL <https://arxiv.org/abs/2502.07978>.
- [126] J. Ma, E. Zhang, H. Yang, D. Li, C. Xu, G. Wang, et al. Robot Learning from Human Videos: A Survey. *arXiv preprint arXiv:2604.27621*, 2026. URL <https://arxiv.org/abs/2604.27621>.
- [127] A. Run and Z. Ding. In-Context Reinforcement Learning under Non-Stationarity: A Survey, 2026. URL <https://arxiv.org/abs/2607.11906>.
- [128] B. Hou, G. Li, J. Jia, T. An, X. Guo, S. Leng, H. Geng, Y. Ze, T. Harada, P. Torr, O. Mees, M. Pollefeys, Z. Liu, J. Wu, P. Abbeel, J. Malik, Y. Du, and J. Yang. World Model for Robot Learning: A Comprehensive Survey. *arXiv preprint arXiv:2605.00080*, 2026. URL <https://arxiv.org/abs/2605.00080>.
- [129] Z. Wang, B. Wang, H. Zhang, T. Du, T. Chen, G. Sun, Y. He, Z. Shen, W. Ye, and A. Li. Vision-Language-Action in Robotics: A Survey of Datasets, Benchmarks, and Data Engines. *Transactions on Machine Learning Research*, 2026. URL <https://arxiv.org/abs/2604.23001>.
- [130] Y. Domae, K. Shirai, H. Oh, R. Nakajo, T. Motoda, K. Makihara, M. Murooka, T. Yagi, Y. Bando, and R. Hanai. The Embodiment Gap in Robot Foundation Models. *Transactions on Machine Learning Research*, 2026. URL <https://arxiv.org/abs/2608.18433>.
- [131] X. Liu, D. Pathak, and K. Kitani. REvolveR: Continuous Evolutionary Models for Robot-to-robot Policy Transfer. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 13995–14007, 2022. URL <https://proceedings.mlr.press/v162/liu22p.html>.

- [132] R. Doshi, H. Walke, O. Mees, S. Dasari, and S. Levine. Scaling Cross-Embodied Learning: One Policy for Manipulation, Navigation, Locomotion and Aviation. *arXiv preprint arXiv:2408.11812*, 2024. URL <https://arxiv.org/abs/2408.11812>.
- [133] S. Gupta, D. H. Salunkhe, and A. Billard. Demonstrate once, execute on many: Kinematic intelligence for cross-robot skill transfer. *Science Robotics*, 11(113):eaea1995, 2026. doi: 10.1126/scirobotics.aea1995. URL <https://doi.org/10.1126/scirobotics.aea1995>.
- [134] R. E. Fikes and N. J. Nilsson. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2(3–4):189–208, 1971. doi: 10.1016/0004-3702(71)90010-5. URL <https://www.sciencedirect.com/science/article/pii/0004370271900105>.
- [135] R. A. Brooks. A Robust Layered Control System for a Mobile Robot. Technical Report A.I. Memo 864, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, 1985. URL <https://people.csail.mit.edu/brooks/papers/AIM-864.pdf>.
- [136] O. Khatib. A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, 1987. doi: 10.1109/JRA.1987.1087068. URL https://khatib.stanford.edu/publications/pdfs/Khatib_1987_RA.pdf.
- [137] L. P. Kaelbling and T. Lozano-Pérez. Hierarchical Task and Motion Planning in the Now. In *IEEE International Conference on Robotics and Automation*, pages 1470–1477, 2011. doi: 10.1109/ICRA.2011.5980391. URL <https://people.csail.mit.edu/tlp/pdf/2011/hpnICRA11Final.pdf>.
- [138] A. Marzintotto, M. Colledanchise, C. Smith, and P. Ögren. Towards a Unified Behavior Trees Framework for Robot Control. In *IEEE International Conference on Robotics and Automation*, 2014. URL <https://www.csc.kth.se/~ccs/Publications/icra14b.html>.
- [139] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-End Training of Deep Visuomotor Policies. *arXiv preprint arXiv:1504.00702*, 2015. URL <https://arxiv.org/abs/1504.00702>.
- [140] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine. QT-Opt: Scalable Deep Reinforcement Learning for Vision-Based Robotic Manipulation. *arXiv preprint arXiv:1806.10293*, 2018. URL <https://arxiv.org/abs/1806.10293>.
- [141] A. Sharma, A. M. Ahmed, R. Ahmad, and C. Finn. Self-Improving Robots: End-to-End Autonomous Visuomotor Reinforcement Learning, 2023. URL <https://arxiv.org/abs/2303.01488>.
- [142] O. Kwon, P. Ortega-Kral, A. Buckner, and J. Oh. Training-Free Action Correction for VLA Model Failures via Language Feedback. *arXiv preprint arXiv:2608.29967*, 2026. URL <https://arxiv.org/abs/2608.29967>.
- [143] T. Lesort, V. Lomonaco, A. Stoian, D. Maltoni, D. Filliat, and N. Díaz-Rodríguez. Continual Learning for Robotics: Definition, Framework, Learning Strategies, Opportunities and Challenges, 2019. URL <https://arxiv.org/abs/1907.00182>.
- [144] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey. Meta-Learning in Neural Networks: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5149–5169, 2022. doi: 10.1109/TPAMI.2021.3079209. URL <https://arxiv.org/abs/2004.05439>.
- [145] Y. Hu, S. Liu, Y. Yue, G. Zhang, B. Liu, F. Zhu, J. Lin, H. Guo, S. Dou, Z. Xi, S. Jin, J. Tan, Y. Yin, J. Liu, Z. Zhang, Z. Sun, Y. Zhu, H. Sun, B. Peng, Z. Cheng, X. Fan, J. Guo, X. Yu, Z. Zhou, Z. Hu, J. Huo, J. Wang, Y. Niu, Y. Wang, Z. Yin, C. Hu, Y. Liao, Q. Li, K. Wang, W. Zhou, Y. Liu, D. Cheng, Q. Zhang, T. Gui, S. Pan, Y. Zhang, P. Torr, Z. Dou, J.-R. Wen, X. Huang, Y.-G. Jiang, and S. Yan. Memory in the Age of AI Agents. *arXiv preprint arXiv:2512.13564*, 2025. URL <https://arxiv.org/abs/2512.13564>.
- [146] W. Yu, J. Tan, C. K. Liu, and G. Turk. Preparing for the Unknown: Learning a Universal Policy with Online System Identification. *arXiv preprint arXiv:1702.02453*, 2017. URL <https://arxiv.org/abs/1702.02453>.
- [147] A. Zhou, E. Jang, D. Kappler, A. Herzog, M. Khansari, P. Wohlhart, Y. Bai, M. Kalakrishnan, S. Levine, and C. Finn. Watch, Try, Learn: Meta-Learning from Demonstrations and Reward. *arXiv preprint arXiv:1906.03352*, 2019. URL <https://arxiv.org/abs/1906.03352>.
- [148] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, et al. BC-Z: Zero-Shot Task Generalization with Robotic Imitation Learning. *arXiv preprint arXiv:2202.02005*, 2022. URL <https://arxiv.org/abs/2202.02005>.
- [149] J. Pari, N. M. Shafullah, S. P. Arunachalam, and L. Pinto. The Surprising Effectiveness of Representation Learning for Visual Imitation. In *Robotics: Science and Systems*, 2022. URL <https://www.roboticsproceedings.org/rss18/p010.pdf>.
- [150] S. Bahl, A. Gupta, and D. Pathak. Human-to-Robot Imitation in the Wild. *arXiv preprint arXiv:2207.09450*, 2022. URL <https://human2robot.github.io/resources/paper.pdf>.
- [151] H. Fei, Y. Zhou, J. Li, X. Li, Q. Xu, B. Li, S. Wu, et al. On Path to Multimodal Generalist: General-Level and General-Bench, 2025. URL <https://arxiv.org/abs/2505.04620>.
- [152] J. Park, J. Park, T. Kim, S. Choi, D. Han, and S. Yun. Retrieve, Don’t Retrain: Extending Vision Language Action Models to New Tasks at Test Time, 2026. URL <https://arxiv.org/abs/2606.15631>.
- [153] Y. Yin, Z. Wang, Y. Sharma, D. Niu, T. Darrell, and R. Herzig. In-Context Learning Enables Robot Action Prediction in LLMs. *arXiv preprint arXiv:2410.12782*, 2024. URL <https://arxiv.org/abs/2410.12782>.

- [154] K. Sridhar, S. Dutta, D. Jayaraman, and I. Lee. REGENT: A Retrieval-Augmented Generalist Agent That Can Act In-Context in New Environments. In *International Conference on Learning Representations*, 2025. URL <https://arxiv.org/abs/2412.04759>.
- [155] S. Park, H. Bharadhwaj, and S. Tulsiani. DemoDiffusion: One-Shot Human Imitation using pre-trained Diffusion Policy. *arXiv preprint arXiv:2506.20668*, 2025. URL <https://arxiv.org/abs/2506.20668>.
- [156] X. Zhang and A. Boularias. One-Shot Imitation Learning with Invariance Matching for Robotic Manipulation. In *Robotics: Science and Systems*, 2024. URL <https://www.roboticsproceedings.org/rss20/p134.html>.
- [157] S. James, M. Bloesch, and A. J. Davison. Task-Embedded Control Networks for Few-Shot Imitation Learning. *arXiv preprint arXiv:1810.03237*, 2018. URL <https://arxiv.org/abs/1810.03237>.
- [158] S. Dasari and A. Gupta. Transformers for One-Shot Visual Imitation. *arXiv preprint arXiv:2011.05970*, 2020. URL <https://arxiv.org/abs/2011.05970>.
- [159] G. Chen, M. Wang, Q. Shao, Z. Zhou, W. Mao, T. Cui, M. Zhu, Y. Deng, L. Yang, Z. Zhang, Y. Yang, H. Chen, and Y. Yue. See Once, Then Act: Vision-Language-Action Model with Task Learning from One-Shot Video Demonstrations, 2025. URL <https://arxiv.org/abs/2512.07582>.
- [160] R. Huang, C. Zeng, W. Tang, J. Cai, C. Lu, and P. Cai. Mimic Intent, Not Just Trajectories. In *Robotics: Science and Systems*, 2026. URL <https://www.roboticsproceedings.org/rss22/p206.html>.
- [161] Q. She, H. Yu, Z. Li, and P. Tan. MatchingPolicy: Correspondence-Aware Policy Enables Cross-Object In-Context Learning, 2026. URL <https://arxiv.org/abs/2608.16715>.
- [162] J. Yang, W. Huang, J. Zhang, M. Hu, and H. Guo. In-Context VLA: Endowing Vision-Language-Action Models with Language via In-Context Post-Training and Agentic Tool Use. *arXiv preprint arXiv:2608.05738*, 2026. URL <https://arxiv.org/abs/2608.05738>.
- [163] S. Xu, Y. Wang, Z. Wang, D. Zhu, C. Xia, C. Du, D. Liu, T. Huang, and C. Xu. StellaVLA: In-Context Structured Demonstration for Generalizable Vision-Language-Action Models. *arXiv preprint arXiv:2608.11671*, 2026. URL <https://arxiv.org/abs/2608.11671>.
- [164] M. Chang and S. Gupta. One-shot Visual Imitation via Attributed Waypoints and Demonstration Augmentation. *arXiv preprint arXiv:2302.04856*, 2023. URL <https://arxiv.org/abs/2302.04856>.
- [165] Y. Zhu, A. Lim, P. Stone, and Y. Zhu. Vision-based Manipulation from Single Human Video with Open-World Object Graphs. *arXiv preprint arXiv:2405.20321*, 2024. URL <https://arxiv.org/abs/2405.20321>.
- [166] L. Fu, J. Yu, K. El-Refaï, E. Kou, H. Xue, H. Huang, et al. CaP-X: A Framework for Benchmarking and Improving Coding Agents for Robot Manipulation. *arXiv preprint arXiv:2603.22435*, 2026. URL <https://arxiv.org/abs/2603.22435>.
- [167] F. Torabi, G. Warnell, and P. Stone. Behavioral Cloning from Observation. *arXiv preprint arXiv:1805.01954*, 2018. URL <https://arxiv.org/abs/1805.01954>.
- [168] S. H. Allu, Jishnu Jaykumar P, N. Khargonkar, T. Summers, J. Yao, and Y. Xiang. Hrt1: One-shot human-to-robot trajectory transfer for mobile manipulation, 2025. URL <https://arxiv.org/abs/2510.21026>.
- [169] M. Vecerik, C. Doersch, Y. Yang, T. Davchev, Y. Aytar, G. Zhou, R. Hadsell, L. Agapito, and J. Scholz. RoboTAP: Tracking Arbitrary Points for Few-Shot Visual Imitation, 2023. URL <https://arxiv.org/abs/2308.15975>.
- [170] P. Vitiello, K. Dreczkowski, and E. Johns. One-Shot Imitation Learning: A Pose Estimation Perspective. In *Conference on Robot Learning*, 2023. URL <https://proceedings.mlr.press/v229/vitiello23a.html>.
- [171] O. Biza, S. Thompson, K. R. Pagidi, A. Kumar, E. van der Pol, R. Walters, T. Kipf, J.-W. van de Meent, L. L. S. Wong, and R. Platt. One-shot Imitation Learning via Interaction Warping. In *Conference on Robot Learning*, 2023. URL <https://proceedings.mlr.press/v229/biza23a.html>.
- [172] C. Tang, A. Xiao, Y. Deng, T. Hu, W. Dong, H. Zhang, D. Hsu, and H. Zhang. FUNCTO: Function-Centric One-Shot Imitation Learning for Tool Manipulation, 2025. URL <https://arxiv.org/abs/2502.11744>.
- [173] Y. Liu, J. Mao, J. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling. One-Shot Manipulation Strategy Learning by Making Contact Analogies, 2024. URL <https://arxiv.org/abs/2411.09627>.
- [174] C. de Farias, L. Figueredo, R. Laha, M. Adjigble, B. Tamadazte, R. Stolkin, S. Haddadin, and N. Marturi. GIFT: Geometry-Induced Functional Transfer for Category-level Object Manipulation, 2025. URL <https://arxiv.org/abs/2503.15371>.
- [175] Z. Zhu, Y. Chen, X. Liang, H. Bae, Y. Wang, Z. Yuan, R. Wu, and H. Dong. Sparse Meets Dense: Correspondence Guided Robotic Manipulation with Rigid-Deformable Interactions, 2026. URL <https://arxiv.org/abs/2608.01083>.
- [176] J. Oh, K. Kawaharazuka, and K. Okada. VLAff: Vision-Language-Affordance Model for Unified Actionable Affordances. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2026. URL <https://arxiv.org/abs/2608.05215>. Accepted.
- [177] V. Wichitwechkarn, E. Williams, C. Fox, and R. Choudhary. Annotation-Free One-Shot Imitation Learning for Multi-Step Manipulation Tasks, 2025. URL <https://arxiv.org/abs/2509.24972>.
- [178] S. Wang, J. Shi, Z. Fu, X. He, F. Liu, C. Yang, Y. Zhou, Z. Fei, J. Gong, J. Fu, M. Z. Shou, X. Huang, X. Qiu, and Y.-G. Jiang. World Action Models: The Next Frontier in Embodied AI. *arXiv preprint arXiv:2605.12090*,

2026. URL <https://arxiv.org/abs/2605.12090>.
- [179] A. Srinivas, A. Jabri, P. Abbeel, S. Levine, and C. Finn. Universal Planning Networks. *arXiv preprint arXiv:1804.00645*, 2018. URL <https://arxiv.org/abs/1804.00645>.
 - [180] X. Li, X. He, L. Zhang, M. Wu, X. Li, and Y. Liu. A Comprehensive Survey on World Models for Embodied AI. *arXiv preprint arXiv:2510.16732*, 2025. URL <https://arxiv.org/abs/2510.16732>.
 - [181] T. Nguyen, W. Yuan, S. Wei, H. Li, D. Seita, and Y. Wang. ICLR: In-Context Imitation Learning with Visual Reasoning. *arXiv preprint arXiv:2603.07530*, 2026. URL <https://arxiv.org/abs/2603.07530>.
 - [182] J. Son, J. Kim, K. Kam, J. Coholich, S. J. Kim, J. Kim, C. D. Kim, J. Cho, D. Fox, and Z. Kira. SeeTraceAct: Visibility-Aware Latent Planning from Cross-Embodiment Demonstration Videos, 2026. URL <https://arxiv.org/abs/2606.02745>.
 - [183] S. Choi, J. Jung, S. Kim, Y. Lee, and Y. Yu. PonderPounce: A Pretrained MLLM as an Episode Context Engine for Robot Control, 2026. URL <https://arxiv.org/abs/2608.24115>.
 - [184] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei. VoxPoser: Composable 3D Value Maps for Robotic Manipulation with Language Models. *arXiv preprint arXiv:2307.05973*, 2023. URL <https://arxiv.org/abs/2307.05973>.
 - [185] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. J. Ruano, K. Jeffrey, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, M. Yan, and A. Zeng. Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. *arXiv preprint arXiv:2204.01691*, 2022. URL <https://arxiv.org/abs/2204.01691>.
 - [186] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, P. Sermanet, N. Brown, T. Jackson, L. Luu, S. Levine, K. Hausman, and B. Ichter. Inner Monologue: Embodied Reasoning through Planning with Language Models. *arXiv preprint arXiv:2207.05608*, 2022. URL <https://arxiv.org/abs/2207.05608>.
 - [187] C.-Y. Lin, J.-W. Chen, H.-E. Chang, H.-A. Chen, S.-H. Chang, C.-P. Huang, F.-E. Yang, M.-H. Chen, Y.-T. Chen, Y.-C. F. Wang, and S.-H. Sun. PhysCaP: Grounding Code-as-Policy Agent with Physics-Informed Exploration, 2026. URL <https://arxiv.org/abs/2608.21031>.
 - [188] Y. Feng, N. Leung, J. Wang, L. Yang, H. Qi, and P. Culbertson. A minimalist retargeting-guided reinforcement learning recipe for dexterous manipulation, 2026. URL <https://arxiv.org/abs/2607.11874>.
 - [189] X. Fang, B.-R. Huang, J. Mao, J. Shone, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling. Keypoint Abstraction using Large Models for Object-Relative Imitation Learning, 2024. URL <https://arxiv.org/abs/2410.23254>.
 - [190] S. Haldar and L. Pinto. Point Policy: Unifying Observations and Actions with Key Points for Robot Manipulation, 2025. URL <https://arxiv.org/abs/2502.20391>.
 - [191] T. Li, S. Sun, S. S. Aditya, and N. Figueroa. Elastic Motion Policy: An Adaptive Dynamical System for Robust and Efficient One-Shot Imitation Learning, 2025. URL <https://arxiv.org/abs/2503.08029>.
 - [192] W. Yuan, J. Duan, V. Blukis, W. Pumacay, R. Krishna, A. Murali, A. Mousavian, and D. Fox. Robopoint: A vision-language model for spatial affordance prediction for robotics, 2024. URL <https://arxiv.org/abs/2406.10721>.
 - [193] D. Qu, H. Song, Q. Chen, Y. Yao, X. Ye, Y. Ding, Z. Wang, J. Gu, B. Zhao, D. Wang, and X. Li. Spatialvla: Exploring spatial representations for visual-language-action model. In *Robotics: Science and Systems*, 2025. URL <https://arxiv.org/abs/2501.15830>.
 - [194] F. Chaumette and S. Hutchinson. Visual servo control, part i: Basic approaches. *IEEE Robotics & Automation Magazine*, 13(4):82–90, 2006. doi: 10.1109/MRA.2006.250573. URL https://web.mit.edu/amcp/OldFiles/drg/Chaumette_Part_I.pdf.
 - [195] Y. Fan, Y. Cao, P. Gao, H. Gao, G. Guo, Z. Zhang, J. Tan, B. Chen, Z. Zou, and X. Wang. Retrieve in Time, Correct in Frequency, 2026. URL <https://arxiv.org/abs/2608.04527>.
 - [196] L. Li, R. Yang, and D. Luo. FOCAL: Efficient Fully-Offline Meta-Reinforcement Learning via Distance Metric Learning and Behavior Regularization. *arXiv preprint arXiv:2010.01112*, 2020. URL <https://arxiv.org/abs/2010.01112>.
 - [197] H. Yuan and Z. Lu. Robust Task Representations for Offline Meta-Reinforcement Learning via Contrastive Learning. *arXiv preprint arXiv:2206.10442*, 2022. URL <https://arxiv.org/abs/2206.10442>.
 - [198] H. Li, F. Shen, D. Chen, L. Yang, X. Wang, J. Shi, Z. Bing, Z. Liu, and A. Knoll. ReMem-VLA: Empowering Vision-Language-Action Model with Memory via Dual-Level Recurrent Queries, 2026. URL <https://arxiv.org/abs/2603.12942>.
 - [199] E. Cherepanov, N. Kachaev, D. Zelezetsky, A. Bulatov, A. Pshenitsyn, Y. Kuratov, A. Skrynnik, A. I. Panov, and A. K. Kovalev. μ VLA: On Recurrent Memory for Partially Observable Manipulation in VLA Models, 2026. URL <https://arxiv.org/abs/2606.12497>.

- [200] G. Yang, Z. Tu, Y. Yang, S. Mao, J. Dong, T. Chen, J. Peng, J. Xiong, J. Cao, J. Dai, W. Zhou, Y. Mu, and T. Wang. EventVLA: Event-Driven Visual Evidence Memory for Long-Horizon Vision-Language-Action Policies, 2026. URL <https://arxiv.org/abs/2606.20092>.
- [201] H. Oh, A. M. Salcedo-Vázquez, I. G. Ramirez-Alpizar, and Y. Domae. Robust Instant Policy: Leveraging Student’s t-Regression Model for Robust In-context Imitation Learning of Robot Manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 7973–7980, 2025. doi: 10.1109/IROS60139.2025.11247005. URL <https://arxiv.org/abs/2506.15157>.
- [202] J. Grigsby, L. Fan, and Y. Zhu. AMAGO: Scalable In-Context Reinforcement Learning for Adaptive Agents. In *International Conference on Learning Representations*, 2024. URL <https://arxiv.org/abs/2310.09971>.
- [203] A. Elawady, G. Chhablani, R. Ramrakhya, K. Yadav, D. Batra, Z. Kira, and A. Szot. ReLIC: A Recipe for 64k Steps of In-Context Reinforcement Learning for Embodied AI, 2024. URL <https://arxiv.org/abs/2410.02751>.
- [204] R. Shah, Y. Li, F. Bello, Y. Zhu, and R. Martín-Martín. Memory Retrieval in Visuomotor Policies for Long-Horizon Robot Control. *arXiv preprint arXiv:2606.25136*, 2026. URL <https://arxiv.org/abs/2606.25136>.
- [205] J. Yang, Y. Lu, Y. Su, Y. Zhong, Y. Xie, Y. Zhang, et al. TemporalFlow-VLA: Learning Physically Grounded Execution History for Long-Horizon Robot Manipulation. *arXiv preprint arXiv:2608.26821*, 2026. URL <https://arxiv.org/abs/2608.26821>.
- [206] H. Gao, Z. Ni, X. Wen, S. Xu, and R. Li. AGM: Achievement-Grounded Memory for Closed-Loop Agents with Frozen VLA Policies. *arXiv preprint arXiv:2608.29537*, 2026. URL <https://arxiv.org/abs/2608.29537>.
- [207] Y. Zhang, H. Zhang, F. Gao, X. Li, Z. Liu, C. Zhu, J. Qiu, Y. Yan, J. Liu, W. Tang, Z. Fang, Y. Nie, C. Wei, Y. Wang, W. Ding, and C. Yu. Harness VLA: Steering Frozen VLAs into Reliable Manipulation Primitives via Memory-Guided Agents, 2026. URL <https://arxiv.org/abs/2607.08448v4>.
- [208] V. Giridhar, A. Khandelwal, J. A. Collins, I. Georgiev, and A. Garg. Beyond Imitation: Self-Improving Robot Policies via Off-Policy Q-Planning. *arXiv preprint arXiv:2608.21204*, 2026. URL <https://arxiv.org/abs/2608.21204>.
- [209] H. Shi, W. Li, B. Xie, Y. Wang, R. Zhou, T. Wang, X. Zhang, P. Luo, and G. Huang. MemoryVLA++: Temporal Modeling via Memory and Imagination in Vision-Language-Action Models, 2026. URL <https://arxiv.org/abs/2606.09827>.
- [210] W. Chen, J. S. Bhatia, C. Glossop, N. Mathihalli, R. Doshi, A. Tang, D. Driess, K. Pertsch, and S. Levine. Steerable Vision-Language-Action Policies for Embodied Reasoning and Hierarchical Control. *arXiv preprint arXiv:2602.13193*, 2026. URL <https://arxiv.org/abs/2602.13193>.
- [211] H. Liu, X. Li, S. Yao, P. Shi, T. Zhou, J.-B. Huang, F. Huang, and J. Mao. Guava: An Effective and Universal Harness for Embodied Manipulation, 2026. URL <https://arxiv.org/abs/2606.18363>.
- [212] L. Galanti, D. Shah, and T. Dao. Addressing the Orchestration Gap in Generalist Robots via Physical Agency, 2026. URL <https://arxiv.org/abs/2607.21725>.
- [213] H. Singh, R. J. Das, M. Han, P. Nakov, and I. Laptev. MALMM: Multi-Agent Large Language Models for Zero-Shot Robotics Manipulation, 2024. URL <https://arxiv.org/abs/2411.17636>.
- [214] A. Curtis, N. Kumar, J. Cao, T. Lozano-Pérez, and L. P. Kaelbling. Trust the PRoC3S: Solving Long-Horizon Robotics Problems with LLMs and Constraint Satisfaction. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pages 1362–1383, 2025. URL <https://proceedings.mlr.press/v270/curtis25a.html>.
- [215] G. Sarch, L. Jang, M. J. Tarr, W. W. Cohen, K. Marino, and K. Fragkiadaki. VLM Agents Generate Their Own Memories: Distilling Experience into Embodied Programs of Thought. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://papers.nips.cc/paper_files/paper/2024/hash/8ac50fd0a4eeeb1f077b17bb7c5353c3-Abstract-Conference.html.
- [216] R. Ju, X. Wang, X. Ding, Y. Yang, H. Wu, S. Jiang, Q. Zhang, H. Wen, X. Li, W. Wang, K. Li, Y. Liu, H. Dai, W. Wang, and T. Cao. EmbodiSkill: Skill-Aware Reflection for Self-Evolving Embodied Agents, 2026. URL <https://arxiv.org/abs/2605.10332>.
- [217] S. Xie, Y. Zhang, R. Wang, and X. Chen. Uni-Skill: Building Self-Evolving Skill Repository for Generalizable Robotic Manipulation, 2026. URL <https://arxiv.org/abs/2603.02623>.
- [218] V. Kumar. Act-Observe-Rewrite: Multimodal Coding Agents as In-Context Policy Learners for Robot Manipulation, 2026. URL <https://arxiv.org/abs/2603.04466>.
- [219] J. Hu, M. Shridhar, C. Lu, D. Shah, H.-T. L. Chiang, J. Tan, and A. Xie. What Matters in Orchestrating Robot Policies: A Systematic Study of Hierarchical VLA Agents, 2026. URL <https://arxiv.org/abs/2606.10267>.
- [220] R. Li, Y. Zhou, Y. Zhu, K. Chen, J. Wang, S. Wang, K. Hu, M. Yu, B. Jiang, Z. Su, J. Ma, X. He, Y. Shen, Y. Yang, G. Ren, M. Yao, W. Wang, and Y. Mu. RoboClaw: An Agentic Framework for Scalable Long-Horizon Robotic Tasks, 2026. URL <https://arxiv.org/abs/2603.11558>.
- [221] Y. Ye, Y. Fu, Y. Lv, B. Hou, J. Cen, L. Kong, D. Zheng, T. Chen, J. Liu, Z. Cao, Y. Lou, W. Chow, X. Sun, Y. Wang, K. Ge, X. Chi, X. Zhang, Z. Pang, Y. Zhong, S. Han, Z. Lu, W. Yuan, Q. Chen, M. Y. Wang, Y. Mu, Z. Liu, J. Yang, P. Luo, and S. Zhang. Data Pyramid for Embodied Manipulation: A Survey, 2026. URL <https://arxiv.org/abs/2607.24744>.

- [222] Y. Wang, P. Lin, X.-H. Chen, H. Yuan, Z. Liang, Y. Huang, A. Chen, Z. Lei, J. Zhang, T. Zhang, H. Li, T. Zhang, C. Xiao, Z. Jiao, and Q. Jin. Ego2Robot: Scalable Robot Data Synthesis from Egocentric Human Data, 2026. URL <https://arxiv.org/abs/2608.02580>.
- [223] Simple AI, Y. Wei, J. Ma, J. Wang, W. Zhou, Y. Zuo, K. Rui, M. Li, J. Zhang, Z. Pan, X. Wang, H. Jia, H. Du, Z. Zeng, J. Ma, G. Qin, D. Zhang, and X. Li. HiFi-UMI: Learning Deployable Manipulation Policies from High-Fidelity UMI Data Alone, 2026. URL <https://arxiv.org/abs/2607.25895>.
- [224] Zhaxizhuoma, K. Liu, C. Guan, Z. Jia, Z. Wu, X. Liu, T. Wang, S. Liang, P. Chen, P. Zhang, H. Song, D. Qu, D. Wang, Z. Wang, N. Cao, Y. Ding, B. Zhao, and X. Li. Fastumi: A scalable and hardware-independent universal manipulation interface with dataset, 2024. URL <https://arxiv.org/abs/2409.19499>.
- [225] H. Ha, Y. Gao, Z. Fu, J. Tan, and S. Song. Umi on legs: Making manipulation policies mobile with manipulation-centric whole-body controllers. In *Conference on Robot Learning*, 2024. URL <https://arxiv.org/abs/2407.10353>.
- [226] S. Luo, Y. Li, Y. Hu, C. Yu, C. Xu, J. Zhang, G. Yao, T. Huang, R. He, and Z. Wang. Omniumi: Towards physically grounded robot learning via human-aligned multimodal interaction, 2026. URL <https://arxiv.org/abs/2604.10647>.
- [227] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison. RL Bench: The Robot Learning Benchmark & Learning Environment. *arXiv preprint arXiv:1909.12271*, 2019. URL <https://arxiv.org/abs/1909.12271>.
- [228] B. Liu, Y. Zhu, C. Gao, Y. Feng, Q. Liu, Y. Zhu, and P. Stone. LIBERO: Benchmarking Knowledge Transfer for Lifelong Robot Learning. *arXiv preprint arXiv:2306.03310*, 2023. URL <https://arxiv.org/abs/2306.03310>.
- [229] J. Gu, F. Xiang, X. Li, Z. Ling, X. Liu, T. Mu, Y. Tang, S. Tao, X. Wei, Y. Yao, X. Yuan, P. Xie, Z. Huang, R. Chen, and H. Su. Maniskill2: A unified benchmark for generalizable manipulation skills. In *International Conference on Learning Representations*, 2023. URL <https://arxiv.org/abs/2302.04659>.
- [230] T. Chen, Z. Chen, B. Chen, Z. Cai, Y. Liu, Z. Li, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation, 2025. URL <https://arxiv.org/abs/2506.18088>.
- [231] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL <https://arxiv.org/abs/1706.03762>.
- [232] G. Reddy. The mechanistic basis of data dependence and abrupt learning in an in-context classification task. *arXiv preprint arXiv:2312.03002*, 2023. URL <https://arxiv.org/abs/2312.03002>.
- [233] R. Yu, P. Zhang, S. Liu, B. Liu, M. Kang, S. Li, L. Shi, E. Ma, P. Yang, C. Pan, J. Chen, D. Liu, R. Sun, M. Guo, B. Zhang, H. Zhou, Z. Xu, V. Chen, H. Huang, J. Wang, D. Kuzi, A. Zhai, H. Su, R. Gan, L. Liang, H. Wang, and Q. Wang. Wall-OSS-0.5 Technical Report, 2026. URL <https://arxiv.org/abs/2605.30877v2>.
- [234] P. de Haan, D. Jayaraman, and S. Levine. Causal Confusion in Imitation Learning. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://papers.neurips.cc/paper_files/paper/2019/hash/947018640bf36a2bb609d3557a285329-Abstract.html.
- [235] M. Torne, A. Tang, Y. Liu, and C. Finn. Learning Long-Context Diffusion Policies via Past-Token Prediction. *arXiv preprint arXiv:2505.09561*, 2025. URL <https://arxiv.org/abs/2505.09561>.
- [236] D. Kalashnikov, J. Varley, Y. Chebotar, B. Swanson, R. Jonschkowski, C. Finn, S. Levine, and K. Hausman. MT-Opt: Continuous Multi-Task Robotic Reinforcement Learning at Scale, 2021. URL <https://arxiv.org/abs/2104.08212>.
- [237] F. Ebert, Y. Yang, K. Schmeckpeper, B. Bucher, G. Georgakis, K. Daniilidis, C. Finn, and S. Levine. Bridge Data: Boosting Generalization of Robotic Skills with Cross-Domain Datasets, 2021. URL <https://arxiv.org/abs/2109.13396>.
- [238] AgiBot World Team. AgiBot World 2026. Official dataset documentation; accessed 13 September 2026, 2026. URL <https://huggingface.co/datasets/agibot-world/AgibotWorld2026>.
- [239] AGIBOT. AGIBOT Releases AGIBOT WORLD 2026 Theme 3 for Real-World Robot Reinforcement Learning. Official release, 1 September 2026, September 2026. URL <https://agibot.com/article/231/detail/95.html>.
- [240] K. Wu, C. Hou, J. Liu, Z. Che, X. Ju, Z. Yang, M. Li, Y. Zhao, Z. Xu, G. Yang, S. Fan, X. Wang, F. Liao, Z. Zhao, G. Li, Z. Jin, L. Wang, J. Mao, N. Liu, P. Ren, Q. Zhang, Y. Lyu, M. Liu, J. He, Y. Luo, Z. Gao, C. Li, C. Gu, Y. Fu, D. Wu, X. Wang, S. Chen, Z. Wang, P. An, S. Qian, S. Zhang, and J. Tang. RoboMIND: Benchmark on Multi-embodiment Intelligence Normative Data for Robot Manipulation, 2024. URL <https://arxiv.org/abs/2412.13877>.
- [241] C. Hou, K. Wu, J. Liu, Z. Che, D. Wu, F. Liao, G. Li, J. He, Q. Feng, Z. Jin, C. Gu, Z. Liu, N. Han, X. Mi, Y. Lv, Y. Fu, G. Dai, L. Gu, T. Li, Y. Zhang, Y. Zhang, X. Wang, S. Fan, M. Li, Z. Zhao, N. Liu, Z. Xu, P. Ren, J. Ji, H. Liu, K. Cheng, S. Zhang, and J. Tang. RoboMIND 2.0: A Multimodal, Bimanual Mobile Manipulation Dataset for Generalizable Embodied Intelligence, 2025. URL <https://arxiv.org/abs/2512.24653>.
- [242] S. Wu, X. Liu, S. Xie, P. Wang, X. Li, B. Yang, Z. Li, K. Zhu, H. Wu, Y. Liu, Z. Long, R. Xu, Y. Wang, C. Liu, D. Wang, Z. Ni, X. Yang, Y. Liu, R. Feng, L. Zhang, D. Huang, C. Jin, A. Yin, X. Wang, Z. Sun, J. Zhao, M. Du, M. Cao, X. Chen, H. Cheng, X. Zhang, Y. Fu, N. Chen, C. Chi, S. Chen, H. Lyu, X. Hao, Y. Wang,

- B. Lei, D. Liu, X. Yang, Y. Jiao, T. Pan, Y. Zhang, S. Wang, Z. Zhang, X. Liu, J. Zhang, C. Meng, Z. Zhang, J. Gao, S. Wang, X. Leng, Z. Xie, Z. Zhou, P. Huang, W. Yang, Y. Guo, Y. Zhu, S. Zheng, H. Cheng, X. Ding, Y. Yue, H. Wang, C. Chen, J. Pang, Y. Qian, H. Geng, L. Gao, H. Li, B. Fang, G. Huang, Y. Yang, H. Dong, H. Wang, H. Zhao, Y. Mu, D. Hu, H. Zhao, T. Huang, S. Zhang, Y. Lin, Z. Wang, and G. Yao. RoboCOIN: An Open-Sourced Bimanual Robotic Data Collection for Integrated Manipulation, 2025. URL <https://arxiv.org/abs/2511.17441>.
- [243] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar. RoboAgent: Generalization and Efficiency in Robot Manipulation via Semantic Augmentations and Action Chunking. *arXiv preprint arXiv:2309.01918*, 2023. URL <https://arxiv.org/abs/2309.01918>.
- [244] J. Xu, Q. Li, Y. Shen, Y. Ren, T. Davies, S. He, Z. Wang, Y. Yang, R. Cheng, and H. Dong. Scaling Bimanual Household Manipulation from 1,500 hours of Demonstrations to On-Policy Corrections, 2026. URL <https://arxiv.org/abs/2609.03591>.
- [245] R. Shah, S. Liu, Q. Wang, Z. Jiang, S. Kumar, M. Seo, et al. MimicDroid: In-Context Learning for Humanoid Robot Manipulation from Human Play Videos. *arXiv preprint arXiv:2509.09769*, 2025. URL <https://arxiv.org/abs/2509.09769>.
- [246] H. Li, G. Zhao, Y. Liu, H. Hou, G. Ye, T. Fang, C. Liu, S. Huang, J. Liu, X. Wang, and H. Li. ACE-Ego-0: Unifying Egocentric Human and Robotic Data for VLA Pretraining, 2026. URL <https://arxiv.org/abs/2606.17200>.
- [247] X. R. Team, J. Guo, P. Jin, J. Li, P. Li, Y. Li, F. Liu, W. Peng, O. Qin, Y. Su, N. Sun, Q. Sun, R. Suo, H. Wang, Y. Wang, R. Wu, C. Xia, L. Zhang, J. Zhao, G. Chen, W. Chen, X. He, B. Li, Q. Li, Z. Li, H. Qu, W. Song, D. Xiang, Y. Xie, P. Xu, H. Ye, W. Ye, H. Zhao, and Q. Zhou. Xiaomi-Robotics-1: Scaling Vision-Language-Action Models with over 100K Hours of Real-World Trajectories, 2026. URL <https://arxiv.org/abs/2607.15330>.
- [248] R. Punamiya, S. Kareer, Z. Liu, J. Citron, R.-Z. Qiu, X. Cai, A. Gavryushin, J. Chen, D. Liconti, L. Y. Zhu, P. Aphiwetsa, B. Li, A. Cheluva, P. Kuppili, Y. Liu, D. Patel, A. Gao, H.-Y. Chung, R. Co, R. Zbizika, J. Liu, X. Xu, H. Xiong, G. Chen, S. Oliani, W. Xuan, C. Yang, X. Wang, J. Fort, R. Newcombe, J. Gao, J. Chong, G. Matsuda, A. Doriwala, M. Pollefeys, R. Katzschnmann, X. Wang, S. Song, J. Hoffman, and D. Xu. EgoVerse: An Egocentric Human Dataset for Robot Learning from Around the World, 2026. URL <https://arxiv.org/abs/2604.07607>.
- [249] Z. Li, B. Yang, C. Miao, K. Zhu, H. Chen, Q. Guan, Z. Wu, W. Zhan, Y. Sun, Z. Huang, Z. Shan, Z. Jin, J. Hong, T. Wang, Y. Feng, Y. Liu, Y. Wang, Y. Yang, Z. Zhou, M. Luo, H. Cheng, B. Zhang, J. Li, J. Cai, G. Yao, J. Zhang, C. Lin, R. Xu, L. Yu, C. Shen, C. Shen, and Z. Li. Open-AoE: An Open Egocentric Manipulation Dataset and Toolchain for Embodied Learning, 2026. URL <https://arxiv.org/abs/2607.14183>.
- [250] Y. Cao, H. Xie, B. Wen, R. Yao, Y. Liu, Y. Huang, Z. Liao, Y. Wang, H. Liu, X. Tian, D. Su, L. Zhuo, D. Tao, X. Wang, L. Pan, and Z. Liu. ACE-Data-0: Human-Centric Ambient Capture as Embodied Data Engine, 2026. URL <https://arxiv.org/abs/2607.28625>.
- [251] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar. MimicPlay: Long-Horizon Imitation Learning by Watching Human Play. *arXiv preprint arXiv:2302.12422*, 2023. URL <https://arxiv.org/abs/2302.12422>.
- [252] OpenAI. Public summary of training content for GPT-6 Astra. EU AI Act training-content summary, version 1, September 2026. URL <https://cdn.openai.com/pdf/gpt-6-astra-eu-ai-act-public-summary-of-training-content.pdf>. Updated September 2, 2026; Section 1.3, page 2.
- [253] A. Raventós, M. Paul, F. Chen, and S. Ganguli. Pretraining task diversity and the emergence of non-Bayesian in-context learning for regression. *arXiv preprint arXiv:2306.15063*, 2023. URL <https://arxiv.org/abs/2306.15063>.
- [254] S. K. S. Ghasemipour, A. Wahid, J. Tompson, P. Sanketi, and I. Mordatch. Self-Improving Embodied Foundation Models, 2025. URL <https://arxiv.org/abs/2509.15155>.
- [255] H. Qian, Y. Chen, Y. Xie, K. Ren, P. Chanrungrameekul, G. Wang, B. Wen, C. Wei, and K. Hang. RoboTok: An Internet-Scale Data Engine for Human Demonstration Retrieval and Dexterous Manipulation Learning, 2026. URL <https://arxiv.org/abs/2609.03199>.
- [256] N. Lin, T. Ohkawa, S. Chen, R. Wen, Z. Li, L. Huang, Z. Zhu, Y. Bao, Y. Li, M. Cai, X. Ma, W. Xu, and Y. Sato. SiMDex: Mining Similar Egocentric Videos for Cross-Embodiment Dexterous Manipulation, 2026. URL <https://arxiv.org/abs/2608.04196>.
- [257] H. Zhou, R. Wang, Y. Tai, Y. Deng, G. Liu, and K. Jia. You Only Teach Once: Learn One-Shot Bimanual Robotic Manipulation from Video Demonstrations. In *Robotics: Science and Systems*, 2025. URL <https://www.roboticsproceedings.org/rss21/p149.html>.
- [258] H. Zhou and K. Jia. One-Shot Real-World Demonstration Synthesis for Scalable Bimanual Manipulation. In *Robotics: Science and Systems*, 2026. URL <https://www.roboticsproceedings.org/rss22/p001.html>.

- [259] S. Yang, W. Yu, J. Zeng, J. Lv, K. Ren, C. Lu, D. Lin, and J. Pang. Novel Demonstration Generation with Gaussian Splatting Enables Robust One-Shot Manipulation. In *Robotics: Science and Systems*, 2025. URL <https://www.roboticsproceedings.org/rss21/p146.html>.
- [260] A. Bonardi, S. James, and A. J. Davison. Learning One-Shot Imitation from Humans without Humans. *arXiv preprint arXiv:1911.01103*, 2019. URL <https://arxiv.org/abs/1911.01103>.
- [261] H. Kang, X. Ye, Y. Liu, S. H. Mantri, L. Mao, J. Fleming, et al. SimWorld Studio: Automatic Environment Generation with Evolving Coding Agent for Embodied Agent Learning. *arXiv preprint arXiv:2605.09423*, 2026. URL <https://arxiv.org/abs/2605.09423>.
- [262] W. Wang, Z. Li, Y. Lei, D. Martinez, X. Gao, Y. Jiang, et al. SUN: Persistent Programs For Language-Grounded Control-to-Learning-to-Real Policies. *arXiv preprint arXiv:2608.31167*, 2026. URL <https://arxiv.org/abs/2608.31167>.
- [263] S. Nasiriany, S. Nasiriany, A. Maddukuri, and Y. Zhu. RoboCasa365: A Large-Scale Simulation Framework for Training and Benchmarking Generalist Robots. *arXiv preprint arXiv:2603.04356*, 2026. URL <https://arxiv.org/abs/2603.04356>.
- [264] P. Zhao, J. Lin, T. Rupprecht, A. Akbari, C. Yang, R. Chowdhury, et al. PhyWorld: Physics-Faithful World Model for Video Generation. *arXiv preprint arXiv:2605.19242*, 2026. URL <https://arxiv.org/abs/2605.19242>.
- [265] Y. Liu, F. Feng, L. Kong, W. Lu, J. Tang, K. Zhang, et al. World Action Verifier: Self-Improving World Models via Forward-Inverse Asymmetry. *arXiv preprint arXiv:2604.01985*, 2026. URL <https://arxiv.org/abs/2604.01985>.
- [266] M. Shi, L. Chen, J. Chen, Y. Lu, C. Liu, G. Ren, P. Luo, D. Huang, M. Yao, and H. Li. Is Diversity All You Need for Scalable Robotic Manipulation?, 2025. URL <https://arxiv.org/abs/2507.06219v1>.
- [267] S. Ross, G. J. Gordon, and J. A. Bagnell. 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, pages 627–635, 2011. URL <https://proceedings.mlr.press/v15/ross11a.html>.
- [268] M. Laskey, J. Lee, R. Fox, A. Dragan, and K. Goldberg. DART: Noise Injection for Robust Imitation Learning, 2017. URL <https://arxiv.org/abs/1703.09327>.
- [269] R. Hoque, A. Balakrishna, E. Novoseller, A. Wilcox, D. S. Brown, and K. Goldberg. ThriftyDagger: Budget-Aware Novelty and Risk Gating for Interactive Imitation Learning, 2021. URL <https://arxiv.org/abs/2109.08273>.
- [270] D. Li, J. Lei, H. Wang, L. Liu, Y. Yang, Z. Wang, et al. Learning Actionable Manipulation Recovery via Counterfactual Failure Synthesis. *arXiv preprint arXiv:2603.13528*, 2026. URL <https://arxiv.org/abs/2603.13528>.
- [271] Y. Li, Z. Zhou, Y. Chen, Y. Guo, J. Liu, S. Zhang, et al. Hi-WM: Human-in-the-World-Model for Scalable Robot Post-Training. *arXiv preprint arXiv:2604.21741*, 2026. URL <https://arxiv.org/abs/2604.21741>.
- [272] B. Sukhija, O. Groth, M. Shridhar, T. Hertweck, M. Bloesch, M. Wulfmeier, et al. EXIMO: VLM Guided Exploration of VLA Policies. *arXiv preprint arXiv:2608.19891*, 2026. URL <https://arxiv.org/abs/2608.19891>.
- [273] Physical Intelligence et al. $\pi_{0.6}^*$: a VLA That Learns From Experience, 2025. URL <https://arxiv.org/abs/2511.14759>.
- [274] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. *arXiv preprint arXiv:2304.13705*, 2023. URL <https://arxiv.org/abs/2304.13705>.
- [275] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine. Octo: An Open-Source Generalist Robot Policy. *arXiv preprint arXiv:2405.12213*, 2024. URL <https://arxiv.org/abs/2405.12213>.
- [276] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. RT-1: Robotics Transformer for Real-World Control at Scale. *arXiv preprint arXiv:2212.06817*, 2022. URL <https://arxiv.org/abs/2212.06817>.
- [277] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choromanski, T. Ding, D. Driess, A. Dubey, C. Finn, P. Florence, C. Fu, M. G. Arenas, K. Gopalakrishnan, K. Han, K. Hausman, A. Herzog, J. Hsu, B. Ichter, A. Irpan, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, L. Lee, T.-W. E. Lee, S. Levine, Y. Lu, H. Michalewski, I. Mordatch, K. Pertsch, K. Rao, K. Reymann, M. Ryoo, G. Salazar, P. Sanketi, P. Sermanet, J. Singh, A. Singh, R. Soricut, H. Tran, V. Vanhoucke, Q. Vuong, A. Wahid, S. Welker, P. Wohlhart, J. Wu, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. *arXiv preprint arXiv:2307.15818*, 2023. URL <https://arxiv.org/abs/2307.15818>.

- [278] L. Wang, X. Chen, J. Zhao, and K. He. Scaling Proprioceptive-Vision Learning with Heterogeneous Pre-trained Transformers, 2024. URL <https://arxiv.org/abs/2409.20537>.
- [279] S. Liu, L. Wu, B. Li, H. Tan, H. Chen, Z. Wang, K. Xu, H. Su, and J. Zhu. RDT-1B: a Diffusion Foundation Model for Bimanual Manipulation. *arXiv preprint arXiv:2410.07864*, 2024. URL <https://arxiv.org/abs/2410.07864>.
- [280] S. Kareer, K. Pertsch, J. Darpinian, J. Hoffman, D. Xu, S. Levine, C. Finn, and S. Nair. Emergence of Human to Robot Transfer in Vision-Language-Action Models, 2025. URL <https://arxiv.org/abs/2512.22414>.
- [281] Physical Intelligence et al. $\pi_{0.7}$: a Steerable Generalist Robotic Foundation Model with Emergent Capabilities, 2026. URL <https://arxiv.org/abs/2604.15483>.
- [282] A. D. Vuong, M. N. Vu, D. An, and I. Reid. Action Tokenizer Matters in In-Context Imitation Learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 13490–13496, 2025. doi: 10.1109/IROS60139.2025.11246836. URL <https://arxiv.org/abs/2503.01206>.
- [283] F. J. Fateh, A. S. Ali, M. Popattia, U. Nizamani, A. Konin, M. Z. Zia, and Q.-H. Tran. A Hierarchical Spatiotemporal Action Tokenizer for In-Context Imitation Learning in Robotics, 2026. URL <https://arxiv.org/abs/2604.15215>.
- [284] Y. Hong, J. Liu, H. Yin, M. Li, L. Guibas, L. Fei-Fei, et al. ESI-Bench: Towards Embodied Spatial Intelligence that Closes the Perception-Action Loop. *arXiv preprint arXiv:2605.18746*, 2026. URL <https://arxiv.org/abs/2605.18746>.
- [285] H. J. Lillemark, B. Huang, F. Zhan, Y. Du, and T. A. Keller. Flow Equivariant World Models: Memory for Partially Observed Dynamic Environments. *arXiv preprint arXiv:2601.01075*, 2026. URL <https://arxiv.org/abs/2601.01075>.
- [286] Z. Zheng, J. Yu, X. Peng, jun shi, M. Li, C. Zhang, et al. Mem-World: Memory-Augmented Action-Conditioned World Models for Persistent Robot Manipulation. *arXiv preprint arXiv:2606.18960*, 2026. URL <https://arxiv.org/abs/2606.18960>.
- [287] Y. Lee, H. Ju, and G. Kim. LT-Mem: Volatility-Aware Spatio-Temporal Memory for Lifelong Scene Understanding. *arXiv preprint arXiv:2608.19059*, 2026. URL <https://arxiv.org/abs/2608.19059>.
- [288] M. Liao, C. Ye, J. Fei, and W. Lin. DSG: Dynamic 3D Scene Graph Construction for Embodied Agents in Changing Indoor Environments. *arXiv preprint arXiv:2609.00619*, 2026. URL <https://arxiv.org/abs/2609.00619>.
- [289] R. Tang, C. Zou, G. Wu, J. Wei, W. Chen, and J. Zhu. HitMem: Hierarchical Temporal 3D Memory with Multi-Modal Context-Aware Retrieval for Dynamic Environments. *arXiv preprint arXiv:2609.00950*, 2026. URL <https://arxiv.org/abs/2609.00950>.
- [290] L. Li, M. Zhu, Z. Zhao, H. Zhao, K. Liu, L. Zhong, et al. Where to Look: Can Foundation Models Reach a Target Viewpoint Through Active Exploration? *arXiv preprint arXiv:2606.01247*, 2026. URL <https://arxiv.org/abs/2606.01247>.
- [291] V. Bhat, S. Chen, A. Zook, X. Yang, S. Birchfield, V. Blukis, et al. PROBE: Manipulation-Grounded Visual Question Answering with VLM Agents. *arXiv preprint arXiv:2608.17129*, 2026. URL <https://arxiv.org/abs/2608.17129>.
- [292] Z. Mai, C. Wang, C. Zeng, S. Mao, H. Zhang, S. Zhou, and C. Yang. CR-VLA-Force: Learning Control-aware Compliance VLA Model for Robust Contact-rich Robotic Manipulation, 2026. URL <https://arxiv.org/abs/2609.05832v1>.
- [293] S. Jain, S. Sachdeva, and R. Paul. Enabling Long(er) Horizon Imitation for Manipulation Tasks by Modeling Subgoal Transitions. In *Conference on Robot Learning*, 2025. URL <https://proceedings.mlr.press/v305/jain25b.html>.
- [294] F. Ebert, S. Dasari, A. X. Lee, S. Levine, and C. Finn. Robustness via Retrying: Closed-Loop Robotic Manipulation with Self-Supervised Learning. In *Conference on Robot Learning*, 2018. URL <https://proceedings.mlr.press/v87/ebert18a.html>.
- [295] A. Bahety, P. Mandikal, B. Abbatematteo, and R. Martín-Martín. ScrewMimic: Bimanual Imitation from Human Videos with Screw Space Projection. In *Robotics: Science and Systems*, 2024. URL <https://www.roboticsproceedings.org/rss20/p078.html>.
- [296] H. Zhang, Y. Lu, B. Wang, X. Kang, Y.-L. Kuo, Z. Cheng, et al. Foresight: Failure Detection for Long-Horizon Robotic Manipulation with Action-Conditioned World Model Latents. *arXiv preprint arXiv:2606.23085*, 2026. URL <https://arxiv.org/abs/2606.23085>.
- [297] T. Matsushima, N. Kondo, Y. Iwasawa, K. Nasuno, and Y. Matsuo. Modeling Task Uncertainty for Safe Meta-Imitation Learning. *Frontiers in Robotics and AI*, 7:606361, 2020. doi: 10.3389/frobt.2020.606361. URL <https://doi.org/10.3389/frobt.2020.606361>.
- [298] K. Liang, Z. Zhang, and J. F. Fisac. Introspective Planning: Aligning Robots’ Uncertainty with Inherent Task Ambiguity. *arXiv preprint arXiv:2402.06529*, 2024. URL <https://arxiv.org/abs/2402.06529>.
- [299] J. Zhou, B. Zhao, Y. Zhang, Y. Chen, W. Chen, Q. Li, et al. TacForcing: Streaming Action Generation with Execution-Time Tactile Feedback. *arXiv preprint arXiv:2608.25798*, 2026. URL

- <https://arxiv.org/abs/2608.25798>.
- [300] J. Zhang, J. Wang, H. Hu, T. Chen, Y. Chen, C. Fan, and C. Zhang. MetaCURE: Meta Reinforcement Learning with Empowerment-Driven Exploration. *arXiv preprint arXiv:2006.08170*, 2020. URL <https://arxiv.org/abs/2006.08170>.
 - [301] M. Cakmak and A. L. Thomaz. Designing Robot Learners that Ask Good Questions. In *ACM/IEEE International Conference on Human-Robot Interaction*, 2012. URL <https://homes.cs.washington.edu/~mcakmak/pdfs/2012/cakmak2012hri.pdf>.
 - [302] Y. Zhang, D. Liu, K. Ye, C. Song, X. Li, M. Hariri, V. Singh, Y. Yin, and V. Chaudhary. Imagining Recovery: Inference-Time Counterfactual Realignment for Vision-Language-Action Models, 2026. URL <https://arxiv.org/abs/2608.14822>.
 - [303] S. Xu, R. Jin, H. Zhou, B. Yue, G. Qiao, Y. Tai, et al. From Reaction to Anticipation: Proactive Failure Recovery through Agentic Task Graph for Robotic Manipulation. *arXiv preprint arXiv:2605.11951*, 2026. URL <https://arxiv.org/abs/2605.11951>.
 - [304] S. Chen, H. Hadfield, A. Zook, M. A. Uy, C. H. Song, E. Coumans, X. Yang, F. Ladhak, Q. Qu, S. Birchfield, J. Tremblay, and V. Blukis. VoLo: A Physical Orchestrator for Open-Vocabulary Long-Horizon Manipulation, 2026. URL <https://arxiv.org/abs/2606.07723>.
 - [305] H. Li, Z. Jiao, X. Liu, H. Liu, and Z. Zheng. In-situ Value-aligned Human-Robot Interactions with Physical Constraints. *arXiv preprint arXiv:2508.07606*, 2025. URL <https://arxiv.org/abs/2508.07606>.
 - [306] B. Jing, M. Wang, R. Hao, C. Ge, H. Shen, J. He, et al. SoftVTBench: A Deformation-Aware Visuo-Tactile Dataset and Benchmark for Deformable-Object Manipulation. *arXiv preprint arXiv:2608.18701*, 2026. URL <https://arxiv.org/abs/2608.18701>.
 - [307] Z. Zhu, L. Guo, Y. Mei, B. Pang, Z. Zhang, X. He, et al. When Robots Do the Chores: A Benchmark and Agent for Long-Horizon Household Task Execution. *arXiv preprint arXiv:2605.14504*, 2026. URL <https://arxiv.org/abs/2605.14504>.
 - [308] Y. Zhan, B. He, Z. Tang, and C. Xiao. PACE-Bench: Benchmarking Physics Adaptation via Code Evolution in Dynamic Environments. *arXiv preprint arXiv:2608.14441*, 2026. URL <https://arxiv.org/abs/2608.14441>.
 - [309] T. Jiang, X. Tan, S. Wheeler, L. Sun, T. W. Ayalew, and M. Walter. What Are We Actually Benchmarking in Robot Manipulation?, 2026. URL <https://arxiv.org/abs/2606.04233>.
 - [310] S. Fei, S. Wang, J. Shi, Z. Dai, J. Cai, P. Qian, L. Ji, X. He, S. Zhang, Z. Fei, J. Fu, J. Gong, and X. Qiu. LIBERO-Plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025. URL <https://arxiv.org/abs/2510.13626>.
 - [311] T. Yu, D. Quillen, Z. He, R. Julian, A. Narayan, H. Shively, A. Bellathur, K. Hausman, C. Finn, and S. Levine. Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning. *arXiv preprint arXiv:1910.10897*, 2019. URL <https://arxiv.org/abs/1910.10897>.
 - [312] W. Gu, A. Gupta, A. Sah, R. Varghese, L. S. Vanam, P. Adireddi, P. Karkus, and N. Gopalan. Monkey See, Can Monkey Do? A Benchmark for Evaluating Robot Skill Learning by Observation, 2026. URL <https://arxiv.org/abs/2609.08209>. Version 1, accessed 11 September 2026.
 - [313] E. Cherepanov, N. Kachaev, A. K. Kovalev, and A. I. Panov. Memory, Benchmark & Robots: A Benchmark for Solving Complex Tasks with Reinforcement Learning. *arXiv preprint arXiv:2502.10550*, 2025. URL <https://arxiv.org/abs/2502.10550>.
 - [314] D. Rong, Y. Shi, C. Ma, J. Cao, Z. Wang, Z. Zhang, et al. H2R-Bench: Benchmarking Human-to-Robot Manipulation Video Generation in World Models. *arXiv preprint arXiv:2608.13049*, 2026. URL <https://arxiv.org/abs/2608.13049>.
 - [315] J. Kwok, C. Agia, R. Sinha, M. Foutter, S. Li, I. Stoica, A. Mirhoseini, and M. Pavone. RoboMonkey: Scaling Test-Time Sampling and Verification for Vision-Language-Action Models, 2025. URL <https://arxiv.org/abs/2506.17811>.
 - [316] X. Cai, Y. Cai, B. Chen, J. Chen, Z. Chen, S. Feng, et al. $\$ \tau_0 \$$ -VLA: a Hierarchical Robot Foundation Model with World-Model-Guided Test-Time Computation. *arXiv preprint arXiv:2608.16885*, 2026. URL <https://arxiv.org/abs/2608.16885>.
 - [317] M. Sato, Y. Iwasawa, Y. Tang, and S. Kuroki. SAIL: Test-Time Scaling for In-Context Imitation Learning with VLM. *arXiv preprint arXiv:2603.08269*, 2026. URL <https://arxiv.org/abs/2603.08269>.
 - [318] S. M. Xie, A. Raghunathan, P. Liang, and T. Ma. An explanation of in-context learning as implicit Bayesian inference. *arXiv preprint arXiv:2111.02080*, 2021. URL <https://arxiv.org/abs/2111.02080>.
 - [319] Y. Chen, C. Zhao, Z. Yu, K. McKeown, and H. He. Parallel Structures in Pre-training Data Yield In-Context Learning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8582–8592, 2024. doi: 10.18653/v1/2024.acl-long.465. URL <https://aclanthology.org/2024.acl-long.465/>.
 - [320] S. Min, X. Lyu, A. Holtzman, M. Artetxe, M. Lewis, H. Hajishirzi, and L. Zettlemoyer. Rethinking the role of demonstrations: What makes in-context learning work? In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, 2022. doi: 10.18653/v1/2022.emnlp-main.759.

- URL <https://aclanthology.org/2022.emnlp-main.759/>.
- [321] S. Garg, D. Tsipras, P. Liang, and G. Valiant. What can transformers learn in-context? a case study of simple function classes. *arXiv preprint arXiv:2208.01066*, 2022. URL <https://arxiv.org/abs/2208.01066>.
 - [322] L. Shen, A. Mishra, and D. Khashabi. Position: Do Pretrained Transformers Learn In-Context by Gradient Descent? In *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 44712–44740, 2024. URL <https://proceedings.mlr.press/v235/shen24d.html>.
 - [323] R. Schaeffer, B. Miranda, and S. Koyejo. Are emergent abilities of large language models a mirage? *arXiv preprint arXiv:2304.15004*, 2023. URL <https://arxiv.org/abs/2304.15004>.
 - [324] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, et al. Scaling Instruction-Finetuned Language Models. *arXiv preprint arXiv:2210.11416*, 2022. URL <https://arxiv.org/abs/2210.11416v5>.
 - [325] C.-P. Huang, Y. Man, Z. Yu, M.-H. Chen, J. Kautz, Y.-C. F. Wang, and F.-E. Yang. Fast-ThinkAct: Efficient Vision-Language-Action Reasoning via Verbalizable Latent Planning, 2026. URL <https://arxiv.org/abs/2601.09708v2>.
 - [326] Y. Duan, Y. Liu, Z. Tang, H. Chen, J. Zhou, Y. Liu, B. Xu, Y. Wu, S. Chen, Y. Zhou, H. Wang, X. Yu, S. Han, X. Zhu, L. Zhou, B. Lu, W. Zhou, J. Liu, N. Fang, J. Tian, R. Chen, Y. Li, K. Zuo, K. Zhang, J. Qiu, C. He, G. Li, B. Zhou, Z. Liu, Z. Wen, J. Kang, X. Zhou, and F. Wu. The Last AI Built by Humans: Toward Genuine Recursive Self-Improvement, 2026. URL <https://arxiv.org/abs/2609.11873v1>.
 - [327] C. Luo, Z. Zeng, M. Jia, Y. Du, and C. Sun. Self-Improving Loops for Visual Robotic Planning, 2025. URL <https://arxiv.org/abs/2506.06658>.
 - [328] Y. Guo, T. Lee, L. X. Shi, J. Chen, P. Liang, and C. Finn. VLAW: Iterative Co-Improvement of Vision-Language-Action Policy and World Model. *arXiv preprint arXiv:2602.12063*, 2026. URL <https://arxiv.org/abs/2602.12063>.
 - [329] X. Liu, Z. Bai, H. Ci, K. Y. Ma, and M. Z. Shou. World-VLA-Loop: Closed-Loop Learning of Video World Model and VLA Policy. *arXiv preprint arXiv:2602.06508*, 2026. URL <https://arxiv.org/abs/2602.06508>.
 - [330] H. Bi, Z. Zhou, Y. Tang, J. Pang, S. Huang, H. Liu, R. Wang, S. Huang, Y. Wang, Y. Cheng, R. Zhao, Z. Li, H. Tan, X. Liu, J. Wan, J. Liu, M. Zhao, F. Bao, and J. Zhu. Motus2: A Self-Evolving General World Model for Dexterous Manipulation, 2026. URL <https://arxiv.org/abs/2608.30237>.
 - [331] A. Zala, J. Cho, H. Lin, J. Yoon, and M. Bansal. EnvGen: Generating and Adapting Environments via LLMs for Training Embodied Agents, 2024. URL <https://arxiv.org/abs/2403.12014>.
 - [332] M. Faldor, J. Zhang, A. Cully, and J. Clune. OMNI-EPIC: Open-endedness via Models of human Notions of Interestingness with Environments Programmed in Code, 2024. URL <https://arxiv.org/abs/2405.15568>.
 - [333] A. Wagenmaker, Z. Zhou, and S. Levine. Behavioral Exploration: Learning to Explore via In-Context Adaptation. *arXiv preprint arXiv:2507.09041*, 2025. URL <https://arxiv.org/abs/2507.09041>.
 - [334] L. Zintgraf, K. Shiarlis, M. Igl, S. Schulze, Y. Gal, K. Hofmann, and S. Whiteson. VariBAD: A Very Good Method for Bayes-Adaptive Deep RL via Meta-Learning, 2019. URL <https://arxiv.org/abs/1910.08348>.
 - [335] X. Li, T. Lai, X. Huang, R. Pang, S. Shen, J. Chen, Z. Pan, C. Xu, F. Gao, and Y. Cao. TADreamer: Zero-Shot Language-Guided 3D Navigation for Terrestrial-Aerial Bimodal Robots via Video Imagination. *arXiv preprint arXiv:2609.19824*, 2026. URL <https://arxiv.org/abs/2609.19824>.
 - [336] R. Wang, H.-L. Hsu, S. Mehta, J. Kim, Z. Dou, and M. Pajic. GAVEL: Graph World Models for Verified and Efficient Long-Horizon LLM Task Planning. *arXiv preprint arXiv:2609.19315*, 2026. URL <https://arxiv.org/abs/2609.19315>.
 - [337] J. Huang, Y. Hu, Z. Li, R. Qi, Y. Xiao, Y. Wu, T. Ba, Z. Zhang, and Y. Zhang. StageGuard: Learning Stage Transitions for Long-Horizon Robot Tasks via Agentic Distillation. *arXiv preprint arXiv:2609.20791*, 2026. URL <https://arxiv.org/abs/2609.20791>.
 - [338] W. Xie, M. Hong, Y. Wang, Y. Ji, Y. Zhou, S. Wu, S. Sirejiding, H. Zhou, Y.-C. Chen, M. Ling, Y. Ding, and H. Lu. Navi-Agent: Unlocalized Monocular Navigation Agent. *arXiv preprint arXiv:2609.20388*, 2026. URL <https://arxiv.org/abs/2609.20388>.
 - [339] W. Li, Y. Chen, W. Li, and Q. Wu. LIFD: Anchored Diffusion for 3D-Aware Scene Memory in Robotic Manipulation. *arXiv preprint arXiv:2609.19796*, 2026. URL <https://arxiv.org/abs/2609.19796>.
 - [340] K. Agrawal, M. A. Rahman, R. A. Yeh, and Z. Kingston. SkipVLA: Skipping VLA Steps with Classical Planning for Fast Robot Manipulation. *arXiv preprint arXiv:2609.20648*, 2026. URL <https://arxiv.org/abs/2609.20648>.
 - [341] S. Kuroki and Y. Tang. Learning and Transferring Closed-Loop Robot Software. *arXiv preprint arXiv:2609.19906*, 2026. URL <https://arxiv.org/abs/2609.19906>.
 - [342] Z. Han, Y. Zeng, J. Zhang, Z. Zhao, Y. Wang, Y. Jin, S. Li, S. Chen, W. Huang, R. Li, H. Shen, and H. Dong. HIL-UMI: Bringing Human-in-the-Loop Post-Training of Vision-Language-Action Models to Universal Manipulation Interface. *arXiv preprint arXiv:2609.20659*, 2026. URL <https://arxiv.org/abs/2609.20659>.
 - [343] Z. Li. CoreSense: Traceable Failure Recall and Conflict-Aware Belief Gating for Auditable Robot Decisions. *arXiv preprint arXiv:2609.19512*, 2026. URL <https://arxiv.org/abs/2609.19512>.

- [344] Z. Zhang, J. Jin, Y. Wang, Z. Zhang, H. Diao, L. Wang, and H. Lu. VABench: Measuring Embodied Spatial Intelligence through Visual Demonstrations, Active Perception, and Metric Control. *arXiv preprint arXiv:2609.19554*, 2026. URL <https://arxiv.org/abs/2609.19554>.
- [345] H. Liu, Q. Li, Y. Ding, J. Wen, K. Ren, Y. Ma, T. Jian, Z. Wang, Z. Yu, X. Li, and Z. Song. REVOLVE: An Automated Closed-Loop Framework for Evolving Robot Manipulation with Minimal Human Intervention. *arXiv preprint arXiv:2609.14633*, 2026. URL <https://arxiv.org/abs/2609.14633>.
- [346] A. Simeonov, Y. Du, Y.-C. Lin, A. Rodriguez Garcia, L. P. Kaelbling, T. Lozano-Perez, and P. Agrawal. SE(3)-Equivariant Relational Rearrangement with Neural Descriptor Fields. In *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 835–846, 2023. URL <https://proceedings.mlr.press/v205/simeonov23a.html>.
- [347] S. Thompson, O. Biza, and G. Konidaris. One-Shot Cross-Geometry Skill Transfer through Part Decomposition, 2026. URL <https://arxiv.org/abs/2604.15455>.
- [348] Noematrix Team. RoboRSI: Stable, Efficient, and Reusable Robot Self-Evolution in Complex Real-World Environments. Research Blog, September 2026. URL <https://lab.noematrix.ai/blog/2-roborsi/>.
- [349] Y. Chen, Z. Huai, S. Li, Y. Wang, H. Zhang, Y. Zhang, H. Chen, J. Gong, Y.-G. Jiang, and X. Qiu. ETA: A New Agentic Paradigm for Embodied Tasks, 2026. URL <https://arxiv.org/abs/2608.03924>.
- [350] J. Huang, Y. Hu, Z. Li, R. Qi, Y. Xiao, Z. Zhang, M. Coates, T. Cao, and Y. Zhang. RoboHarness: Memory-Driven Orchestration of Heterogeneous Robot Policies for Long-Horizon Planning, 2026. URL <https://arxiv.org/abs/2607.18060>.
- [351] B. Zhou, Z. Zhang, J. Wang, and Z. Lu. NOLO: Navigate Only Look Once. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 17162–17169. IEEE, October 2025. doi: 10.1109/iros60139.2025.11246428. URL <http://dx.doi.org/10.1109/IROS60139.2025.11246428>.
- [352] D. Buoso, L. Robinson, G. Averta, P. Torr, T. Franzmeyer, and D. De Martini. Select2Plan: Training-Free ICL-Based Planning Through VQA and Memory Retrieval. *IEEE Robotics and Automation Letters*, 10(11): 11267–11274, November 2025. ISSN 2377-3774. doi: 10.1109/lra.2025.3606790. URL <http://dx.doi.org/10.1109/LRA.2025.3606790>.
- [353] M. Elnoor, K. Weerakoon, G. Seneviratne, R. Xian, T. Guan, M. K. M Jaffar, V. Rajagopal, and D. Manocha. VLM-GroNav: Robot Navigation Using Physically Grounded Vision-Language Models in Outdoor Environments. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2391–2398. IEEE, May 2025. doi: 10.1109/icra55743.2025.11128264. URL <http://dx.doi.org/10.1109/ICRA55743.2025.11128264>.
- [354] P. Anderson, A. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, and A. R. Zamir. On Evaluation of Embodied Navigation Agents, 2018. URL <https://arxiv.org/abs/1807.06757>.
- [355] D. Pathak, P. Mahmoudieh, G. Luo, P. Agrawal, D. Chen, Y. Shentu, E. Shelhamer, J. Malik, A. A. Efros, and T. Darrell. Zero-Shot Visual Imitation. In *ICLR*, 2018. URL <https://arxiv.org/abs/1804.08606>.
- [356] A. Kumar, S. Gupta, D. Fouhey, S. Levine, and J. Malik. Visual Memory for Robust Path Following. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/66368270ffd51418ec58bd793fd2d9b1b-Paper.pdf.
- [357] H. Yoo, N. Kim, J. Park, and S. Oh. Path-Following Navigation Network Using Sparse Visual Memory. In *2020 20th International Conference on Control, Automation and Systems (ICCAS)*, pages 883–886. IEEE, October 2020. doi: 10.23919/iccas50221.2020.9268247. URL <http://dx.doi.org/10.23919/ICCAS50221.2020.9268247>.
- [358] N. Savinov, A. Dosovitskiy, and V. Koltun. Semi-parametric topological memory for navigation. In *International Conference on Learning Representations (ICLR)*, 2018. URL <https://arxiv.org/abs/1803.00653>.
- [359] D. Shah, B. Osinski, B. Ichter, and S. Levine. LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision, and Action. In K. Liu, D. Kulic, and J. Ichnowski, editors, *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pages 492–504. PMLR, 14–18 Dec 2023. URL <https://proceedings.mlr.press/v205/shah23b.html>.
- [360] P. Chen, X. Sun, H. Zhi, R. Zeng, T. H. Li, G. Liu, M. Tan, and C. Gan. A²Nav: Action-Aware Zero-Shot Robot Navigation by Exploiting Vision-and-Language Ability of Foundation Models, 2023. URL <https://arxiv.org/abs/2308.07997>.
- [361] H. Li, X. Dong, H. Jiang, Y. Zhou, and X. Ma. CMMR-VLN: Vision-and-Language Navigation via Continual Multimodal Memory Retrieval, 2026. URL <https://arxiv.org/abs/2603.07997>.
- [362] A. Liu, B. Liu, H. Ding, Y. Jiang, Y. Chen, F. Tang, C. Leng, H. Zhang, and J. Cheng. HAM-VLN: Harnessing Hierarchical Agentic Memory for Zero-Shot Vision-and-Language Navigation, 2026. URL <https://arxiv.org/abs/2607.29600>.
- [363] B. Huang, X. Ding, F. Chen, K. Li, W. Sun, H. Wu, Y. Liu, and T. Cao. Zeva-Ego: Egocentric Mid-Training with In-Context Causal Learning for Robot Manipulation, 2026. URL <https://arxiv.org/abs/2609.24411>.

- [364] M. Wu, K. Yao, S. Luan, M. Tomizuka, and Y. Chen. RopeFormer: Cross-Trial Adaptation from Interaction History for Dynamic Rope Manipulation, 2026. URL <https://arxiv.org/abs/2609.23432>.
- [365] M.-H. Guo, Z.-H. Mo, J.-J. Wang, Y. Zhang, K. Wang, Y.-X. Deng, J.-P. Zhang, Y. Rao, and S.-M. Hu. Transferring the Intelligence of VLMs to Robotic Control, 2026. URL <https://arxiv.org/abs/2609.22966>.
- [366] C. Xiao, H. Zhan, S. Krishna, P. Meng, M. Z. Irshad, S. Zakharov, and D. Held. H2RBench: A Real-to-Sim Benchmark for Evaluating Human-to-Robot Transfer, 2026. URL <https://arxiv.org/abs/2609.24778>.
- [367] K. Soni, A. Sehgal, V. Dave, and E. Rueckert. Sandwich-Residuals: Parameter-Efficient Test-time Adaptation of World Models, 2026. URL <https://arxiv.org/abs/2609.21740>.
- [368] D. Lee, S. Golestaneh, J. Singh, Z. Zhong, A. Kapoutsis, and C. Chamzas. MetaPusher: Meta Learning and Planning for Nonprehensile Manipulation of Unseen Objects with Rapid Online Adaption, 2026. URL <https://arxiv.org/abs/2609.21122>.
- [369] Z. Gao, Y. Zhan, M. Khoshnazar, J. Schäfer, and M. Beetz. KnowDemo: Knowledge-Guided Robot Demonstration Generation from Human Videos, 2026. URL <https://arxiv.org/abs/2609.21229>.
- [370] G. Jena, K. Wanaskar, V. Jain, A. Chadha, V. Sharma, and A. Das. Weights or Skills? A Survey of Robot-Learning Techniques: from Action-Predicting Weights to Robots that Write their Own Skills, 2026. URL <https://arxiv.org/abs/2608.01851>.
- [371] Z. Lu, H. Zhai, G. Wang, H. Zeng, J. Yang, J. Liu, L. Cheng, Y. Zhang, Y. Qiu, Z. Cheng, I. Laptev, D. Xu, B. Riviere, G. Loianno, E. Xing, and X. Zuo. World-action models for robot learning and control: A survey, 2026. URL <https://arxiv.org/abs/2609.16074>.
- [372] W. Li. In-Context Learning from Demonstrations for Robotic Manipulation: A Survey. Preprints.org, manuscript 202609.0780, version 1, September 2026. URL <https://www.preprints.org/manuscript/202609.0780>.
- [373] R. Shao, W. Li, L. Zhang, R. Zhang, Z. Liu, R. Chen, and L. Nie. Large VLM-based vision-language-action models for robotic manipulation: A survey, 2025. URL <https://arxiv.org/abs/2508.13073>.
- [374] P. Zhao, L. Wang, M. Wang, C. Chen, F. Zhou, and H. Huang. EvoEmpirBench: Dynamic Spatial Reasoning with Agent-ExpVer. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(43):36564–36572, 2026. doi: 10.1609/aaai.v40i43.40979. URL <https://ojs.aaai.org/index.php/AAAI/article/view/40979>.
- [375] Y. Zhang, J. Su, H. Huang, Y. Chang, T. Zhou, X. Xu, Y. Xu, Y. Li, Z. Li, and Y.-C. Chen. WorldLines: Benchmarking and Modeling Long-Horizon Stateful Embodied Agents. *arXiv preprint arXiv:2606.18847*, 2026. URL <https://arxiv.org/abs/2606.18847>.
- [376] H. Huang, J. Shi, Y. Li, and Y. Chen. Affordance Agent Harness: Verification-Gated Skill Orchestration. *arXiv preprint arXiv:2605.00663*, 2026. URL <https://arxiv.org/abs/2605.00663>.
- [377] L. Wu, Y. Zhao, J. Zhang, S. Chen, W. Chen, Z. Chen, T. Xu, Q. He, H. Hu, H. Huang, Y. Wei, W. Li, Y. Li, and Y.-C. Chen. RoboStressBench: Benchmarking VLM Robustness to Physical Visual Stress in Embodied Scenes. *arXiv preprint arXiv:2606.00828*, 2026. URL <https://arxiv.org/abs/2606.00828>.
- [378] H. Huang, H. H. Chen, M. Luo, J. Du, S. Xu, Z. Chen, Y. Huang, Y. Li, and Y.-C. Chen. No Place to Hide: Benchmarking Video Hallucination with Background-Controlled Pairs. *arXiv preprint arXiv:2606.31933*, 2026. URL <https://arxiv.org/abs/2606.31933>.
- [379] Z. Li, Y. Zhang, W. Li, H. Huang, C. Wang, S. Deng, Y. Wei, T. Zhang, B. Xie, B. Zhou, Y. Chang, K. Zhou, Y.-C. Chen, J. Cheng, and Y. Li. Robo-Harness K1: Harnessing Robot-Use Agents via Perception Augmentation, 2026. URL <https://arxiv.org/abs/2609.29389>.
- [380] Z. Li, Y. Zhang, H. Huang, B. Zhou, W. Li, C. Wang, Y. Chang, Y. Wei, T. Zhang, Y.-C. Chen, K. Zhou, Y. Li, and J. Cheng. RACaP: Agentic Reasoning, Acting, and Coding as Policies for Evolvable Robot Learning, 2026. URL <https://arxiv.org/abs/2609.29394>.
- [381] T. D. Pala, N. Majumder, B. Goh, R. Yee, J. Yang, L. Chen, and S. Poria. MemBodied: Recurrent Associative Memory for Vision-Language-Action Models, 2026. URL <https://arxiv.org/abs/2609.28256>.
- [382] H. You, Z. Shen, Y. Liu, Z. Zheng, L. Zha, K. Yamazaki, M. Zhang, S. Huang, J. Sun, Q. Chen, L. He, K. Liu, H. Chang, K. Fragkiadaki, D. Shah, M. Schwager, P. Henderson, I. Abraham, and C. Xu. EmbodiedSWE: Coding Agents for Long Horizon Dexterous Robotics, 2026. URL <https://arxiv.org/abs/2609.27308>.
- [383] J. Liu, S. Yao, J. Gu, and L. Xu. NavProbe: Evidence-Grounded Reasoning with Active Memory Retrieval for Zero-Shot Navigation, 2026. URL <https://arxiv.org/abs/2609.27526>.
- [384] Z. Li, S. Lin, Y. Shao, J. Zhang, X. Shi, S. Li, and Q. Wu. Talk2Escape: Conversational Grounding for Vision-and-Language Navigation, 2026. URL <https://arxiv.org/abs/2609.28296>.
- [385] J. Zhang, Y. Xiang, and C. Chen. Memory That Changes Action Is Not Memory That Guides It: Counterfactual Auditing of History-Conditioned Robot Policies, 2026. URL <https://arxiv.org/abs/2609.27247>.
- [386] Y. Wang, G. Zhang, and J. Wang. Task-Prototype Guided Flow Matching for Few-Shot Generalization in Vision-Language Robot Manipulation, 2026. URL <https://arxiv.org/abs/2609.27780>.
- [387] D. Yi, P. Sun, C. Rong, J. Wang, X. Dai, X. Yue, and X. Lin. Watch, Recall, Act: Always-On Robots in Concurrent Embodied Streams, 2026. URL <https://arxiv.org/abs/2609.28429>.
- [388] Y. Zhang, H. Huang, Y. Chang, J. Su, B. Zhou, Y. Xu, W. Chen, T. Zhou, C. Wang, T. Zhang, Y. Wei, W. Li, S. Deng, Y. Li, Y.-C. Chen, and Z. Li. World Action Agent: Harnessing VLMs for Robot Manipulation via

- World Action Rehearsal, 2026. URL <https://arxiv.org/abs/2609.29964>.
- [389] Knowin AI. GLOW: A generative learning framework for general-purpose embodied intelligence. Online technical report, September 2026. URL <https://knowinai.com/tech.html>. Accessed September 26, 2026.
 - [390] Y. Liu, J. Mao, D. Hsu, L. P. Kaelbling, and T. Lozano-Pérez. Rapid: Robot agentic programming from demonstrations, 2026. URL <https://arxiv.org/abs/2609.30249>.
 - [391] M. Adang, T. Chen, L. Osterberg, A. Swann, and M. Schwager. Training-free behavior cloning, 2026. URL <https://arxiv.org/abs/2609.30134>.
 - [392] J. Zhang, R. Liu, Y. Zhang, and Y. Yang. Traceflow: Guiding frozen flow-matching robot policies with success and failure traces, 2026. URL <https://arxiv.org/abs/2609.20646>.
 - [393] H. Zhang, C. Lin, T.-H. Wang, Z. Xu, and C. Gan. Self-adaptive vla for robust robot deployment, 2026. URL <https://arxiv.org/abs/2609.30092>.
 - [394] H. Xu, Z. He, H. Wang, J. Xu, and H. Dong. Mimir: A neuro-symbolic memory system with dynamic grounding for embodied agents in interactive environments, 2026. URL <https://arxiv.org/abs/2608.04933>.
 - [395] J. B. Jedlicki, T. Dieudonné, and H. Yang. Occ4m: Object-centric 4d memory for spatiotemporal reasoning in long-horizon manipulation, 2026. URL <https://arxiv.org/abs/2609.28798>.
 - [396] Q. Chen, J. Kang, R. Lin, Z. Zhang, E. Feng, C. Zheng, X. Dong, and J. Lin. Sparsenav: Instruction-conditioned sparse semantic perception for training-free vision-language navigation, 2026. URL <https://arxiv.org/abs/2609.26408>.
 - [397] N. Dashora, D. Chen, I. Shenfeld, J. Marangola, P. Agrawal, and M. Simchowitz. Workspace models: Lightweight robotic memory via saliency-driven supervision, 2026. URL <https://arxiv.org/abs/2609.20820>.
 - [398] X. Wang, W. Wu, M. Zhang, Z. Wang, K. Shao, J. Luan, Y. Li, Q. Li, S. Gu, H. Zhou, S. Shi, F. Ni, S. Lu, W. Meng, K. Li, J. Wu, K. Zhao, S. Guo, G. Li, Y. Tang, Z. Zhang, Y. Xie, and H. Qu. Harnesspai: An evolving harness for physical ai, 2026. URL <https://arxiv.org/abs/2609.29166>.
 - [399] C. Guo, Y. Xie, B. Tan, Z. Chang, Z. Yin, Q. Ma, Y. Wang, C. Liang, and Z. Zhang. Robofollow: Unveiling the instruction following mirage in embodied agents, 2026. URL <https://arxiv.org/abs/2609.25636>.
 - [400] Y. Li, C. Zhao, Z. Wang, J. Wang, C. Shao, Y. Lin, H. Shen, and J. Zhang. Roborecover: Benchmarking robot policy recovery under execution deviations, 2026. URL <https://arxiv.org/abs/2609.28952>.
 - [401] S. Zhao, H. Xie, W. Zhao, C. Zhang, H. Wang, C. Wang, Q. Liu, and S. Zhang. Memory as plans: World-action modeling with memory-grounded planning, 2026. URL <https://arxiv.org/abs/2609.11561>.
 - [402] L. Li, L. Chen, Kwunhang Wong, J. Lei, S. Jin, S. Du, C. Zhang, S. Ma, W. Zhang, J. Xiao, Kwang-Ting, and Cheng. Causal-history test-time scaling for failure recovery in autoregressive world-action models, 2026. URL <https://arxiv.org/abs/2609.18016>.
 - [403] Y. Hu, F. Chen, X. Cao, and R. Detry. 2am: Grounding agent-side memory as guidance for steerable action models in long-horizon manipulation, 2026. URL <https://arxiv.org/abs/2609.11308>.
 - [404] R. Li, W. He, C. Lu, X. Chen, and X.-Y. Zhang. HIRE: History-Conditioned Interaction Reasoning and High-Rate Execution for Visually Aliased Precision Manipulation, 2026. URL <https://arxiv.org/abs/2609.30828>.
 - [405] S. M. Kim, J. Seo, H. Heo, J. Lee, Y. Lee, and Y. M. Kim. POIL: Point-based One-Shot Imitation Learning with Stable Dynamical Systems, 2026. URL <https://arxiv.org/abs/2609.30404>.
 - [406] Z. Ravichandran, J. Diller, F. Cladera, V. Murali, G. J. Pappas, and V. Kumar. Actively Resolving Contextual Uncertainty for Underspecified Tasks in Natural Language, 2026. URL <https://arxiv.org/abs/2609.30428>.
 - [407] M. Merler, B. Li, J. Roy, Y. Liang, Q. Wang, Y. Huang, and T. Silver. Coding Agents for Generalized Task and Motion Planning Problems, 2026. URL <https://arxiv.org/abs/2609.30233>.
 - [408] R. Yang, M. Yu, and C. Fan. Representation-Guided Generation and Integration of Executable Programs for Robot Manipulation, 2026. URL <https://arxiv.org/abs/2609.31337>.
 - [409] S. Choi, S. Ye, N. Bui, A. Shrivastava, K. Ryu, D. Tirumala, M. Wulfmeier, and N. Mehr. HuGo: LLMs as Whole-Body Policy Code Designers for Humanoid Loco-Manipulation, 2026. URL <https://arxiv.org/abs/2609.30594>.
 - [410] C. Li, Y. Di, Y. Li, R. Zhang, M. Li, and X. Ji. DualManip: Agentic Dynamic Manipulation via Dual-Path Semantic Reasoning and Geometric Adaptation, 2026. URL <https://arxiv.org/abs/2609.31112>.
 - [411] J. Siburian, C. C. Beltran-Hernandez, T. Matsushima, Y. Iwasawa, M. Hamaya, and M. Nishimura. PHASE: Compliance-Enabled Tactile Phase Retrieval for Few-Shot Insertion Learning, 2026. URL <https://arxiv.org/abs/2609.30889>.
 - [412] M. Abouagour and B.-C. Min. Audit Before You Commit: Locating Belief Failures in Active Identification for One-Shot Manipulation, 2026. URL <https://arxiv.org/abs/2609.30608>.

A Complete Reference Index

Table A1 links every reference to Figure 9. Full titles, dates, primary sources, and classification details accompany the reference data.

Table A1 Reference index for Figure 9, grouped by first public year. Method families: C, context-conditioned policies; G, geometric transfer; W, world-model-based control; S, skill/agent execution. Other groups record data, evaluation, foundations, and training references.

Year	Family	Count	References
2016	C	1	[2]
2017	C	4	[1, 3, 7, 146]
	S	1	[26]
2018	C	4	[8, 157, 355, 356]
	W	2	[179, 294]
	S	1	[358]
2019	C	4	[4, 147, 260, 334]
2020	C	5	[158, 196, 297, 300, 357]
	G	1	[32]
2021	C	3	[12, 39, 149]
2022	C	6	[5, 6, 100, 148, 197, 276]
	G	2	[33, 346]
	W	1	[19]
	S	5	[27, 28, 185, 186, 359]
2023	C	6	[57, 114, 202, 243, 274, 277]
	G	4	[34, 169–171]
	W	1	[20]
	S	9	[29, 46, 84, 86, 96, 164, 184, 251, 360]
2024	C	16	[13–16, 80, 118, 119, 132, 153, 154, 189, 203, 275, 278, 279, 351]
	G	4	[35, 156, 173, 295]
	S	11	[50, 76, 87, 117, 165, 213–215, 298, 352, 353]
2025	C	18	[17, 40, 42, 43, 61, 120, 155, 159, 190, 193, 201, 235, 245, 273, 280, 282, 315, 333]
	G	8	[36, 37, 93, 168, 172, 174, 177, 191]
	W	2	[21, 327]
	S	5	[9, 30, 51, 293, 305]
2026	C	56	[10, 18, 41, 44, 45, 54, 55, 77, 88, 91, 92, 97–99, 105, 160–163, 176, 181–183, 195] [198–200, 204, 205, 208, 233, 244, 247, 270, 281, 283, 292, 299, 317, 325, 328, 329, 339, 340, 342, 363, 364, 381] [386, 387, 389, 391–393, 397, 404]
	G	5	[38, 133, 175, 347, 405]
	W	16	[11, 22–25, 152, 209, 302, 316, 330, 335, 336, 367, 368, 401, 402]
	S	59	[31, 47–49, 52, 53, 89, 90, 94, 142, 166, 187, 206, 207, 210–212, 216–220, 262, 272] [290, 291, 303, 304, 307, 337, 338, 341, 343–345, 348– 350, 361, 362, 365, 376, 379, 380, 382–384, 388] [390, 394–396, 398, 403, 406–410]
Data resources			
		44	[56, 58–60, 62–75, 221–226] [237–242, 246, 248–250, 255–259, 261, 263, 331, 332, 369]
Evaluation			
		39	[81–83, 85, 95, 101–104, 106–110, 151, 227–230, 284, 296, 306, 308, 309] [310–314, 354, 366, 374, 375, 377, 378, 385, 399, 400, 412]
Foundations and surveys			
		40	[111–113, 121–130, 134–138, 143–145, 178, 180, 194] [231, 232, 234, 253, 301, 318–323, 326, 370–373]
Training and representations			
		29	[78, 79, 115, 116, 131, 139–141, 150, 167, 188, 192, 236, 252, 254, 264–269, 271, 285, 286] [287–289, 324, 411]

Table A1 Complete reference assignments (continued).

Year	Family	Count	References
Total		412	

B Method and Resource Comparisons

The following tables provide per-study detail for the main-text classifications. They follow the progression from demonstration correspondence to retained experience and the data and comparisons needed to evaluate it.

Table B1 Correspondence mechanisms compared by the entities they relate and their use in control. The same matched points can condition action inference or specify a reference for a separate executor.

Method	Correspondence	Control interface
Correspondence conditions or corrects action generation		
MatchingPolicy [161]	Demonstration–query 3D points	Graph-conditioned gripper motion
RTCF [195]	Current and recorded progress	Bounded correction
Correspondence specifies an executable motion or constraint		
RoboTAP [169]	Stage-specific points	Visual-servo reference tracking
FUNCTO [172]	Grasp; function; center	Function-preserving tool motion
MAGIC [173]	Shape; contact geometry	Retarget; simulate
SemAnCorr [38]	Surfaces; local frames	Grasp; relative motion
Sparse Meets Dense [175]	Changing contact pairs	Constraint; pose updates

The same matched geometry can have different computational roles. Table B1 separates correspondence that conditions action inference from correspondence that supplies an executable motion reference.

Table B2 Memory mechanisms compared by representation, access, and update regime. Groups (a)–(d) retain evidence with fixed neural parameters; group (e) adapts weights. A fixed test bank supports retrieval without adding records.

Method	Retained evidence	Access and update
(a) Demonstration records		
KAT [15]	Keypoint–action examples	Load before rollout
ICRT [13]	Demonstrations; own trajectory	Read prefix; append execution
RICL [17]	Retrieved demonstrations	Query-dependent selection
ReCAP [152]	Trajectory retrieval pool	Extend pool; retrieve each step
RTCF [195]	Successful trajectories	Phase retrieval; fixed test bank
TraceFlow [392]	Success and failure traces	Progress-aligned guidance; rollout admission
(b) Interaction histories and inferred state		
RMA [39]	Recent state–action history	Online dynamics estimation
AMAGO [202]	Trial history and feedback	Append across episodes
ReLIC (navigation) [203]	Multi-episode visual history	Attend across trials
LocoFormer [40]	Recurrent rollout state	Carry across trials
ICWM [41]	Action–effect probes	Probe before execution
(c) Execution memory		
MemoryVLA [42]	Perceptual; cognitive tokens	Retrieve; merge
MemER [43]	Historical visual keyframes	Retrieve; predict subtask
MEM [44]	Recent views; semantic history	Recent and long-term

Table B2 Memory mechanisms (continued).

Method	Retained evidence	Access and update
ReMem-VLA [198]	Recurrent memory queries	Frame and chunk updates
MemoryVLA++ [209]	Memory bank; imagined future	Retrieve; predict
μ VLA [199]	Recurrent memory tokens	Carry state; reset per episode
EventVLA [200]	Anchors; event keyframes	Event-triggered writes
SimpleMemVLA [45]	Timestamped video window	Resample per decision
AGM [206]	Verified subgoal achievements	Completion-verified update
Workspace Models [397]	Latent history tokens	Saliency-supervised compression
Mimir; OCC4M [394, 395]	Task state; persistent entities	Goal binding; feedback-supported updates
2AM [403]	Agent-side interaction history	Compile current intent into 2D cues
(d) Reusable guidance		
DROC [46]	Distilled correction knowledge	Feedback update; retrieval
(e) Parameter adaptation through fast-weight memory		
RoboTTT [10]	Video and rollout history	Sequential weight updates
WAM-TTT [11]	Human demonstration video	Adapt; freeze

Table B2 compares what each system retains, when it accesses that evidence, and how the retained state changes. The comparison links decision-relevant information to its storage and update mechanism.

Table B3 Artifacts revised under fixed neural weights. Bullets mark updates; dashes mark undocumented updates for the specified branch. Plans specify procedures, guidance stores lessons, and skills are callable routines. Programs include controller settings; oversight includes context construction and verification. ZS denotes the zero-shot video branch (Section 3.6).

System and branch	Plan	Guidance	Skills	Program	Oversight
(a) Constructing and revising the current procedure					
Inner Monologue [186]	•	—	—	—	—
Code as Policies [27]	—	—	•	•	—
ProgPrompt [28]	—	—	—	•	—
VoxPoser (fixed dynamics) [184]	•	—	—	•	—
PRoC3S [214]	—	—	—	•	—
MALMM [213]	•	—	—	•	—
Show-Harness (video, ZS) [31]	•	—	—	—	—
GPT-Policy [94]	—	—	—	•	—
ICPI [9]	—	—	—	•	—
AOR [218]	—	—	—	•	—
ENPIRE (code branch) [89]	—	—	—	•	—
(b) Acquiring examples, guidance, and reusable behavior					
DROC [46]	•	•	—	•	—
ICAL (TEACH, retrieval) [215]	•	•	—	•	—
LRL [50]	—	—	•	•	•
LYRA [51]	—	—	•	•	—
Uni-Skill (deployment) [217]	—	•	•	•	—
CaP-Agent0 (CaP-X) [166]	—	—	•	•	—
RATs [52]	•	•	•	•	—
RoboRSI (code skills) [348]	•	—	•	•	—
ASPIRE [47]	—	•	—	•	—

Table B3 Artifacts revised under fixed neural weights (continued).

System and branch	Plan	Guidance	Skills	Program	Oversight
ARCHITECT [48]	–	•	–	•	–
Teach and Grow (pilot) [53]	–	–	•	–	–
Harness VLA (memory) [207]	•	•	–	–	–
Hi-VLA (episode summary) [219]	–	•	–	–	–
(c) Revising context construction and execution oversight					
SHAPER [49]	–	•	–	–	•
Zetta [90]	–	•	•	–	•

Table B3 distinguishes revisions to procedures, retained guidance, and callable skills. These artifacts can change together: RoboRSI retains both revised plans and executable routines.

Table B4 Data resources compared by demonstrator, motion supervision, and role in contextual control. Motion may be measured, estimated, retargeted, simulated, or synthesized; human $\rightarrow$ robot denotes robot imagery synthesized from human footage. The final column identifies potential uses of the available records in contextual learning. Est., Sim., and Synth. denote estimated, simulated, and synthesized.

Source	Visual record	Motion supervision	Available learning evidence
Robot and cross-domain recordings			
Open X-Emb. [57]	Robot	Robot control	Robot action targets
DROID [58]	Robot	Robot control	Robot action targets
AgiBot World [60]	Robot	Robot control	Stages; corrections
BridgeData V2 [56]	Robot	Robot control	Robot action targets
RoboSet [243]	Robot	Robot control	Task groups
PrimeBot, XR-2 [244]	Human + robot	Robot; gripper	Separate corpora
RH20T [59]	Human + robot	Robot control	Human-robot pairs
ICRT data [13]	Robot	Robot control	Robot demonstrations
MimicPlay [251]	Human + robot	Robot control	Separate play
Instrumented human recordings			
UMI [62]	Human	Gripper motion	Motion correspondence
YUBI [66]	Human	Gripper motion	Motion correspondence
HiFi-UMI [223]	Human	Gripper motion	Motion correspondence
ACE-Data-0 [250]	Human	Human motion	Contact; procedure
Human activity video			
Ego4D [68]	Human	No robot actions	Behavior pretraining
Ego-Exo4D [70]	Human	No robot actions	Cross-view
EPIC-KITCHENS-100 [69]	Human	No robot actions	Behavior pretraining
Something-Something [71]	Human	No robot actions	Action clips
EgoVerse [248]	Human	Est. hand motion	Behavior pretraining
Open-AoE [249]	Human	Est. hand motion	Action segments
Synthesis			
MimicGen [74]	Robot	Sim. control	Query synthesis
HumanGen [24]	Human + robot	Robot control	Synth. support
Ego2Robot [222]	Human $\rightarrow$ robot	Retargeted motion	Synth. query

Table B4 connects data provenance to action correspondence. Measured robot commands, estimated human motion, and synthesized robot imagery offer different supervision even when they depict the same task.

C Related Survey Perspectives

Table C1 compares related surveys by scope and organizing principle. ICL for robots connects their perspectives through the evidence and interfaces that change execution.

Table C1 Related surveys by scope, organizing axes, and analytical emphasis.

Review focus	Organizing axes	Analytical emphasis
Demonstration and contextual adaptation		
Learning from demonstration (2009) [122]	Demonstrator, problem space, policy derivation	Acquiring examples and deriving policies
Learning from demonstration (2020) [123]	Teaching interface; policy, reward, or plan	Matching teaching and learning to the task
In-context learning [124]	Training, prompting, mechanisms	Examples as language-model context
In-context RL [125]	Pretraining; context; models	History-based adaptation with fixed weights
Nonstationary ICRL [127]	Shift type; timing; evidence	Validity of past evidence after changes
Demonstration-driven manipulation ICL [372]	Context, inference target, adaptation, transfer	How examples and priors support manipulation
Data and embodiment		
Human-video learning [126]	Task, observation, and action transfer	Acquiring skills from video-derived evidence
VLM-based VLA models [373]	Monolithic/hierarchical models	From multimodal features to actions
VLA data infrastructure [129]	Data, benchmarks, data engines	Data and evaluation for robot learning
The embodiment gap [130]	Shared structure; adaptation stage	Reuse and robot-specific requirements
Predictive control and reusable skills		
World models for robot learning [128]	Policy coupling, simulation, representations	Prediction for learning, planning, and evaluation
World Action Models [178]	Cascaded/joint models; conditioning, decoding	Coupling future-state prediction and action generation
World-action learning and control [371]	Transition models, architectures, training	Predictive control across robot applications
Weights or Skills? [370]	Parameters; skills; revision	Capability representation and reuse
ICL for robots (this survey)	Context-to-action interfaces and transfer assumptions	Teaching, physical realization, and experience reuse across tasks