Robustness of Robotic Manipulation: Foundations and Frontiers

Journal Title
XX(X):1–24
©The Author(s) 2025
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI:
www.sagepub.com/
SAGE

Yifei Dong[1,2], Zhanyi Sun[3], Lujie Yang[4], Manuel Baum[5], Kei Ikemura[2], Shuran Song[3], Florian T. Pokorny[2], and Xianyi Cheng[1]

Abstract

Humans and animals exhibit remarkable robustness in physical manipulation, yet robots remain far behind. Progress toward human-level manipulation robustness is hindered by the absence of a unified and systematic understanding: different subfields frame robustness in distinct ways, often leaving the concept ambiguous and limiting deeper analysis as well as communication across research areas. This review paper presents a systematic study of manipulation robustness. We begin with a formal definition, characterizing robustness as the degree to which a manipulation system can achieve its goal in the presence of uncertainty and variation. Building on this definition, we introduce general formulations of manipulation robustness from probabilistic and control-theoretic perspectives. We then synthesize the guiding principles and concrete mechanisms of manipulation robustness across perception, planning, control, policy learning, and hardware, illustrating each mechanism through representative works, including foundational and recent studies. In addition, we review existing metrics and evaluation methods for quantifying manipulation robustness. Finally, we distill broader lessons for designing robust manipulation systems and discuss open problems and future directions toward achieving human-level robustness in robotic manipulation.

Keywords

Manipulation and Grasping, Review, Robustness, Manipulation under Uncertainty

1 Introduction

Manipulation in the real world is inherently uncertain (Mason 2018). Object pose may be partially observed, friction and compliance may be poorly modeled, contact events are discontinuous and difficult to predict, and task instances may vary substantially across episodes. Robust manipulation is therefore not simply a matter of executing a nominal plan accurately in an idealized environment. Instead, it is the ability to *achieve task goals despite such uncertainty and variation*.

Biological systems provide compelling evidence that robust manipulation is possible. Humans often handle these uncertainties and variations without explicit deliberation, relying instead on sensorimotor control strategies that combine prediction with rapid feedback correction (Flanagan et al. 2006). For instance, humans regulate grip forces to prevent objects from slipping in the hand. When grasping objects such as fruits (Fig. 1), humans apply grip forces with a safety margin above the expected slip threshold. Under varying or dynamic conditions, this margin is adjusted accordingly (Hadjiosif and Smith 2015), preventing grasp failure even in uncertain environments.

Robust behaviors are also generally observed in animals. Sea otters, for example, juggle pebbles on their bellies (Foster-Turley and Markowitz 1982). They maintain this behavior despite turbulent disturbances and across wide variations in object shape, mass, and body posture. Even far less complex organisms exhibit robust interaction with their environment. For example, coordinated ciliary motion in paramecium generates fluid flows that draw and capture food particles (Fenchel 1980). Without complex visual perception, the organism can handle uncertainty in particle location using only mechanosensory cilia in highly dynamic fluid environments. These examples suggest that robustness in physical interaction is a fundamental property of biological systems (Kitano 2004).

In contrast, today's robotic systems manipulate objects with a level of robustness still below that of humans and many animals (Mason 2018; Billard and Kragic 2019). Even routine behaviors such as pick-and-place, which humans perform effortlessly, remain challenging for robots to execute reliably in the presence of contact uncertainty (Rodriguez 2021). Although impressive demonstrations from academia and industry have shown remarkable capabilities under controlled conditions, many of these systems degrade significantly when deployed outside carefully engineered environments. Ultimately, robotic manipulation must operate in open, time-varying environments characterized by significant uncertainty.

This evident gap between the robustness of robotic and human manipulation motivates this review. A major

[1]Duke University, USA
[2]KTH Royal Institute of Technology, Stockholm, Sweden
[3]Stanford University, USA
[4]Massachusetts Institute of Technology, USA
[5]Organifarms, Germany

Corresponding author:
Yifei Dong, KTH Royal Institute of Technology, Stockholm, Sweden.
Email: yifeid@kth.se

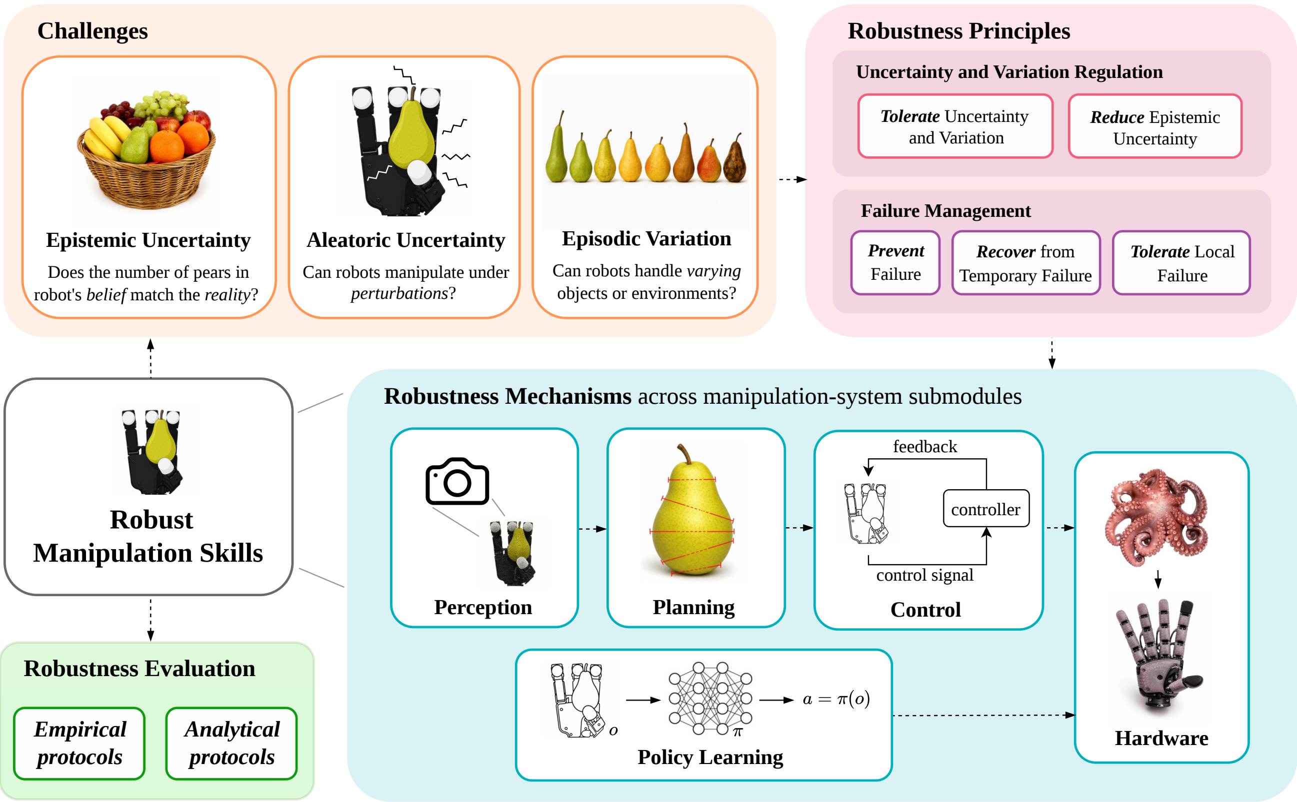

Figure 1. Overview of the survey on the robustness of robotic manipulation, illustrated using a pear grasping task. Robotic manipulation faces three core challenges: epistemic uncertainty, aleatoric uncertainty, and episodic variation. The survey is organized around two robustness principles: *uncertainty and variation regulation*, which concerns reducing epistemic uncertainty or tolerating uncertainty and variation, and *failure management*, which concerns preventing failure, recovering from temporary failure, or tolerating local failure. We discuss robustness mechanisms across five submodules of a manipulation system: perception, planning, control, policy learning, and hardware. Several icons or illustrations are original and conceptually inspired by (Lu et al. 2025; Fruit Hunters 2026; OnlineDelivery.in 2024; Berkeley AI Research 2019; Cao et al. 2022; Sankar et al. 2025).

obstacle to progress is the absence of a deep understanding and overarching framework for manipulation robustness. Each subfield of robotics discusses robustness in its own terms, often restricted to a narrow domain and leaving the concept implicit or underspecified (Moos et al. 2022; Skogestad and Postlethwaite 2005; Ghazi-Zahedi et al. 2017). To address this issue, we propose a shared conceptual language and a taxonomy to connect existing research threads and enable the accumulation and accessibility of insights across subfields. We also identify several foundational but unanswered questions: What explains the extraordinary robustness observed in human and animal manipulation? What mechanisms and principles enable robust manipulation? How should robustness be quantified and measured? This review calls attention to these questions and encourages the field to pursue a deeper and more systematic understanding of robust manipulation skills.

Robust manipulation skills appear not to arise from a single mechanism, but rather from the coordination of mechanisms across perception, planning, control, policy learning, and hardware, calling for joint research across subfields (Billard and Kragic 2019). This perspective distinguishes the present review from prior surveys on robotic manipulation, such as data-driven grasp synthesis (Bohg et al. 2013), robot learning for manipulation (Kroemer et al. 2021), foundation models for robotics (Firoozi et al. 2025), etc. While these works have made significant contributions to broad areas of manipulation, robustness is typically treated only incidentally rather than as the central object of study. To our knowledge, there has not yet been a comprehensive review that systematically examines the principles, mechanisms, and evaluation methods underlying robust manipulation under real-world uncertainty and variation.

This review makes the following *contributions*: (i) it proposes a task-centered definition of manipulation robustness; (ii) it formulates manipulation robustness from both probabilistic and control-theoretic perspectives; (iii) it organizes robustness mechanisms across perception, planning, control, policy learning, and hardware under a set of guiding principles; and (iv) it synthesizes evaluation protocols and open challenges for future research.

In the following, we begin by introducing a definition (Section 2) and a mathematical formulation (Section 3) of manipulation robustness. We then present robustness principles (Section 4), followed by a systematic review of existing robustness mechanisms (Section 5). Next, we analyze current metrics and evaluation protocols (Section 6). We then synthesize key insights and observations (Section 7). Finally, we highlight open challenges and suggest future directions toward achieving human-level manipulation robustness (Section 8).

2 Definition of Manipulation Robustness

This section introduces the concept of robustness progressively. We begin with general definitions of robustness across disciplines, narrow the scope to robotics, and finally formalize the notion of manipulation robustness.

2.1 Robustness Definitions

Despite its frequent use, the notion of robustness is often interpreted differently across contexts and disciplines. Several domain-specific efforts have attempted to formalize robustness in biology (Kitano 2004), machine learning (Braiek and Khomh 2025), and robotics (Baum 2024). We provide representative definitions below for illustration.

Definition 1. Biological robustness. Robustness is a property that allows a system to maintain its functions despite external and internal perturbations (Kitano 2004).

Definition 2. Machine learning robustness. The robustness of machine learning models denotes the capacity of a model to sustain stable predictive performance in the face of variations and changes in the input data (Braiek and Khomh 2025).

Definition 3. Robotic robustness. Robustness specifies the degree to which a behavior can fulfill its task despite being challenged by an adversity (Baum 2024).

Synthesizing these perspectives reveals a common theme: robustness is inherently *context-dependent*. It is not an absolute property of a system, but a relational one. It is defined only with respect to what the system is intended to achieve and the conditions under which it operates. Without context, claims of robustness become ill-posed and difficult to compare across domains or methods.

To make this dependence explicit, we identify four foundational dimensions that together define the context of robustness: *goal*, *challenge*, *mechanism*, and *evaluation*. With these dimensions, we define robustness as follows:

Definition 4. Robustness. Robustness refers to a system's ability to achieve a given *goal* under specified *challenges*, enabled by particular *mechanisms*, and assessed through quantitative *evaluation*.

This structure applies broadly across disciplines. For example, in machine learning, robustness typically refers to maintaining predictive performance (*goal*) (Braiek and Khomh 2025) under distribution shifts or adversarial perturbations (*challenge*) (Madry et al. 2018), achieved through techniques such as adversarial training (*mechanism*), and evaluated using quantitative metrics such as worst-case accuracy, or performance degradation under controlled perturbations (*evaluation*) (Hendrycks and Dietterich 2019).

2.2 Robustness in Robotics

What distinguishes robotics is the embodied physical interactions and tight integration of subsystems. *Goals*: Because robotic systems integrate sensing, mechanics, planning, control, actuation, and learning, robustness must often be addressed both within and across the components. Thus, goals may be defined at different levels of abstraction, ranging from high-level task completion (e.g., object transport or assembly) to component-level objectives such as state estimation accuracy, collision avoidance, or trajectory tracking performance. *Challenges*: Direct interaction with the physical world exposes robots to uncertainty in perception and contact dynamics, sensing noise and external disturbances, as well as variations in system parameters and environmental conditions. *Mechanisms*: Robustness is realized through complementary passive and active strategies spanning hardware and algorithms. Passive strategies, such as compliant hardware and soft materials, physically tolerate disturbances. Active strategies, such as motion planning, feedback control, and learning methods, proactively mitigate model mismatch or compensate for noise or disturbances. *Evaluation*: Robustness in robotics is commonly evaluated on task execution success rates, complemented by more fine-grained metrics such as stability measure, control convergence, or theoretical guarantees, defined at the level of individual tasks or system components.

2.3 Robotic Manipulation Robustness

Manipulation is characterized as "an agent's control of its environment through selective contact" (Mason 2018), which makes manipulation robustness a distinct subclass of robotic robustness, with unique challenges arising from contact interactions among the robot, manipulated objects, and the environment.

Manipulation is dominated by multi-body contact interactions, which induce dynamics that are high-dimensional, hybrid, discontinuous, constrained, and therefore difficult to model accurately. These interaction dynamics constitute a major *challenge*, as small variations in geometry, contact conditions, or material properties can lead to qualitatively different outcomes. Perception introduces additional challenges in manipulation because effective contact interaction depends on estimating interaction-relevant states that are largely external to the robot, such as object pose, insertion depth, local contact geometry, material properties, and deformation. Many of these quantities are only partially observable, visually ambiguous, or occluded during contact. Therefore, it has substantially higher demands on perception than tasks such as locomotion, where the robot primarily regulates its own state. Together, these factors complicate the design of robust *mechanisms* in manipulation planning, control, and learning.

Moreover, real-world manipulation encompasses a broad spectrum of task specifications, in which robots must achieve diverse objectives across domestic, industrial, and open-world environments. Unlike locomotion, manipulation relies on multi-body contact interactions among the robot, manipulated objects, and the environment that actively change the state of the world, leading to a combinatorial diversity of manipulation *goals*. Finally, the discontinuous and task-dependent nature of contact interactions makes the definition of general *evaluation* criteria and the establishment of formal robustness guarantees particularly difficult.

We refine Definition 4 for manipulation by specifying the four foundational dimensions as follows:

- *Goals*: Robotic manipulation robustness is defined relative to manipulation goals, which involve achieving desired object states through contact interactions, potentially within task-specific tolerance levels.
- *Challenges*: Challenges arise from uncertainty and variation, which we categorize into epistemic uncertainty, aleatoric uncertainty, and episodic variation.
- *Mechanisms*: Robustness mechanisms improve task goal achievement by tolerating or mitigating uncertainty and variation in contact interactions, preventing failures, or enabling recovery when failures occur.
- *Evaluation*: Evaluation methods define quantitative criteria for measuring manipulation performance and robustness, enabling principled comparison across approaches.

Among these dimensions, challenges play a central role. The importance of uncertainty has been recognized since the earliest days of robotics in the 1950s (Goertz 1952), which remains a central theme in robotics research (Mason 2012). Specifically, epistemic uncertainty stems from incomplete or inaccurate knowledge of system properties and can, in principle, be reduced through additional data or modeling. Aleatoric uncertainty reflects irreducible randomness, such as observation noise or stochastic transition disturbances. Episodic variation refers to systematic differences in the physical world that remain fixed within a single task execution but change across task episodes, such as variations in object identity, environment configuration, or robot embodiment (Cui and Trinkle 2021). An intuitive analogy is that of a robot searching for a light switch in a dark room. Episodic variation: across different rooms, the position or shape of the switch may vary. Epistemic uncertainty: within a given room, the initial belief about the switch location is uncertain but can be refined through interaction, such as touching the wall. Aleatoric uncertainty: even with perfect knowledge, however, randomness such as sensor noise, slips of the hand, or minor actuation errors may still occur.

3 Formulation of Manipulation Robustness

Given the definition in Section 2, we formulate the general problem of manipulation robustness using a unified partially observable stochastic control model. We then instantiate this model through probabilistic and control-theoretic perspectives, respectively. The unified model serves as an analytical framework to structure the discussion in the remainder of the paper. It is intended as a conceptual tool for analysis rather than a prescriptive model; it does not imply that a robot must explicitly represent these processes to achieve robust manipulation.

3.1 Unified Formulation

We formalize a manipulation episode as a discrete-time control process over a finite horizon T, denoted by the sequence $\{(x_t, u_t, y_t)\}_{t=0}^{T}$. Here, $x_t \in \mathcal{X}$ is the system state (encompassing both robot and object generalized coordinates), $u_t \in \mathcal{U}$ is the control input, and $y_t \in \mathcal{Y}$ is the partial observation. Note that we present here a time-discretized model, though in principle a continuous-time formulation can also be adopted.

The evolution of this system is driven by underlying deterministic functions subjected to noise:

$$x_{t+1} = f(x_t, u_t, w_t; \theta), \tag{1}$$

$$y_t = g(x_t, v_t; \theta). \tag{2}$$

In this framework, manipulation robustness is defined by how well a system achieves a target goal $\mathcal{G}$ despite three fundamental challenges:

Aleatoric Uncertainty The variables $w_t \in \mathcal{W}$ and $v_t \in \mathcal{V}$ represent exogenous process and measurement noise. This captures inherent, irreducible stochasticity in the physical world, such as thermal sensor noise, unmodeled micro-impacts, or random slip during contact.

Epistemic Uncertainty The parameter $\theta \in \Theta$ dictates the true physical and perceptual properties of the task, such as object mass, surface friction, and camera calibration. However, the robot rarely has perfect knowledge of θ or the true state x_t. It operates using internal estimates $\hat{\theta}$ and $\hat{x}_t$. The discrepancy between reality (θ, x_t) and the robot's belief $(\hat{\theta}, \hat{x}_t)$ constitutes epistemic uncertainty. Unlike aleatoric noise, this uncertainty can theoretically be reduced through exploration or learning.

Episodic Variation While epistemic uncertainty concerns what the robot does not know, episodic variation describes the objective physical differences between separate executions of a task. We define an episodic variation ν as a specific draw of the environment parameters, initial state, and goal region:

$$\nu = (\theta, x_0, \mathcal{G}) \in \mathcal{N}, \tag{3}$$

where $\mathcal{N}$ denotes the space of all possible task episodes. Note that the goal $\mathcal{G} \subset \mathcal{X}$ represents a desired set of states, especially the desired state of the manipulated object. This compact notation covers many terminal manipulation objectives, but it does not by itself encode temporal ordering among subgoals. Ordered long-horizon tasks, such as opening a door, entering a room, and then closing the door, would require a richer formulation, e.g., a sequence of goal sets $(\mathcal{G}_1, \ldots, \mathcal{G}_K)$ or a temporal task specification.

3.2 Probabilistic View

Learning-based and probabilistic planning methods treat the noise variables (w_t, v_t) and episodic variations (ν) as probability distributions. Consequently, the dynamics and observation models are viewed as stochastic transitions,

$$x_{t+1} \sim \mathcal{T}(\cdot \mid x_t, u_t; \theta), \; y_t \sim \mathcal{O}(\cdot \mid x_t; \theta). \tag{4}$$

Robustness in this paradigm is typically formulated as a Partially Observable Markov Decision Process (POMDP) and evaluated as the expected performance of a policy π across a distribution of episodes $\rho_{\mathcal{N}}$:

$$\Gamma(\pi) = \mathbb{E}_{\nu \sim \rho_{\mathcal{N}}} \big[J_\nu(\pi) \big], \tag{5}$$

where $J_\nu(\pi)$ is the expected return of the trajectory under the true parameters θ, despite the policy acting on its imperfect belief $\hat{\theta}$. A policy is defined as a mapping $\pi : \mathcal{B} \to \mathcal{U}$ from beliefs to actions, where a belief $b_t \in \mathcal{B}(\mathcal{X})$ is a probability distribution over the state space $\mathcal{X}$ conditioned on the history of observations and actions.

Table 1. Overview of the manipulation robustness mechanisms (Learning mechanisms in Table 2).

Submodule	Mechanism	Principle	Section
Perception	Active and Interactive Perception	Uncertainty Reduction	5.1.1
	Invariance in Perceptual Representations	Uncertainty & Variation Toleration	5.1.2
	Multimodality	Local Failure Toleration	5.1.3
Planning	Reasoning over Uncertainties	Uncertainty Reduction	5.2.1
	Motion Strategy and Funneling	Uncertainty Reduction	5.2.2
	Robustness Margin	Uncertainty & Variation Toleration	5.2.3
	Closure	Uncertainty & Variation Toleration	5.2.4
Control	Predictivity	Failure Prevention	5.3.3
	Active Compliance	Uncertainty & Variation Toleration	5.3.1
	Control Invariance	Uncertainty & Variation Toleration	5.3.2
Hardware	Passive Compliance	Uncertainty & Variation Toleration	5.5.1
	Adhesion	Uncertainty & Variation Toleration	5.5.2
	Morphology Adaptation	Temporary Failure Recovery	5.5.3

3.3 Control-Theoretic View

In contrast, we follow robust control frameworks (e.g., H_∞ or Robust MPC) and model aleatoric noise w_t, v_t and episodic variations ν as unknown, deterministic variables confined within bounded sets $\mathcal{W}$, $\mathcal{V}$, and $\mathcal{N}$. Robustness here is framed as a min-max optimization problem. The goal is to synthesize a policy that guarantees constraint satisfaction and minimizes a cost function $c(x_t, u_t)$ against the worst-case possible realizations of the uncertainties:

$$\min_{\pi} \max_{\nu\in\mathcal{N}, w_t\in\mathcal{W}, v_t\in\mathcal{V}} \sum_{t=0}^{T} c(x_t, u_t). \tag{6}$$

While the probabilistic view maximizes *average* reliability across a distribution of environments, the control-theoretic view provides strict performance and safety guarantees in a *worst-case* robustness sense.

4 Principles of Robust Manipulation

There are diverse mechanisms for achieving robustness in robotic manipulation, spanning perception, planning, control, learning, and hardware. While these mechanisms are often studied separately, many can be understood through several recurring principles. At a high level, robustness mechanisms can be organized along two axes. The first concerns how systems cope with uncertainty and variation: either by *reducing* epistemic uncertainty or by *tolerating* uncertainty and variation. The second concerns failure management: whether robustness is achieved primarily through *preventing* failures from occurring, *recovering* from temporary failures or *tolerating* local failures.

4.1 Uncertainty and Variation Regulation

One route to robustness is to *reduce* epistemic uncertainty by improving the robot's knowledge of the environment, task, or system state. This principle underlies active and interactive perception (Section 5.1.1), belief-space planning methods that explicitly reason over uncertainty (5.2.1), and world-model-based learning approaches that learn predictive representations of environment dynamics (5.4.2).

In contrast, a second route is based on *tolerating* the challenges: it accepts that aleatoric uncertainty and episodic variation cannot be eliminated, and instead seeks to make task execution insensitive to them. This principle is perhaps more ubiquitously employed in robotic manipulation. One common strategy is *compliance*, realized both actively through control (Section 5.3.1) and passively through hardware and materials (5.5.1), which reduces sensitivity to contact uncertainty and modeling errors. Another is *invariance*, where behaviors or representations are designed to remain unaffected by uncertainty or variation, as exemplified by closure-based grasping (5.2.4), invariant perceptual representations (5.1.2), and control invariance (5.3.2). Robustness can also arise from *mechanics*-driven strategies that exploit task and environmental structure, such as motion strategies and funnels (5.2.2). Toleration also appears in learning-based methods that generalize across scenes, tasks, and environments through data diversity or policy design (5.4.1, 5.4.2), as well as in hardware mechanisms such as adhesion (5.5.2), which leverage favorable contact physics to tolerate uncertainty and variation.

These two strategies are not mutually exclusive. Robust manipulation systems typically combine them, reducing uncertainty where possible while remaining effective in the presence of irreducible uncertainty.

4.2 Failure Management

Along the second axis, robustness can be understood through failure management. Here, two complementary strategies emerge: failure prevention, which seeks to avoid failure a priori, and failure recovery or tolerance, which assumes that failures may occur and emphasizes recovery from them or tolerance of local failure. Tasks involving safety-critical operations or irreversible consequences (e.g., an object's fragility) often demand preventative mechanisms, as failures are unacceptable. In contrast, for tasks where temporary failures can be tolerated, such as regrasping a rigid object after a drop, recovery mechanisms may be sufficient. Local failure tolerance further complements these strategies by allowing failures confined to specific components, sensing modalities, or subtasks to be absorbed without causing overall task failure.

Failure prevention aims to avoid entering failure-prone states altogether. This principle appears in a variety of robustness mechanisms. Examples include predictive control methods that down-weight risky rollouts in receding-horizon optimization (Section 5.3.3), closure-based grasp synthesis that remains stable under disturbance wrenches (5.2.4), and compliance-based control and hardware design (5.3.1, 5.5.1), which reduce the likelihood of failure arising from contact modeling errors. Many learning-based manipulation policies similarly emphasize failure avoidance, either by training predominantly on successful demonstrations or by incorporating conservative and safety-aware objectives that discourage risky behaviors and constraint violations (5.4.3).

By contrast, a complementary strategy focuses on restoring task execution once failure has occurred. From this perspective, failures need not be catastrophic; they may be temporary and recoverable, or local and confined to particular components. Recent learning-based systems have demonstrated the ability to acquire recovery behaviors from imperfect or failed demonstrations (Section 5.4.1), or to explicitly learn recovery policies that resume task execution following failure. *Temporary failures* may also arise from distribution shifts or changing task conditions, motivating adaptation mechanisms, either through morphological adaptation (5.5.3) or learning-based adaptation (5.4.4).

Local system failures, on the other hand, refer to failures confined to a particular component, sensing modality, actuator, module, or subtask, without immediately implying failure of the overall manipulation task. Such failures can be mitigated through modularity and redundancy. Modular robotic systems (Yim et al. 2000; Xie et al. 2019) can localize failure and prevent its propagation to the rest of the system. Redundancy provides multiple pathways to task success and allows degradation in one component to be compensated for by others. Examples include multi-sensor fusion (Section 5.1.3), dual end-effectors (Zeng et al. 2022), and alternative task-level plans (Siméon et al. 2004). Both modularity and redundancy are also common principles in biological systems; for example, ants dynamically reassign roles during cooperative transport, using redundant individuals to maintain collective performance despite local failures (Gelblum et al. 2015).

Guided by these two axes of robustness principles—uncertainty and variation regulation, and failure management—we next review the mechanisms that enable robust manipulation.

5 Mechanisms of Robust Manipulation

In this section, we survey existing approaches to robustness in robotic manipulation. We organize the literature into five submodules of a robotic system—perception, planning, control, learning, and hardware (Fig. 1)—and summarize the mechanisms in Table 1 and Table 2. Within each submodule, we discuss the mechanisms through the lens of the robustness principles introduced in Section 4, while noting that each mechanism is listed under its *dominant principle* even though it may reflect multiple principles in practice. We also highlight the challenges they address, representative works, and illustrative manipulation tasks. Owing to the rapid recent progress and growing attention to robot learning, the policy learning section (5.4) covers a comparatively larger body of recent work. At the same time, we emphasize foundational contributions, particularly in planning and control, to modern robotic manipulation.

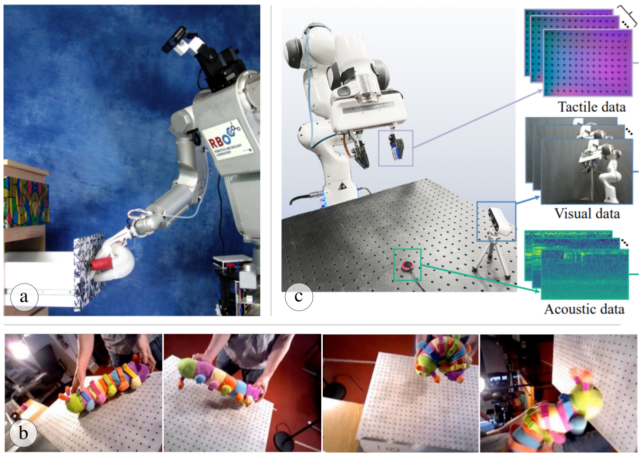

Figure 2. Examples of robustness mechanisms in perception. (a) Interactive perception: the robot interacts with the environment to improve perceptual understanding (Martín-Martín 2018). (b) Invariance in perceptual representation: the perception of the toy remains invariant to changes in camera pose, object pose or deformation, and lighting conditions (Florence et al. 2018). (c) Multimodal sensing: perception combining tactile, visual, and acoustic sensing modalities (Li et al. 2023). Figures are used with permission from the authors, or adapted under CC BY.

5.1 Perception

In the context of manipulation, the goal of perception is to acquire sufficient information for a robot to interact effectively with objects in its environment. Perception is particularly challenged by epistemic uncertainty, since task-relevant properties and system dynamics are often initially unknown. Important information, such as the number of objects in a scene, their motion constraints, or friction properties, may only become available through interaction with the environment. Consequently, active and interactive perception (Section 5.1.1) constitute important mechanisms for *reducing* uncertainty. At the same time, perception should maintain beliefs only over variables that are relevant to the task while remaining invariant to irrelevant features (5.1.2); such invariance is an effective mechanism for *tolerating* episodic variation across scenes. From a failure-management perspective, multimodal perception provides redundancy across sensing modalities and can mitigate *local failures* within the perception system (5.1.3).

5.1.1 Active and Interactive Perception Epistemic uncertainty in robot perception can be *reduced* by exploiting a robot's ability to act in and interact with its environment (Aloimonos et al. 1988). A major challenge is occlusion, which can often be addressed simply by moving a camera to acquire a different viewpoint (Bajcsy 1988), for example, to support grasp execution (Morrison et al. 2019) (Fig. 2-a). By actively selecting viewpoints that reveal previously hidden information, perceptual uncertainty about the environment is reduced, facilitating subsequent manipulation. Beyond contact-free viewpoint selection, robots can also physically interact with their environment to acquire information. This paradigm, known as interactive perception (Bohg et al. 2017; Xiong et al. 2025), is likewise observed in animals that employ exploratory behaviors to gather information about their surroundings (Chappell et al. 2012). Many task-relevant object properties are difficult or impossible to infer from passive observation alone. For example, lifting a milk carton reveals its mass distribution, while removing the lid of a box may reveal its contents. In cluttered environments, robots may need to manipulate obstructing objects to expose a target object (Xu et al. 2015). Interactive perception can also reveal motion constraints and kinematic structures of articulated objects (Katz and Brock 2008). Across these examples, the common principle is the active reduction of epistemic uncertainty through action and interaction.

5.1.2 Invariance in Perceptual Representations Perception for manipulation is challenging because objects need to be reliably detected from different viewing angles and lighting conditions (Fig. 2-b). In this sense, invariant perceptual representations help *tolerate* episodic variations in viewpoint, appearance, and environmental conditions. Before deep learning approaches started to dominate computer vision (Krizhevsky et al. 2012), computer vision pipelines were often based on features such as SIFT (Lowe 2004), SURF (Bay et al. 2006), or ORB (Rublee et al. 2011), which were designed to be invariant to rotation, scale, and possibly other transforms. In deep learning based approaches, such invariance is often achieved using data-augmentation (Krizhevsky et al. 2012) and pooling mechanisms (LeCun et al. 2002). More recently, foundation models such as DINO (Oquab et al. 2024) have demonstrated the ability to learn task-agnostic visual representations that are robust to clutter, occlusion, and illumination changes. At an even higher level of abstraction, language-based scene representations exhibit strong invariance to visual variations, contributing to the robustness and generalization capabilities observed in vision-language-action models (Zitkovich et al. 2023).

5.1.3 Multimodality Robotic manipulation can leverage multiple sensing modalities, including visual, tactile, auditory, and proprioceptive information, to perceive and model the environment, thereby enabling more effective interaction with it (Fig. 2-c). These modalities are both complementary and partially redundant in their functions. As a result, failure in one modality need not lead to failure of the overall perception system; other modalities can still provide sufficient information to support task execution. Humans, for example, often rely on touch when vision is unavailable. In robotics, vision provides rich global scene awareness but is susceptible to occlusion, whereas tactile sensing offers detailed local information about properties such as stiffness, texture, friction, and contact geometry. Audition can capture transient interaction events, such as the sound of a snap-fit assembly (Li et al. 2023). These modalities have been successfully combined in prior work to improve perception robustness (Izatt et al. 2017; Homberg et al. 2019). Beyond multimodal sensing, redundancy can also arise within a single sensing modality, as exemplified by the large number of whiskers in rodents or the compound eyes of insects. From a failure-management perspective, such redundancy helps mitigate *local failures* and prevents them from compromising overall perceptual functionality.

5.2 Planning

Manipulation planning operates on information provided by perception and is therefore inevitably affected by perceptual imperfections. In particular, imperfect perception and state estimation introduce *epistemic uncertainty*: a mismatch between the true robot–environment parameters θ and the internal model $\hat{\theta}$ (dynamics uncertainty), and consequently between the true state x_t and its estimate $\hat{x}_t$ (state uncertainty). Under such mismatches, trajectories computed by the planner may fail to achieve the desired goal $\mathcal{G}$. Manipulation planning must also contend with aleatoric uncertainty in the dynamics and observation models, as well as episodic variation across task instances.

Existing robust planning mechanisms improve manipulation performance through two complementary strategies. On the one hand, epistemic uncertainty can be *reduced* by explicitly reasoning about it (Section 5.2.1) or by exploiting task mechanics through motion strategies and funnels (5.2.2). On the other hand, planning can *tolerate* uncertainty and disturbances by incorporating robustness margins (5.2.3) or by establishing forms of *closure* around target objects (5.2.4).

5.2.1 Reasoning over Uncertainties Even high-quality sensors and advanced perception algorithms cannot fully eliminate epistemic uncertainty. A direct way to mitigate its effect is to represent uncertainty explicitly as a probability distribution over possible states—commonly referred to as a belief—and to update this belief online using new observations (Fig. 3-a). Formally, the agent maintains a belief $b_t \in \mathcal{B}(\mathcal{X})$ over the state x_t, conditioned on the history $h_t = (y_{0:t}, u_{0:t-1})$ under the internal model $\hat{\theta}$, i.e., $b_t(x) = \mathbb{P}(x_t = x \mid h_t, \hat{\theta})$. Planning then operates in belief space through a policy $u_t = \pi(b_t)$, aiming to improve expected task performance $\Gamma(\pi)$ across episodic variations ν.

Belief representations enable exploration that actively *reduces* uncertainty by trading off immediate task performance against information gathering. They also support cautious control, allowing robot policies to account for the reliability of available information. For instance, Platt et al. (2010) applied belief dynamics to grasping, synthesizing locally optimal feedback policies in belief space with replanning. Similarly, Jankowski et al. (2024) incorporated pose and contact-dynamics uncertainty into a belief-space formulation and validated it on planar pushing without sensory feedback. At longer horizons, belief reasoning extends to uncertainty-aware task and motion planning; Kaelbling and Lozano-Pérez (2013) integrated belief propagation with symbolic decision making to address incomplete knowledge of object properties and locations in mobile manipulation.

5.2.2 Motion Strategy and Funneling Besides reasoning over uncertainties in belief space, another way to *reduce* uncertainty is to exploit task mechanics. In *motion strategies*, also termed sensorless manipulation by Erdmann and

Mason (2002), contact dynamics and gravity provide passive negative feedback that drives objects toward desired configurations, often without exteroceptive sensing. This contrasts with sensor strategies, which rely on rich observations but are susceptible to perceptual errors. Motion strategies are effective in diverse settings, including tray-tilting to orient objects of unknown pose using only gravity and contact (Erdmann and Mason 2002), push–grasp behaviors in clutter that funnel the target into the gripper while displacing distractors (Dogar et al. 2012), and industrial part feeders whose fences enforce consistent orientation under pose and contact variability (Peshkin and Sanderson 2002). More generally, environmental contact and gravity as task mechanics can *reduce* uncertainties, or collapse episodic variations into a predictable, smaller set of outcomes (often termed *extrinsic dexterity* (Dafle et al. 2014)). For example, pressing a loose pair of chopsticks against a tabletop passively aligns their tips through the collision.

A representative abstraction of motion strategy is *funneling* (Mason 1985). A funnel is a region of attraction $\mathcal{F} \subset \mathcal{X}$ such that, for $x_0 \in \mathcal{F}$ and bounded disturbances $w_t \in \mathcal{W}, v_t \in \mathcal{V}$, the resulting execution converges to the goal set $\mathcal{G}$ (i.e., $x_T \in \mathcal{G}$), in other words *reducing* the pose or geometry uncertainty and model mismatch. Such strategies exploit repeated contacts and environmental constraints to steer outcomes without precise sensing, as in vibrating bowl feeders used for industrial part feeding. The funneling concept traces to pre-image backchaining (Lozano-Perez et al. 1984), which recursively constructs sets of states from which motions succeed under bounded execution errors. It was later extended by LQR-Tree methods (Tedrake 2009), where verified funnels cover the reachable state space to enable robust trajectory planning under model uncertainty. More recently, composing simple in-hand manipulation funnels has shown surprising robustness to variations in object mass and pose (Bhatt et al. 2021) (Fig. 3-b).

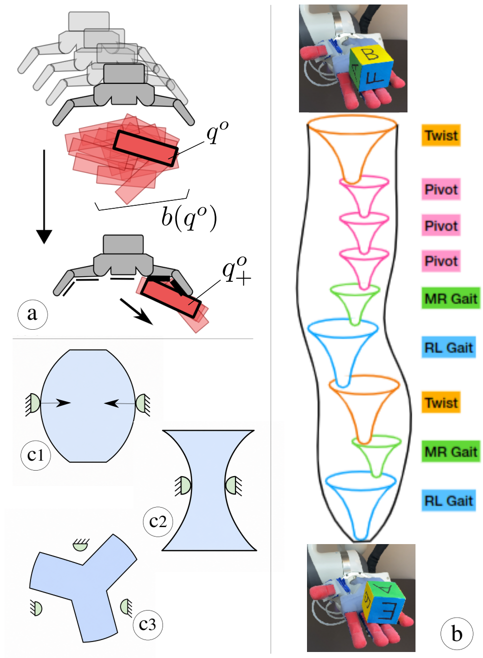

Figure 3. Examples of robust planning mechanisms. (a) Reasoning over uncertainties: planning in belief space, where the belief of the box position $b(q^o)$ is updated to $b(q^o_+)$ after it is pushed by the robot hand (Koval et al. 2016). (b) Funneling with in-hand reorientation primitives, where primitives are chained to reach the target cube orientation (Bhatt et al. 2021); (c) Force closure (c1), form closure (c2), and caging (c3). Photos in (a) and (b) are adapted under CC BY; figure (c) is inspired by Fox (2020); Rodriguez et al. (2012).

5.2.3 Robustness Margin Beyond explicitly estimating or compensating for uncertainty, an alternative strategy is to *tolerate* uncertainty by planning with a robustness margin. The central idea is to generate motions that remain sufficiently far from constraint boundaries, so that moderate disturbances or model mismatch do not immediately lead to failure. Many manipulation tasks require changing the object's state through interaction while satisfying state and input constraints, $x_t \in \mathcal{X}$ and $u_t \in \mathcal{U}$. These constraints encode collision avoidance, contact consistency, kinematic limits, torque bounds, and related feasibility conditions. Classical methods in this category primarily guarantee feasibility under a nominal model (Berenson et al. 2009; Siméon et al. 2004; Saha and Isto 2007).

However, under epistemic mismatch in θ and stochastic disturbances, executions may deviate from the nominal constraint manifold and violate constraints. Robust manipulation planning addresses this gap by constructing robustified feasible subsets, thereby maintaining a "margin" from constraint violation to tolerate uncertainty. For example, chance-constrained formulations enforce the probabilistic satisfaction of contact and state constraints under stochastic dynamics (Blackmore et al. 2011; Shirai et al. 2022), while set-based approaches employ tightened constraint sets or disturbance-invariant tubes to guarantee feasibility under bounded uncertainty (Mayne et al. 2005). These approaches ensure that constraint satisfaction is preserved within a specified disturbance or risk budget, thereby tolerating uncertainty at the planning level. We will discuss several more works along this line in Section 6.2.2 and how they can be used as robustness evaluation protocols. The rich body of work in robust control theory provides a strong foundation for future developments of this mechanism.

5.2.4 Closure In the spirit of uncertainty tolerance, robustness can also be achieved through *closure*, which imposes geometric or force constraints on object motion and thereby makes the motion resistant to, for example, external disturbances to manipulated objects. In many grasping and fixturing tasks, the goal is to maintain the object within an in-hand stable set. Therefore, closure can be viewed as shaping the robot-object interactions such that the object's reachable configurations remain within a bounded subset $\mathcal{G} \subset \mathcal{X}$ despite bounded disturbances $w_t \in \mathcal{W}$ in the dynamics $x_{t+1} = f(x_t, u_t, w_t; \theta)$. *Force closure* captures equilibrium grasps that can resist arbitrary external wrenches via internal contact forces (Howard and Kumar 1996) (Fig. 3-c1). *Form closure* is a geometric condition in which contact constraints eliminate all object motions even without friction (Bicchi 1995) (Fig. 3-c2). *Caging* (also referred to as "object closure" in some contexts) further relaxes contact requirements by trapping the object within a bounded region of its configuration space without precise force regulation (Pereira et al. 2004) (Fig. 3-c3).

Across these closure methods, the shared principle is to restrict the configuration space so that the object cannot move freely or escape, thereby making the manipulation outcome *tolerant* to perceptual errors, control imprecision, and external disturbances.

5.3 Control

Whereas planning operates at a higher level, often through open-loop decision making, control closes the loop in real time by responding to sensory feedback and correcting deviations that no plan can fully anticipate. A large class of robotic control methods addresses motion control in free space, such as trajectory tracking or point-to-point reaching, where the manipulator operates without contact. Manipulation, however, is fundamentally characterized by the making and breaking of unilateral frictional contacts, which introduce discontinuities and uncertainties. As a result, manipulation control must regulate not only motion but also forces and contact events. Controllers in contact-rich settings must therefore contend with external disturbances and actuation errors, as well as epistemic uncertainty arising from imperfect contact models. Poorly designed controllers can amplify these challenges, leading to instability or excessive contact forces, particularly in interactions with rigid environments. From the perspective of robustness, *compliance* (5.3.1) and *invariance* (5.3.2) primarily serve to *tolerate* uncertainty and variation, whereas *predictive* control (5.3.3) contributes to *failure prevention* by anticipating future contact events and system evolution.

5.3.1 Compliance Compliance balances the control of position and force; higher compliance allows larger position deviations under external forces. It can be realized passively through mechanical design (Section 5.5.1) or actively within the control loop (Mason 2007). Active compliance control regulates how a robot responds to external forces during contact, allowing it to comply rather than rigidly enforcing a motion. This property makes compliance a toleration strategy that bridges the gap between the robot's internal model and the complex, often unpredictable, physics of the real world. Impedance and admittance control represent the two classical approaches. Impedance control enforces compliance from motion to force by commanding motion trajectories that indirectly regulate contact forces (Hogan 1984),

$$M_d(\ddot{x} - \ddot{x}_d) + D_d(\dot{x} - \dot{x}_d) + K_d(x - x_d) = \tau_{\text{ext}}, \quad (7)$$

where $x \in \mathbb{R}^m$ denotes the task-space position (or pose) of the end-effector, x_d is the desired trajectory, and τ_{ext} is the external wrench arising from contact. The matrices M_d, D_d, and K_d are the desired inertia, damping, and stiffness parameters that define a virtual mass–spring–damper system. This equation specifies the closed-loop interaction dynamics: external forces induce bounded motion deviations governed by (M_d, D_d, K_d). On the other hand, admittance control does so from force to motion by computing the motion response to a commanded force (Dimeas and Aspragathos 2015).

Classical active compliance methods typically rely on accurate models or restrictive assumptions. More recently, model-free approaches have emerged, learning compliant manipulation behavior from human demonstrations or through reinforcement learning (Xu et al. 2026; Hou et al. 2025; Kamijo et al. 2024). For example, an adaptive compliance policy can adjust compliance both spatially and temporally from demonstrations for contact-rich tasks (Hou et al. 2025). Future directions may explore combining active compliance with passive mechanical compliance to enhance adaptability under contact uncertainty.

5.3.2 Invariance Invariance ensures that a system executes the behavior prescribed by a nominal controller while remaining within a certified safe set, such as bounds on force tracking error (Polverini et al. 2017) or admissible state regions. Like other active compliance control, invariance-based methods do not reduce epistemic uncertainty in states or parameters. Rather, they achieve robustness by constraining system evolution so that bounded disturbances or modeling errors cannot drive the state outside a prescribed safe set; in this way, disturbances are *tolerated* rather than eliminated. Formally, consider the closed-loop dynamics $x_{t+1} = f(x_t, u_t, w_t; \theta)$ under an output-feedback policy $u_t = \pi_t(b_t)$, where $b_t(x) = \mathbb{P}(x_t = x \mid h_t, \hat{\theta})$. A set $\mathcal{C} \subset \mathcal{X}$ is robustly forward invariant if

$$x_t \in \mathcal{C} \;\Rightarrow\; x_{t+1} \in \mathcal{C}, \quad \forall w_t \in \mathcal{W}.$$

Invariance-based control enforces this condition by modifying or projecting control inputs when the boundary of $\mathcal{C}$ is at risk of being violated, thereby guaranteeing constraint satisfaction despite aleatoric uncertainty and bounded modeling errors.

Practical approaches include reachability analysis, which conservatively overapproximates all possible evolutions to guarantee safety (Holmes et al. 2020), and control barrier functions (CBFs), which enforce invariance through real-time inequality constraints. When combined with control Lyapunov functions, CBF-based methods unify safety and goal achievement in an optimization framework, as demonstrated in tasks such as ball balancing (Wang et al. 2025b).

5.3.3 Predictivity Neuroscience suggests that human manipulation relies on the integration of long-term prediction with short-horizon reactivity (Flanagan et al. 2006). Reactivity provides rapid sensory feedback, particularly tactile and visual, to detect mismatches between expected and actual outcomes and correct errors. Predictivity, on the other hand, enables the anticipation of collisions or contact events. In this way, predictive control primarily *tolerates* uncertainty and variation by proactively compensating for their future effects, thereby helping *prevent* failures before they occur. Purely local reactive control may fail in contact-rich settings, where abrupt changes in dynamics due to contacts can drive the system into unsafe states. To address this, model predictive control (MPC) incorporates a lookahead structure with reactive schemes. It optimizes actions over a finite horizon and executes the first control input. For example, MPC-based hybrid force–motion control is employed to simultaneously regulate end-effector trajectories and contact forces during interaction, explicitly accounting for contact mode transitions within the prediction horizon (Jiang et al. 2025); a contact-implicit MPC is utilized that predicts

Table 2. Overview of learning-based robustness mechanisms for robotic manipulation (Complementary to Table 1).

Mechanism	Method	Principle	Section
Training Distribution	Domain Randomization Data Augmentation Data Scaling and Task Diversity	Variation Toleration	5.4.1
Policy Architecture	Perceptual and Structural Inductive Biases Generative Policy Parameterizations Temporal Abstraction and Hierarchy World-model-based Policies	Uncertainty & Variation Toleration Uncertainty & Variation Toleration Failure Prevention Uncertainty Reduction	5.4.2
Learning Objective	Adversarial Training Objectives Safety- and Constraint-based Objectives Regularized and Conservative Objectives	Uncertainty & Variation Toleration Failure Prevention Failure Prevention	5.4.3
Policy Adaptation	Interactive Human-in-the-loop Imitation Autonomous Continual Improvement	Temporary Failure Recovery	5.4.4

motion and contact forces within a trust region, enabling stable local control under unilateral contact and friction constraints (Suh et al. 2025); tube-based MPC guarantees that the real trajectory remains within a bounded "tube" around the nominal one, ensuring performance under external disturbances and model mismatch (Nubert et al. 2020).

5.4 Policy Learning

In parallel with the classical perception–planning–control pipeline, end-to-end policy learning has emerged as an alternative paradigm for manipulation. Learned policies can adapt to unstructured variations in the open world that are difficult to address analytically. In this sense, policy learning mainly improves robustness by *tolerating* episodic variation and disturbances. Depending on the policy design, it can both *prevent failure* through robust action selection and *recover* after temporary task failure.

Within the formalism of Section 3, these methods instantiate a parameterized policy π_ϕ and train it from data so that the robust performance functionals $\Gamma(\pi)$ remain high across an episodic distribution $\rho_\mathcal{N}(\nu)$. In the rest of this section, we group policy-learning approaches according to four complementary ways (Table 2) in which they achieve robustness: some methods shape the tasks, environments, and perturbations that π_ϕ is trained on (5.4.1); some methods design policy classes and architecture whose inductive biases and temporal structure reduce sensitivity to nuisance variation and partial observability (5.4.2); some methods modify the learning objective to better align training with robust performance (5.4.3); and others allow policies to adapt at deployment through residual corrections, online updates, or test-time adaptation (5.4.4).

5.4.1 Training Distribution In the notation of Section 3, robustness is measured by the functional $\Gamma(\pi)$ evaluated over an episodic distribution $\rho_\mathcal{N}(\nu)$. A branch of approaches acts directly on this distribution: it shapes how training episodes are generated while keeping the policy class Π and the learning objective fixed. These methods aim to *tolerate* variations across episodes by broadening the training distribution over dynamics, initial states, environments, and goals.

A first family of methods uses *domain randomization*, widely used, particularly in reinforcement learning. Instead of a single nominal environment, one defines a family of "worlds" θ (and sometimes x_0 and $\mathcal{G}$) from a broad distribution. Policies trained under such variation can transfer to real manipulation with minimal tuning, as demonstrated for vision-based grasping appearance (Tobin et al. 2017) and for dexterous in-hand manipulation (OpenAI et al. 2019; Andrychowicz et al. 2020). By expanding the training distribution $\rho_\mathcal{N}$ to encompass diverse scenarios, these methods enable the policy to *tolerate* variations in dynamics and appearance.

A second line achieves robustness via *data augmentation* on training trajectories. These methods keep the generator of ν fixed, but expand the empirical distribution $\hat{\rho}_\mathcal{N}$ by transforming demonstrations according to task priors: prior work applies SE(2)/SE(3) transformations to scenes and actions to exploit invariance and equivariance (Florence et al. 2019; Mandlekar et al. 2023; Ameperosa et al. 2025), re-renders fixed trajectories from novel viewpoints to vary appearance (Zhou et al. 2023a; Zhang et al. 2024), and retargets demonstrated trajectories to new object and scene configurations (Yu et al. 2023; Chen et al. 2023; Xue et al. 2025b). Beyond such structured priors, a related class injects calibrated noise into states and actions to mitigate covariate shift and compounding error, making the policy more robust (Laskey et al. 2017; Ke et al. 2021; Simchowitz et al. 2025).

A third class relies on *data scaling and task diversity*, as most visibly demonstrated in recent vision–language–action (VLA) models. Rather than explicitly parameterizing randomization or augmentations, these methods assemble very large, heterogeneous datasets of manipulation episodes spanning many robots, scenes, tasks, and language-specified goals, and train a single policy across this mixture (Brohan et al. 2023; Zitkovich et al. 2023; Open X-Embodiment Collaboration 2024; Kim et al. 2025). Robustness here is treated as generalization over a very broad empirical $\hat{\rho}_\mathcal{N}$ that pools demonstrations and deployments across embodiments, tasks, and environments, with a single π_ϕ expected to maintain high $\Gamma(\pi_\phi)$ throughout this mixture. An important direction in data scaling extends beyond successful task executions to include *recovery* behaviors following *temporary failures*, capturing the adaptive strategies required in unstructured environments. Policies trained only on ideal trajectories may overfit to nominal conditions. For example, Large Behavior Models incorporate hybrid sim-to-real datasets containing both successful executions and recovery episodes (TRI LBM Team et al. 2026).

5.4.2 Policy Architecture A policy is a mapping $\pi : \mathcal{B} \to \mathcal{U}$ from beliefs to actions. A line of approaches pursues robustness by designing policy architecture and data representation: they constrain how information about observations is encoded and how actions are parameterized, without changing the training distribution $\rho_{\mathcal{N}}$ or the learning objective. The goal is to make the mapping $b_t \mapsto u_t$ inherently *tolerant* to nuisance variation, partial observability, and long horizons.

First, *perceptual and structural inductive biases* build a geometric state from raw observations. Keypoint-based representations encode objects and goals as a small set of 2D/3D landmarks and express tasks as geometric relations among these keypoints, so the policy operates in a low-dimensional space aligned with contacts and affordances rather than raw pixels (Manuelli et al. 2019; Qin et al. 2020; Huang et al. 2025). SE(3)-equivariant representations constrain internal features (and sometimes actions) to transform consistently under rigid motions of the workspace, so that rotating or translating the scene induces the same transformation in the encoded state and allows one controller to reuse a strategy across pose-shifted instances of a task (Simeonov et al. 2022; Ryu et al. 2022; Eisner et al. 2024; Wang et al. 2025a). Graph-structured encoders treat objects, robot links, and goals as nodes with edges capturing spatial or relational constraints, and use message passing to learn local interaction rules that generalize to scenes with more objects, new goal configurations, and multi-object rearrangement (Li et al. 2020; Lin et al. 2022; Huang et al. 2023). Across these designs, robustness comes from aligning the learned representation with the geometry and relational structure of the task, so that pose- and relation-preserving changes in the scene tend to yield similar decisions, resulting in stable performance under these variations.

Second, *generative policy parameterizations* implement π_ϕ as a conditional generative process over short action sequences, most commonly using diffusion or flow-matching models for manipulation tasks (Chi et al. 2023; Black et al. 2025). The policy maps beliefs b_t to actions by starting from injected noise and iteratively refining a sample via a supervised denoising or flow-integration process. Empirically, this iterative computation with noise injection induces an inductive bias that improves closed-loop robustness to covariate shift (Pan et al. 2026).

Temporal abstraction and hierarchy represent actions as skills or options that span multiple timesteps. The policy acts in an extended action space of options: a high-level policy selects a skill, and a low-level controller executes it until termination, so the decision-making problem becomes a semi-MDP with a shorter effective horizon (Sutton et al. 1998). Hierarchical RL methods learn both the skills and the high-level policy, often using subgoal states as the interface between levels and off-policy updates to remain sample efficient in continuous control (Nachum et al. 2018). In manipulation, skill-based approaches learn and plan with such temporally extended behaviors, such as predicting long-horizon discrete action sequences from a scene image for sequential physical reasoning (Driess et al. 2021), discovering reusable closed-loop skills from unsegmented demonstrations and training a meta-controller to compose them (Zhu et al. 2022), or chaining separate diffusion models for individual skills to solve unseen long-horizon tasks under constraints (Mishra et al. 2023). Robustness in this class arises from horizon reduction: the high-level policy makes skill-level decisions per episode, which reduces compounding error, supports reuse of skills across task variations, and thereby helps *prevent* failures before they occur.

Lastly, *world-model-based policies* couple the controller with a learned dynamics model in a latent space. They maintain a latent state $z_t = e(o_{0:t}; \theta)$ together with a predictive model $\hat{p}_\psi(z_{t+1} \mid z_t, u_t)$, and use short imagined rollouts in this latent space to reason about future states before committing to an action (Hafner et al. 2019; Ichter and Pavone 2019; Hansen et al. 2022; Hafner et al. 2023). World-model-based policies can *reduce* effective epistemic uncertainty by learning predictive structure over how the environment evolves under the robot's actions. Rather than eliminating uncertainty, the learned model provides an internal approximation of otherwise unknown or partially observed dynamics that can be queried during decision making. In manipulation, such latent models are used for planning and runtime monitoring to score candidate actions under many predicted futures and reject those that would enter unsafe, constraint-violating, or high-risk regions (Liu et al. 2024; Nakamura et al. 2025; Sun and Song 2026).

5.4.3 Learning Objective While distribution- and architecture-centric approaches shape what the policy sees and how it represents it, in this part we introduce methods that instead shape what the policy optimizes for, embedding worst-case performance, safety, or conservatism directly into the learning objective.

Adversarial training objectives make robustness explicit by optimizing the min–max criteria of Eq. (6) rather than a nominal expected return. The policy is trained against an adversary that chooses disturbances, environment parameters, or competing behaviors to minimize task performance while the protagonist maximizes it. In practice, the adversary injects destabilizing forces or physical perturbations to grasps and object poses, so policies are trained on deliberate worst cases rather than randomly sampled disturbances (Pinto et al. 2017b,a; Jian et al. 2021). By repeatedly optimizing against these hard cases, the learned policy becomes less sensitive to uncertainty and variation, thereby improving its ability to *tolerate* them at deployment.

Safety- and constraint-based objectives restrict which policies are admissible when we evaluate robustness. In the POMDP view of Section 3, one augments the control problem with constraint costs and searches for policies that achieve high task performance $\Gamma(\pi)$ while keeping these costs below a threshold, as in constrained MDP methods such as Constrained Policy Optimization (Achiam et al. 2017). Related work uses safety layers or control barrier functions that project or override the learned action when it would violate certified constraints (Dalal et al. 2018; Cheng et al. 2019), thereby *preventing* failures by keeping execution within admissible regions.

Regularized and conservative objectives aim to improve robustness to epistemic uncertainty and limited data coverage by constraining how far the learned policy and value function extrapolate beyond the empirical episodic distribution $\hat{\rho}_{\mathcal{N}}(\nu)$. Offline RL work frames this as "extrapolation error" and shows that robust performance typically requires (i) behavior regularization, which keeps π close to actions well supported by the data, and (ii) pessimistic value estimates that down-weight out-of-distribution actions (Levine et al. 2020; Kumar et al. 2020; Kostrikov et al. 2021; Fujimoto and Gu 2021). In real-robot manipulation, these ideas appear in conservative offline RL that improves over imitation on tabletop tasks using logs from safe operation (Zhou et al. 2023b), in large-scale deployments that rely on heavily regularized off-policy updates for stable long-term waste-sorting with mobile manipulators (Herzog et al. 2023), and in regularized imitation methods that combine optimal-transport trajectory matching with a pull toward demonstration behavior to achieve few-shot real visual manipulation (Haldar et al. 2023). Across these examples, robustness comes from biasing learning toward regions of state–action space that are well covered and conservatively valued, thereby reducing out-of-distribution actions and helping *prevent* failures during execution.

5.4.4 Policy Adaptation

This line of work assumes a base policy π_{base} has been trained under some episodic distribution $\rho_{\mathcal{N}}^{\text{train}}(\nu)$, and focuses on how its behavior is modified using feedback from the deployed environment when the test-time distribution $\rho^{\text{test}}(\nu)$ differs from $\rho_{\mathcal{N}}^{\text{train}}(\nu)$. Unlike distribution-centric approaches, they do not redesign $\rho_{\mathcal{N}}$ up front, but instead close the loop between deployment performance and further adaptation.

First, *interactive human-in-the-loop imitation* refines a pretrained policy at deployment using online expert feedback. Humans can intuitively identify and correct failure-imminent situations or recover from temporary *failure*, providing targeted supervision where pre-trained policies are weakest (Liu et al. 2022). In DAgger-style methods, imitation learning is reduced to online no-regret learning: at each iteration, the current policy is rolled out to generate states, the expert labels those visited states with preferred actions, these pairs are added to an aggregate dataset, and a new policy is trained on the union (Ross et al. 2011). This ensures that the learned stationary policy performs well under the state distribution it induces, rather than only on states seen in expert demonstrations, which directly addresses covariate shift in sequential decision making. Subsequent work adapts this idea to human experts in safety-critical or latency-limited settings (Kelly et al. 2019), to remote teleoperation for contact-rich manipulation (Mandlekar et al. 2020), to "on-the-job" deployment where autonomy and learning run continuously with humans stepping in on hard cases (Liu et al. 2022), and to kinesthetic delta corrections with force-aware residual policy via compliance (Xu et al. 2026).

Second, *autonomous continual improvement* adapts policies during deployment without relying on explicit expert labels. One branch uses online RL on the real robot: a warm-start policy (often from demonstrations or offline data) continues to collect rollouts on the hardware, and policy updates π_ϕ under the actual test-time episodic distribution, with mechanisms such as safety constraints, automatic or cheap resets, and simple reward specification to make this practical for contact-rich manipulation (Levine et al. 2016; Rajeswaran et al. 2018; Johannink et al. 2019; Xu et al. 2023; Luo et al. 2024, 2025). A second branch uses in-context and prompt-based adaptation: sequence models are trained on multi-task data so that, at test time, a short prompt of teleoperated demonstrations or recent trajectories for the current variation ν conditions the model, enabling few-shot adaptation to new objects, layouts, and goals without gradient updates (Duan et al. 2017; Valassakis et al. 2022; Fu et al. 2025).

5.5 Hardware

Manipulation robustness can also be achieved through the robot's embodiment, complementing the contributions of perception, planning, control, and learning. This perspective is grounded in the concept of *morphological computation*, where physical structures perform functions that would otherwise require explicit control (Paul 2006). Through the design of materials, geometry, and actuation, the body can shape the system dynamics f such that a wider range of states x_t and actions u_t lead to successful outcomes, while reducing the consequences of errors in perception, modeling, and execution.

A recurring theme of hardware-based robustness is the *toleration* of object variation and external disturbances through physical design. This can be achieved through passive compliance (Section 5.5.1), which absorbs disturbances and accommodates contact uncertainty, or through adhesive contact mechanisms (5.5.2), which remain effective despite geometric variation and perturbations. In addition, some robotic systems achieve robustness through *morphological adaptation* (5.5.3), reconfiguring their physical structure in response to changing conditions and thereby restoring performance when an existing morphology becomes ineffective.

5.5.1 Passive Compliance

Passive compliance enhances manipulation robustness by physically *tolerating* disturbances and contact uncertainty through compliant, soft surface material, before they propagate through the control loop. It can be viewed as a mechanical counterpart to active compliance (Section 5.3.1) that shapes the interaction dynamics prior to control. By allowing bounded motion in response to external forces, compliant structures reduce the sensitivity of contact interactions to disturbances, geometric variation, and modeling errors. As a result, reliable and adaptive manipulation can often be achieved with reduced reliance on precise sensing, accurate models, and high-bandwidth control.

This principle is widely observed in both biological and robotic systems. Dolphins, for example, use marine sponges as protective tools while foraging on rocky seabeds (Fig. 4-a), thereby tolerating environmental uncertainty and avoiding damaging contacts (Mann et al. 2008). In robotic manipulation, passive compliance is commonly

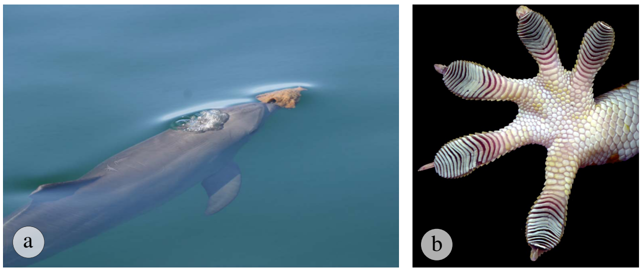

Figure 4. Examples of hardware mechanisms for robust manipulation. Hardware intelligence in robotics is often inspired by biological strategies. (a) Passive compliance: Dolphins wrap marine sponges around their beaks for protection (Mann et al. 2008), while robotic hands benefit from compliant surface material. (b) Adhesion: gecko feet inspire adhesive grippers, such as Song et al. (2017). Figure (a) is adapted under CC BY; image courtesy of (b): Prof. Kellar Autumn, Lewis and Clark College (Autumn 2006).

used to tolerate misalignment and reduce impact forces in assembly tasks (Drake 1978). Similar principles underlie soft grippers based on compliant materials or tendon-driven structures, which passively conform to object geometry under uncertainty (Deimel and Brock 2013). Hybrid designs further combine compliance with controllable stiffening, as exemplified by granular-jamming grippers, enabling compliant exploration followed by adaptive stiffening to securely grasp the object (Shintake et al. 2018; Amend et al. 2012).

5.5.2 Adhesion Adhesive material modifies robot-object contact surfaces to sustain strong tangential (shear) forces with minimal normal force, which is particularly suitable for delicate, thin, or flat objects. By enabling surface-based attachment rather than relying solely on friction-limited point contacts, adhesion alters the contact dynamics and makes manipulation more *tolerant* to uncertainty in contact location, object pose, and contact modeling. Common adhesion mechanisms include gecko-inspired dry adhesion and vacuum-based suction (Shintake et al. 2018). Gecko-inspired adhesives exploit van der Waals forces to generate strong shear interactions across a wide range of surface textures without requiring active control (Fig. 4-b). Such systems have been used to grasp bulky objects in microgravity (Jiang et al. 2017), demonstrating robustness to object misalignment and uncertain contact conditions. Vacuum-based suction is another widely adopted adhesion mechanism. Besides *adhesion*, many suction cups incorporate bellows that provide *passive compliance*, allowing the system to accommodate pose errors and conform to local surface curvature. Multi-affordance manipulators that combine suction cups with parallel-jaw grippers (Zeng et al. 2022) further exploit mechanical *redundancy*, improving robustness and extending applicability across a broader range of objects and environments.

5.5.3 Morphology Adaptation The environment and task requirements encountered by a robot are often not fixed. As a result, a morphology that performs well under one set of conditions may become ineffective when the environment, object properties, or task requirements change. Some robotic systems address this challenge by adapting their embodiment, modifying geometry, stiffness distribution, or kinematic structure to suit new conditions. In this way, morphological adaptation can be viewed as a mechanism for recovering from *temporary failure* caused by changing environments or task requirements (Liang et al. 2026). Complementary to passive compliance and adhesion, which shape interactions locally at the contact interface, morphological adaptation reshapes the embodiment at a global level.

In soft robotic hands, for example, the bending location and range can be adjusted by inserting a stiff rod into the center of each finger (Pagoli et al. 2021). Such reconfiguration enables the end-effector to switch between grasping modes and maintain performance across different object types without requiring complex sensing or control. PolyBot (Yim et al. 2000) further illustrates *modular* reconfiguration, where detachable modules autonomously rearrange to accommodate new workspace constraints. By adapting their morphology to changing conditions, such systems improve robustness, support fault tolerance, and extend manipulation capabilities in unstructured environments.

6 Evaluation Methods of Manipulation Robustness

Evaluating manipulation robustness requires quantitative criteria that reflect a system's capability to achieve its task objectives under challenges. In this section, we first present *empirical* evaluation protocols that assess how effectively a robot performs physical manipulation tasks in simulation or in the real world. We then survey *analytical* evaluation protocols that assess robustness using grasp-quality measures, margin-based metrics, and other abstract mathematical constructs rather than empirical test-time task performance.

6.1 Empirical Evaluation Protocols

Based on task performance, these evaluation protocols empirically assess whether and how effectively a manipulator achieves the task goal $\mathcal{G}$ despite varying challenges. They provide more direct comparisons for end users (Liconti et al. 2026) and are widely adopted by both empirical (learning-based) and analytical methods.

6.1.1 Goal-Oriented Measures Given a manipulation task specification together with domains of epistemic uncertainty, aleatoric uncertainty, and episodic variation, the most widely used robustness evaluation protocol is to measure empirical task success or failure (Mason 2018). This approach, commonly referred to as *Monte Carlo* robustness tests, executes the system repeatedly under a sampled model mismatch $(\hat{\theta}, \theta)$, stochastic disturbances w_t and v_t, and episodic variations ν, and estimates the probability of achieving the task goal. Concretely, robustness is the expected probability of goal satisfaction under the challenges,

$$\mathbb{E}_{\nu \sim \rho_{\mathcal{N}}, w_{0:T} \sim \rho_{\mathcal{W}}, v_{0:T} \sim \rho_{\mathcal{V}}} \left[\mathbf{1}\{x_T \in \mathcal{G}\} \mid (\hat{\theta}, \theta) \right], \quad (8)$$

often approximated by empirical success rates $n_{\text{success}}/n_{\text{total}}$ computed over multiple trials. This evaluation directly aligns with the robustness formulations introduced in Section 3.

Following this evaluation protocol, examples include success rates in object pose estimation under perception noise and occlusion (Mandlekar et al. 2023); success rates or expected returns in reinforcement learning, which often reduce to goal satisfaction indicators, evaluated across episodic variations such as different initial object placements or object instances (Luo et al. 2025; Andrychowicz et al. 2020); and success under adversarial perturbations, including external disturbances applied to manipulated objects or injected joint-level perturbations during execution (Jian et al. 2021). Variants of this approach include systematic parameter sweeps, where parameters in θ (e.g., inference latency, friction coefficients, lighting conditions) are varied across controlled ranges and performance is plotted as a function of the parameter (Chi et al. 2023).

6.1.2 Stage-wise Measures While the formulation in Eq. (8) focuses on terminal goal satisfaction, there are many manipulation tasks more naturally characterized by *trajectory-level subgoals*. In such cases, success depends not only on reaching a final state but also on correctly executing a sequence of intermediate stages. Consequently, several works adopt stage-wise evaluation metrics, in which a task is decomposed into semantically meaningful phases and success is assessed at each stage. For example, Kang et al. (2026) evaluate fragile-object manipulation policies across three stages: approach, stable grasp, and task completion. Xue et al. (2025a) decompose a paper-cup lifting task into clamping and lifting stages. Evaluators may also design task-specific semantic rubrics to assess policy behavior beyond aggregate success rates (Kress-Gazit et al. 2024). For instance, in a *Flip-and-Serve Pancake* task, relevant criteria may include: "Robot collided with anything?", "Robot flipped pancake?", "Robot picked up pancake?", in addition to overall task success. The stage-wise protocols provide a more fine-grained assessment of policy performance, enabling the identification of failure modes that may be obscured by a single binary success measure.

6.2 Analytical Evaluation Protocols

While success probability under challenges provides an empirical, task-level notion of robustness, it is often insufficiently fine-grained for answering more specific questions about *how* a manipulation system withstands challenges. In many scenarios, one is not only interested in whether a task succeeds, but also in the structure and degree of tolerance to disturbances, uncertainty, or variation. This is where analytical evaluation protocols become valuable: they aim for potential capability, rather than only realized capability, through explicit mathematical constructs. They are structured and interpretable, tailored to specific manipulation mechanisms and failure modes. They are often approximated by the success probability in Eq. (8). Unlike performance-based protocols, which primarily serve as empirical *diagnostic tools* for robustness assessment, analytical protocols can also function as *optimization objectives* embedded within planning, control, or learning methods for the synthesis of robust manipulation behavior. As follows, we introduce several analytical evaluation protocols.

6.2.1 Closure-based Grasp Quality Measures Closure properties characterize how contact forces and geometric constraints prevent an object from deviating from a desired configuration or set of configurations. These properties are typically analyzed from two different perspectives: *local* conditions at the contact and *global* conditions within the object's configuration space.

Local measures derived from force closure or form closure quantify the instantaneous robustness of quasistatic manipulation via wrench-space or geometric analysis. These measures can be evaluated either analytically or empirically. Analytical approaches, commonly referred to as grasp quality measures (Ferrari and Canny 1992), compute robustness using physics-based models to determine the system's ability to resist external disturbances. In contrast, empirical approaches approximate these analytic measures by learning from data. For instance, grasp datasets labeled with analytic quality values are used to train neural networks that predict robustness directly from visual or sensory inputs (Mahler et al. 2017), enabling scalable estimation in unstructured environments.

While local measures assess stability via instantaneous force or geometric conditions, many complex manipulation tasks depend on the global characteristics of the configuration and contact spaces—specifically, how reachable, connected, or "trapped" an object remains within its environment. Topological analysis captures such global invariants, providing a perspective complementary to purely local evaluations (Pokorny and Kragic 2015; Bhattacharya et al. 2015). For example, energy-bounded caging (Mahler et al. 2018) defines robustness in terms of the size or persistence of the configuration-space region that keeps the object contained. In this context, a larger connected component implies a greater structural tolerance to uncertainty or disturbances.

6.2.2 Margin-based Measures "Margin" to failure quantifies robustness by measuring the distance between a system's operating state and the boundary of task failure. Intuitively, systems maintaining larger margins are more robust, as they can tolerate greater disturbances, modeling errors, or stochastic noise before a failure occurs. Consider the task of placing a cup of coffee on a coaster: precise positioning is secondary as long as the liquid level remains safely below the rim. If the goal region $\mathcal{G}$ is defined as the set of states where all liquid remains inside the cup, its complement $\mathcal{G}^c$ corresponds to failure (spillage). While a binary reward only detects the occurrence of failure, a safety margin quantifies the resilience of the current state:

$$\sigma_{\text{margin}} = \text{dist}(x, \partial\mathcal{G}), \tag{9}$$

where $\text{dist}(\cdot)$ denotes a task-appropriate distance metric and $\partial\mathcal{G}$ is the boundary of the goal region. Positive values indicate the system is within $\mathcal{G}$, with larger values implying greater robustness.

In neurophysiology, this distance is often defined through energy, i.e., the energy required for the system to transition from a safe state x to failure $\mathcal{G}^c$ (Hasson et al. 2012). Humans intuitively maintain such margins by relying on simplified internal models rather than exact dynamics, selecting control strategies that preserve safety margins (Sternad and Hasson 2016). Similar concepts have been applied to evaluate robotic manipulation robustness (Dong et al. 2024). Analogous ideas appear in force- and wrench-based robustness metrics, such as grasp stability measures that quantify the maximum external wrench a grasp can resist before slippage (Roa and Suárez 2015). Extending the concept over time leads to the notion of a "safety tube", a region surrounding a nominal trajectory within which the system can remain despite bounded disturbances (Fox et al. 2006; Nubert et al. 2020).

6.2.3 Signal Temporal Logic Measures Signal Temporal Logic (STL) is a formal language for describing time-dependent task requirements over real-valued signals, such as contact force, object position, or velocity — examples include conditions such as maintaining a grasp force below a threshold, reaching a target region within a given time, or keeping an object stable throughout execution. STL further provides quantitative semantics that assign a robustness score measuring how strongly a trajectory satisfies or violates the specified requirements (Maler and Nickovic 2004; Donzé and Maler 2010). In robotics manipulation, STL robustness has been predominantly used as an optimization target — either as a reward signal in reinforcement learning (Kapoor et al. 2020; Li et al. 2017), a loss function for neural predictive control (Meng and Fan 2023), or a planning objective in task-and-motion planning (Takano et al. 2021). Its use as a post-hoc evaluation metric remains rare: Kress-Gazit et al. (2024) have explicitly leveraged STL robustness for grading learned manipulation policies after training, demonstrating how the robustness score reveals not only whether a policy succeeds but how close it is to failure across complex, temporally structured objectives. The STL measures are conceptually related to the margin-to-failure measures discussed above, but generalize them to composite temporal specifications involving sequencing, timing, and conditional constraints. A practical limitation is that STL robustness is scale-dependent: predicates expressed in different physical units (e.g., millimeters versus Newtons) produce robustness values on incomparable scales, so that the min/max aggregation inherent in the semantics can be dominated by the choice of units rather than by task-relevant difficulty (Varnai and Dimarogonas 2020; Dhonthi et al. 2021). Despite these challenges, the broader use of STL measures as an evaluation tool for manipulation robustness remains a promising direction.

6.2.4 Convergence and Divergence Measures In dynamically complex manipulation tasks such as tossing, catching, or balancing, robustness often arises from the intrinsic structure of the system dynamics. Humans routinely exploit such dynamics by inducing self-correcting behaviors in which small perturbations are naturally compensated, allowing motion to remain stable under bounded aleatoric disturbances $w \in \mathcal{W}$. From an evaluation perspective, this form of robustness can be captured by convergence or divergence measures (Bazzi and Sternad 2020), which quantify how trajectories evolve under perturbations. For example, the largest eigenvalue of the symmetric part of the system Jacobian measures the local rate of trajectory divergence or contraction, with negative values indicating stabilizing, disturbance-rejecting dynamics. Such measures can be incorporated into constrained optimization or planning formulations, e.g., as constraints on contact dynamics, to favor actions that induce convergent behavior and naturally funnel executions back toward desired outcomes.

7 Discussions

In this section, we discuss several complementary insights and implications that emerge from the robustness mechanisms surveyed throughout this paper.

7.1 Robustness and Performance Trade-offs

Manipulation performance is inherently context-dependent. A system may exhibit high performance in controlled laboratory settings, yet experience sharp degradation in unstructured real-world environments. Systems optimized for narrow conditions can achieve exceptional speed or precision, but often do so at the expense of robustness when those conditions change. In nature, star-nosed moles exemplify this trade-off: their specialized nasal appendages enable rapid prey detection in wetland tunnels, but such specialization may be fragile under environmental shifts (Catania and Kaas 1996). By contrast, some systems with moderate peak performance may trade efficiency or accuracy for robustness. Compliant robotic grippers, for example, adapt well to variations in object shape or mass, but typically sacrifice positional precision (Deimel and Brock 2013). These observations suggest that progress in manipulation should be evaluated not only by peak performance under ideal conditions, but by the ability to sustain reliable function under imperfect and unpredictable in-the-wild challenges.

7.2 Robustness and Related Concepts

Robustness often intersects with but differs from other fundamental concepts, including safety, stability, generalizability, etc. *Safety* concerns preventing harm to humans, the environment, or the robot itself, and is critical in domains such as autonomous driving, navigation, and human–robot interaction (Brunke et al. 2022). *Stability*, common in control and grasping, refers to maintaining or converging to a desired state under perturbations. Robustness extends beyond these notions, encompassing not just safe or stable operation but the broader pursuit of general task goals under uncertainty or variation.

Robustness and *generalizability* are also frequently discussed interchangeably, especially in robot learning, but they capture different aspects of system behavior. Generalizability concerns the scope over which a learned policy or model can be applied beyond its training conditions, such as new objects, scenes, embodiments, task instructions, or simulation-to-real transfer (Kroemer et al. 2021; Aljalbout et al. 2025; Gao et al. 2026). Robustness, in contrast, concerns the reliability of task achievement under a specified set or distribution of uncertainty and variation. Thus, generalization may serve as one mechanism for robustness when the deployment variations are covered by the generalized competence of the policy, but it is neither necessary nor sufficient: a controller may be robust within a narrow operational condition without generalizing broadly, while a generalist policy may cover many tasks yet remain sensitive to small perturbations, contact uncertainty, or distributional shifts not represented in its evaluation protocol.

7.3 Insights from beyond Manipulation Robustness

Robustness research in domains beyond manipulation offers valuable insights. As mentioned in Section 2, locomotion robustness is particularly relevant, as locomotion can be viewed as a form of "self-manipulation" (Johnson et al. 2016; Mason 2018). Techniques such as adversarial training for controller robustness (Shi et al. 2024), widely explored in locomotion, remain comparatively underutilized in manipulation and represent a promising direction. Similarly, biological inspirations such as gecko-inspired adhesives, originally studied in climbing and locomotion, have motivated gripper designs (Jiang et al. 2017). Despite these overlaps, manipulation presents unique challenges: unlike locomotion, which often reduces interaction to discrete foot–ground contacts, manipulation involves rich, multi-body interactions among the hand, object, and environment, with diverse contact types and combinatorial complexity that significantly complicate robustness analysis and design.

7.4 Beneficial Role of Perturbations

Finally, perturbations, typically viewed as detrimental, can in some cases enhance robustness. In granular media, small vibrations can break clogging arches when pouring sand into a funnel (To et al. 2001). The vibrating bowl feeder discussed in Section 5.2.2 provides another example, where controlled vibrations guide parts toward consistent configurations. Analogously in manipulation, slight, intentional perturbations can dislodge unstable contacts, reduce excessive contact forces, or guide objects out of shallow local minima. Everyday actions such as wiggling a Jenga block or flicking an omelette exemplify how controlled perturbations stabilize interactions rather than destabilize them. These observations highlight that robustness is not always achieved by tackling uncertainties or variations, but sometimes by strategically exploiting them.

8 Challenges and Open Problems

The frameworks and principles presented thus far provide a basis for understanding how robustness arises, yet they also make clear that essential scientific and engineering questions remain unanswered. This section identifies some of the open problems and discusses the challenges that must be overcome to achieve manipulation robustness under real-world challenges.

8.1 Benchmarking and Evaluation of Manipulation Robustness

Effective comparison of manipulation robustness across different methods requires well-designed benchmarks and consistent evaluation methods. Simulation-based benchmarking (Tao et al. 2024; James et al. 2020; Zhu et al. 2020) allows experiments under identical conditions and facilitates large-scale testing, yet faces the persistent challenge of the sim-to-real gap that limits physical realism and transferability, as well as being confined to a very narrow distribution of tasks and environments. Real-world benchmarking efforts, including standardized object sets (Calli et al. 2017), manipulation competitions (Correll et al. 2016; Kasper et al. 2012), cloud-based robotic platforms (Zahid and Pokorny 2024), and community-run distributed infrastructures (Chen et al. 2026; Atreya et al. 2025), aim to address this gap, yet each makes trade-offs among task diversity, scalability, accessibility, authenticity, or consistency. Future progress depends on balancing these aspects through advances in computation, communication, and standardized cross-community protocols. Most importantly, the existing simulation or real-world benchmarks *rarely* treat robustness as an explicit evaluation criterion, instead evaluating performance primarily under relatively fixed conditions. Fair comparison further requires robustness evaluation methods that quantify system performance. While analytical and empirical methods have been developed for specific manipulation tasks (Section 6), a unified framework that systematically captures real-world uncertainty or variation remains an open challenge.

8.2 Analytical and Empirical Robustness

Consistent with the analytical and empirical evaluation protocols discussed in Section 6, these two methodological traditions also appear broadly in robustness research. Analytical robustness methods reflect Plato's philosophy of abstraction, seeking general principles and guarantees through formal modeling and analysis; empirical methods follow Aristotle's practical spirit, emphasizing performance validated through data and experience. The field continues to debate whether progress in robotics lies in one paradigm or their synthesis (Amato et al. 2025), and robustness as a core robotic topic follows the same divide. Data offer robots experiential knowledge that models alone cannot provide, enabling robustness to emerge from interaction, much as humans develop know-how before know-why through sensorimotor learning. Yet unlike human toddlers, robots are not bound to start from scratch. They can inherit models and priors designed by humans, effectively knowing why before knowing how. Such models allow robots to adapt behavior in a data-efficient way while mitigating the limitations of purely data-driven methods, which are often sensitive to aleatoric uncertainty and distribution shift. Consider the game of Jenga: humans learn to play robustly not only through trial and error but also by leveraging internal models of perception, planning, and control. Similarly, the future of manipulation robustness likely lies in unifying data-driven learning and model-based reasoning—integrating know-how and know-why into robot systems.

8.3 Towards Human-Level Manipulation Robustness

Achieving human-level manipulation robustness remains a grand challenge in robotics. Evidence suggests that such robustness emerges from the integration of multiple complementary strategies, such as extensive lifetime experience, accurate world models grounded in physical reasoning, domain-specific embodiment, and rapid adaptation. For example, crows bend hooks to retrieve food, effectively co-designing hardware and control in the wild (Hunt and Gray 2004); such behavior cannot be explained by mechanics or control alone. Humans similarly integrate rich lifetime priors with online exploration and adaptation, inspiring approaches such as reinforcement learning bootstrapped with prior data (Luo et al. 2025). Yet, synthesizing all these strategies in a single robotic system remains rare and challenging. Other principles evident in animals and humans, such as redundancy in planning and sensing, are gaining attention as essential components for manipulation robustness. Reaching animal-level manipulation robustness already exceeds current robotic capabilities, but is prospective; attaining human-level robustness represents a far greater leap and the ultimate aspiration for robotic manipulation.

9 Conclusion

This review makes advances towards a systematic understanding of manipulation robustness by defining the concept, formulating the problem, and revisiting prior work across subfields. We surveyed representative mechanisms and evaluation methods, highlighting the core principles that enable robustness in robotic manipulation and providing insights for future practitioners interested in this topic. Despite decades of progress, achieving human-level manipulation robustness on robotic systems under real-world uncertainties and variations remains a central challenge. Bridging this gap will require not only improved algorithms and hardware, but also clearer formulations of robustness and more consistent evaluation methods. We hope this survey will help organize future research efforts toward robotic manipulation systems that routinely approach the robustness demonstrated by humans and animals in the physical world.

Acknowledgements

The project is partially funded by the European Commission under the Horizon Europe Framework Program project SoftEnable, grant number 101070600. The authors thank Oliver Brock, Matthew T. Mason, Yan Zhang, Yunke Ao, Rafael I. Cabral Muchacho, Shaohang Han, Haoyu Li, Zizhe Zhang, and Jinda Cui for helpful discussions, comments, or proofreading. Part of Fig. 1, 3-c were created using generative AI tools (ChatGPT/OpenAI) and subsequently reviewed and edited by the authors. These figures are original representative illustrations intended for conceptual explanation. The authors verified their technical accuracy and take full responsibility for the final content.

Declaration of conflicting interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

References

Achiam J, Held D, Tamar A and Abbeel P (2017) Constrained policy optimization. In: *International conference on machine learning*. PMLR, pp. 22–31.

Aljalbout E, Xing J, Romero A, Akinola I, Garrett CR, Heiden E, Gupta A, Hermans T, Narang Y, Fox D, Scaramuzza D and Ramos F (2025) The reality gap in robotics: Challenges, solutions, and best practices. *Annual Review of Control, Robotics, and Autonomous Systems* 9.

Aloimonos J, Weiss I and Bandyopadhyay A (1988) Active vision. *International journal of computer vision* 1(4): 333–356.

Amato NM, Hutchinson S, Garg A, Billard A, Rus D, Tedrake R, Park F and Goldberg K (2025) “data will solve robotics and automation: True or false?”: A debate. *Science Robotics* 10(105): eaea7897.

Amend JR, Brown E, Rodenberg N, Jaeger HM and Lipson H (2012) A positive pressure universal gripper based on the jamming of granular material. *IEEE transactions on robotics* 28(2): 341–350.

Ameperosa E, Collins JA, Jain M and Garg A (2025) Rocoda: Counterfactual data augmentation for data-efficient robot learning from demonstrations. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 13250–13256.

Andrychowicz OM, Baker B, Chociej M, Jozefowicz R, McGrew B, Pachocki J, Petron A, Plappert M, Powell G, Ray A, Schneider J, Sidor S, Tobin J, Welinder P, Weng L and Zaremba W (2020) Learning dexterous in-hand manipulation. *The International Journal of Robotics Research* 39(1): 3–20.

Atreya P, Pertsch K, Lee T, Kim MJ, Jain A, Kuramshin A, Neary C, Hu ES, Arora K, Ellis K, Macesanu L, Leonard M, Cho M, Aslan O, Dass S, Wang T, Yuan X, Gupta A, Jayaraman D, Berseth G, Daniilidis K, Mart'in-Mart'in R, Lee Y, Liang P, Finn C and Levine S (2025) Roboarena: Distributed real-world evaluation of generalist robot policies. In: *Conference on Robot Learning*. PMLR, pp. 336–364.

Autumn K (2006) Tokay gecko foot image. https://people.eecs.berkeley.edu/~ronf/Gecko/Interface-slide-adhesion/TokayFoot2-KA.jpg. Copyright © 2006 Kellar Autumn. Accessed: 2026-06-21.

Bajcsy R (1988) Active perception. *Proceedings of the IEEE* 76(8): 966–1005.

Baum M (2024) *Robustness in robotic and biological manipulation*. Technische Universitaet Berlin (Germany).

Bay H, Tuytelaars T and Van Gool L (2006) Surf: Speeded up robust features. In: *European conference on computer vision*. Springer, pp. 404–417.

Bazzi S and Sternad D (2020) Robustness in human manipulation of dynamically complex objects through control contraction metrics. *IEEE robotics and automation letters* 5(2): 2578–2585.

Berenson D, Srinivasa SS, Ferguson D and Kuffner JJ (2009) Manipulation planning on constraint manifolds. In: *2009 IEEE international conference on robotics and automation*. IEEE, pp. 625–632.

Berkeley AI Research (2019) Dexterous manipulation blog image. https://bair.berkeley.edu/static/blog/dex_manip/missing_img1.png. Image from BAIR blog on dexterous manipulation. Accessed: 2026-03-13.

Bhatt A, Sieler A, Puhlmann S and Brock O (2021) Surprisingly robust in-hand manipulation: An empirical study. *Robotics: Science and Systems XVII* .

Bhattacharya S, Ghrist R and Kumar V (2015) Persistent homology for path planning in uncertain environments. *IEEE Transactions on Robotics* 31(3): 578–590.

Bicchi A (1995) On the closure properties of robotic grasping. *The International Journal of Robotics Research* 14(4): 319–334.

Billard A and Kragic D (2019) Trends and challenges in robot manipulation. *Science* 364(6446): eaat8414.

Black K, Brown N, Driess D, Esmail A, Equi MR, Finn C, Fusai N, Groom L, Hausman K, Ichter B, Jakubczak S, Jones T, Ke L, Levine S, Li-Bell A, Mothukuri M, Nair S, Pertsch K, Shi LX, Smith L, Tanner J, Vuong Q, Walling A, Wang H and Zhilinsky U (2025) $\pi 0$: A Vision-Language-Action Flow Model for General Robot Control. *Proceedings of Robotics: Science and Systems* .

Blackmore L, Ono M and Williams BC (2011) Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics* 27(6): 1080–1094.

Bohg J, Hausman K, Sankaran B, Brock O, Kragic D, Schaal S and Sukhatme GS (2017) Interactive perception: Leveraging action in perception and perception in action. *IEEE Transactions on Robotics* 33(6): 1273–1291.

Bohg J, Morales A, Asfour T and Kragic D (2013) Data-driven grasp synthesis—a survey. *IEEE Transactions on robotics* 30(2): 289–309.

Braiek HB and Khomh F (2025) Machine learning robustness: A primer. In: *Trustworthy AI in Medical Imaging*. Elsevier, pp. 37–71.

Brohan A, Brown N, Carbajal J, Chebotar Y, Dabis J, Finn C, Gopalakrishnan K, Hausman K, Herzog A, Hsu J, Ibarz J, Ichter B, Irpan A, Jackson T, Jesmonth S, Joshi N, Julian R, Kalashnikov D, Kuang Y, Leal I, Lee KH, Levine S, Lu Y, Malla U, Manjunath D, Mordatch I, Nachum O, Parada C, Peralta J, Perez E, Pertsch K, Quiambao J, Rao K, Ryoo MS, Salazar G, Sanketi PR, Sayed K, Singh J, Sontakke S, Stone A, Tan C, Tran H, Vanhoucke V, Vega S, Vuong QH, Xia F, Xiao T, Xu P, Xu S, Yu T and Zitkovich B (2023) Rt-1: Robotics transformer for real-world control at scale. *Robotics: Science and Systems XIX* .

Brunke L, Greeff M, Hall AW, Yuan Z, Zhou S, Panerati J and Schoellig AP (2022) Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems* 5(1): 411–444.

Calli B, Singh A, Bruce J, Walsman A, Konolige K, Srinivasa S, Abbeel P and Dollar AM (2017) Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research* 36(3): 261–268.

Cao B, Zhang B, Zheng W, Zhou J, Lin Y and Chen Y (2022) Real-time, highly accurate robotic grasp detection utilizing transfer learning for robots manipulating fragile fruits with widely variable sizes and shapes. *Computers and electronics in agriculture* 200: 107254.

Catania KC and Kaas JH (1996) The unusual nose and brain of the star-nosed mole. *Bioscience* 46(8): 578–586.

Chappell J, Demery ZP, Arriola-Rios V and Sloman A (2012) How to build an information gathering and processing system: Lessons from naturally and artificially intelligent systems. *Behavioural Processes* 89(2): 179–186.

Chen Y, Kimble K, Adelson EH, Asfour T, Chanrungmaneekul P, Chitta S, Chitambar Y, Chen Z, Goldberg K, Kragic D, Li H, Li X, Li Y, Prather A, Pollard N, Roa-Garzon MA, Seney R, Sha S, Wang S, Xiang Y, Zhang K, Zhu Y and Hang K (2026) Manipulationnet: An infrastructure for benchmarking real-world robot manipulation with physical skill challenges and embodied multimodal reasoning. *arXiv preprint arXiv:2603.04363* .

Chen Z, Kiami S, Gupta A and Kumar V (2023) Genaug: Retargeting behaviors to unseen situations via generative augmentation. *Robotics: Science and Systems* .

Cheng R, Orosz G, Murray RM and Burdick JW (2019) End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In: *Proceedings of the AAAI conference on artificial intelligence*, volume 33. pp. 3387–3395.

Chi C, Xu Z, Feng S, Cousineau E, Du Y, Burchfiel B, Tedrake R and Song S (2023) Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research* : 02783649241273668.

Correll N, Bekris KE, Berenson D, Brock O, Causo A, Hauser K, Okada K, Rodriguez A, Romano JM and Wurman PR (2016) Analysis and observations from the first amazon picking challenge. *IEEE Transactions on Automation Science and Engineering* 15(1): 172–188.

Cui J and Trinkle J (2021) Toward next-generation learned robot manipulation. *Science robotics* 6(54): eabd9461.

Dafle NC, Rodriguez A, Paolini R, Tang B, Srinivasa SS, Erdmann M, Mason MT, Lundberg I, Staab H and Fuhlbrigge T (2014) Extrinsic dexterity: In-hand manipulation with external forces. In: *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1578–1585.

Dalal G, Dvijotham K, Vecerik M, Hester T, Paduraru C and Tassa Y (2018) Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757* .

Deimel R and Brock O (2013) A compliant hand based on a novel pneumatic actuator. In: *2013 IEEE international conference on robotics and automation*. IEEE, pp. 2047–2053.

Dhonthi A, Schillinger P, Rozo L and Nardi D (2021) Study of signal temporal logic robustness metrics for robotic tasks optimization. *arXiv preprint arXiv:2110.00339* .

Dimeas F and Aspragathos N (2015) Reinforcement learning of variable admittance control for human-robot co-manipulation. In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 1011–1016.

Dogar M, Hsiao K, Ciocarlie M and Srinivasa S (2012) Physics-based grasp planning through clutter. *Robotics: Science and System* : 57–64.

Dong Y, Cheng X and Pokorny FT (2024) Characterizing manipulation robustness through energy margin and caging analysis. *IEEE Robotics and Automation Letters* 9(9): 7525–7532.

Donzé A and Maler O (2010) Robust satisfaction of temporal logic over real-valued signals. In: *International conference on formal modeling and analysis of timed systems*. Springer, pp. 92–106.

Drake SH (1978) *Using compliance in lieu of sensory feedback for automatic assembly.* PhD Thesis, Massachusetts Institute of Technology.

Driess D, Ha JS and Toussaint M (2021) Learning to solve sequential physical reasoning problems from a scene image. *The International Journal of Robotics Research* 40(12-14): 1435–1466.

Duan Y, Andrychowicz M, Stadie B, Jonathan Ho O, Schneider J, Sutskever I, Abbeel P and Zaremba W (2017) One-shot imitation learning. *Advances in neural information processing systems* 30.

Eisner B, Yang Y, Davchev T, Vecerik M, Scholz J and Held D (2024) Deep se (3)-equivariant geometric reasoning for precise placement tasks. In: *The Twelfth International Conference on Learning Representations*.

Erdmann MA and Mason MT (2002) An exploration of sensorless manipulation. *IEEE Journal on Robotics and Automation* 4(4): 369–379.

Fenchel T (1980) Suspension feeding in ciliated protozoa: functional response and particle size selection. *Microbial Ecology* 6(1): 1–11.

Ferrari C and Canny J (1992) Planning optimal grasps. In: *Proceedings., 1992 IEEE International Conference on Robotics and Automation, 1992.*, volume 3. IEEE, pp. 2290–2295.

Firoozi R, Tucker J, Tian S, Majumdar A, Sun J, Liu W, Zhu Y, Song S, Kapoor A, Hausman K, Ichter B, Driess D, Wu J, Lu C and Schwager M (2025) Foundation models in robotics: Applications, challenges, and the future. *The International Journal of Robotics Research* 44(5): 701–739.

Flanagan JR, Bowman MC and Johansson RS (2006) Control strategies in object manipulation tasks. *Current opinion in neurobiology* 16(6): 650–659.

Florence P, Manuelli L and Tedrake R (2019) Self-supervised correspondence in visuomotor policy learning. *IEEE Robotics and Automation Letters* 5(2): 492–499.

Florence PR, Manuelli L and Tedrake R (2018) Dense object nets: Learning dense visual object descriptors by and for robotic manipulation. In: *Conference on Robot Learning*. PMLR, pp. 373–385.

Foster-Turley P and Markowitz H (1982) A captive behavioral enrichment study with asian small-clawed river otters (aonyx cinerea). *Zoo biology* 1(1): 29–43.

Fox D (2020) Grasping. Lecture slides, CSE 478: Robot Learning, University of Washington. URL https://courses.cs.washington.edu/courses/cse478/20wi/site/resources/lec23_grasping.pdf.

Fox M, Howey R and Long D (2006) Exploration of the robustness of plans. In: *AAAI*. pp. 834–839.

Fruit Hunters (2026) Banana variety box. https://fruithunters.com/products/banana-variety-box. Accessed: 2026-03-13.

Fu L, Huang H, Datta G, Chen LY, Panitch WCH, Liu F, Li H and Goldberg K (2025) In-context imitation learning via next-token prediction. *2025 IEEE international conference on robotics and automation (ICRA)* .

Fujimoto S and Gu SS (2021) A minimalist approach to offline reinforcement learning. *Advances in neural information processing systems* 34: 20132–20145.

Gao J, Belkhale S, Dasari S, Balakrishna A, Shah D and Sadigh D (2026) A taxonomy for evaluating generalist robot manipulation policies. *IEEE Robotics and Automation Letters* .

Gelblum A, Pinkoviezky I, Fonio E, Ghosh A, Gov N and Feinerman O (2015) Ant groups optimally amplify the effect of transiently informed individuals. *Nature communications* 6(1): 7729.

Ghazi-Zahedi K, Deimel R, Montúfar G, Wall V and Brock O (2017) Morphological computation: the good, the bad, and the ugly. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 464–469.

Goertz RC (1952) Fundamentals of general-purpose remote manipulators. *Nucleonics* : 36–42.

Hadjiosif AM and Smith MA (2015) Flexible control of safety margins for action based on environmental variability. *Journal of Neuroscience* 35(24): 9106–9121.

Hafner D, Lillicrap T, Fischer I, Villegas R, Ha D, Lee H and Davidson J (2019) Learning latent dynamics for planning from pixels. In: *International conference on machine learning*. PMLR, pp. 2555–2565.

Hafner D, Pasukonis J, Ba J and Lillicrap T (2023) Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104* .

Haldar S, Mathur V, Yarats D and Pinto L (2023) Watch and match: Supercharging imitation with regularized optimal transport. In: *Conference on Robot Learning*. PMLR, pp. 32–43.

Hansen N, Wang X and Su H (2022) Temporal difference learning for model predictive control. In: *International Conference on Machine Learning, PMLR*.

Hasson CJ, Shen T and Sternad D (2012) Energy margins in dynamic object manipulation. *Journal of Neurophysiology* 108(5): 1349–1365.

Hendrycks D and Dietterich T (2019) Benchmarking neural network robustness to common corruptions and perturbations. In: *International Conference on Learning Representations*.

Herzog A, Rao K, Hausman K, Lu Y, Wohlhart P, Yan M, Lin J, Arenas MG, Xiao T, Kappler D, Ho D, Rettinghouse J, Chebotar Y, Lee KH, Gopalakrishnan K, Julian R, Li A, Fu CK, Wei B, Ramesh S, Holden K, Kleiven K, Rendleman D, Kirmani S, Bingham J, Weisz J, Xu Y, Lu W, Bennice M, Fong C, Do D, Lam J, Bai Y, Holson B, Quinlan M, Brown N, Kalakrishnan M, Ibarz J, Pastor P and Levine S (2023) Deep rl at scale: Sorting waste in office buildings with a fleet of mobile manipulators. *Robotics: Science and Systems (RSS)* .

Hogan N (1984) Impedance control of industrial robots. *Robotics and computer-integrated manufacturing* 1(1): 97–113.

Holmes P, Kousik S, Zhang B, Raz D, Barbalata C, Johnson-Roberson M and Vasudevan R (2020) Reachable sets for safe, real-time manipulator trajectory design. *Robotics: Science and Systems* .

Homberg BS, Katzschmann RK, Dogar MR and Rus D (2019) Robust proprioceptive grasping with a soft robot hand. *Autonomous robots* 43(3): 681–696.

Hou Y, Liu Z, Chi C, Cousineau E, Kuppuswamy N, Feng S, Burchfiel B and Song S (2025) Adaptive compliance policy: Learning approximate compliance for diffusion guided control. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 4829–4836.

Howard WS and Kumar V (1996) On the stability of grasped objects. *IEEE transactions on robotics and automation* 12(6): 904–917.

Huang W, Wang C, Li Y, Zhang R and Fei-Fei L (2025) Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In: *Conference on Robot Learning*. PMLR, pp. 4573–4602.

Huang Y, Conkey A and Hermans T (2023) Planning for multi-object manipulation with graph neural network relational classifiers. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1822–1829.

Hunt GR and Gray RD (2004) The crafting of hook tools by wild new caledonian crows. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 271(suppl_3): S88–S90.

Ichter B and Pavone M (2019) Robot motion planning in learned latent spaces. *IEEE Robotics and Automation Letters* 4(3): 2407–2414.

Izatt G, Mirano G, Adelson E and Tedrake R (2017) Tracking objects with point clouds from vision and touch. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 4000–4007.

James S, Ma Z, Arrojo DR and Davison AJ (2020) Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters* 5(2): 3019–3026.

Jankowski J, Brudermüller L, Hawes N and Calinon S (2024) Robust pushing: Exploiting quasi-static belief dynamics and contact-informed optimization. *The International Journal of Robotics Research* : 02783649251318046.

Jian P, Yang C, Guo D, Liu H and Sun F (2021) Adversarial skill learning for robust manipulation. In: *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 2555–2561.

Jiang H, Hawkes EW, Fuller C, Estrada MA, Suresh SA, Abcouwer N, Han AK, Wang S, Ploch CJ, Parness A and Cutkosky MR (2017) A robotic device using gecko-inspired adhesives can grasp and manipulate large objects in microgravity. *Science Robotics* 2(7): eaan4545.

Jiang Y, Yu M, Zhu X, Tomizuka M and Li X (2025) Robust model-based in-hand manipulation with integrated real-time motion-contact planning and tracking. *arXiv preprint arXiv:2505.04978* .

Johannink T, Bahl S, Nair A, Luo J, Kumar A, Loskyll M, Ojea JA, Solowjow E and Levine S (2019) Residual reinforcement learning for robot control. In: *2019 international conference on robotics and automation (ICRA)*. IEEE, pp. 6023–6029.

Johnson AM, Burden SA and Koditschek DE (2016) A hybrid systems model for simple manipulation and self-manipulation systems. *The International Journal of Robotics Research* 35(11): 1354–1392.

Kaelbling LP and Lozano-Pérez T (2013) Integrated task and motion planning in belief space. *The International Journal of Robotics Research* 32(9-10): 1194–1227.

Kamijo T, Beltran-Hernandez CC and Hamaya M (2024) Learning variable compliance control from a few demonstrations for bimanual robot with haptic feedback teleoperation system. In: *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 12663–12670.

Kang X, Tian T, Lee SW, Huang B, Li Y and Kuo YL (2026) Learning force-regulated manipulation with a low-cost tactile-force-controlled gripper. *arXiv preprint arXiv:2602.10013* .

Kapoor P, Balakrishnan A and Deshmukh JV (2020) Model-based reinforcement learning from signal temporal logic specifications. *arXiv preprint arXiv:2011.04950* .

Kasper A, Xue Z and Dillmann R (2012) The kit object models database: An object model database for object recognition, localization and manipulation in service robotics. *The International Journal of Robotics Research* 31(8): 927–934.

Katz D and Brock O (2008) Manipulating articulated objects with interactive perception. In: *2008 IEEE International Conference on Robotics and Automation*. IEEE, pp. 272–277.

Ke L, Wang J, Bhattacharjee T, Boots B and Srinivasa S (2021) Grasping with chopsticks: Combating covariate shift in model-free imitation learning for fine manipulation. In: *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 6185–6191.

Kelly M, Sidrane C, Driggs-Campbell K and Kochenderfer MJ (2019) Hg-dagger: Interactive imitation learning with human experts. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 8077–8083.

Kim MJ, Pertsch K, Karamcheti S, Xiao T, Balakrishna A, Nair S, Rafailov R, Foster EP, Sanketi PR, Vuong Q, Kollar T, Burchfiel B, Tedrake R, Sadigh D, Levine S, Liang P and Finn C (2025) Openvla: An open-source vision-language-action model. In: *Conference on Robot Learning*. PMLR, pp. 2679–2713.

Kitano H (2004) Biological robustness. *Nature Reviews Genetics* 5(11): 826–837.

Kostrikov I, Nair A and Levine S (2021) Offline reinforcement learning with implicit q-learning. In: *International Conference on Learning Representations*.

Koval MC, Pollard NS and Srinivasa SS (2016) Pre-and post-contact policy decomposition for planar contact manipulation under uncertainty. *The International Journal of Robotics Research* 35(1-3): 244–264.

Kress-Gazit H, Hashimoto K, Kuppuswamy N, Shah P, Horgan P, Richardson G, Feng S and Burchfiel B (2024) Robot learning as an empirical science: Best practices for policy evaluation. *arXiv preprint arXiv:2409.09491* .

Krizhevsky A, Sutskever I and Hinton GE (2012) Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25.

Kroemer O, Niekum S and Konidaris G (2021) A review of robot learning for manipulation: Challenges, representations, and algorithms. *Journal of machine learning research* 22(30): 1–82.

Kumar A, Zhou A, Tucker G and Levine S (2020) Conservative q-learning for offline reinforcement learning. *Advances in neural information processing systems* 33: 1179–1191.

Laskey M, Lee J, Fox R, Dragan A and Goldberg K (2017) Dart: Noise injection for robust imitation learning. In: *Conference on robot learning*. PMLR, pp. 143–156.

LeCun Y, Bottou L, Bengio Y and Haffner P (2002) Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11): 2278–2324.

Levine S, Finn C, Darrell T and Abbeel P (2016) End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research* 17(39): 1–40.

Levine S, Kumar A, Tucker G and Fu J (2020) Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643* .

Li H, Zhang Y, Zhu J, Wang S, Lee MA, Xu H, Adelson E, Fei-Fei L, Gao R and Wu J (2023) See, hear, and feel: Smart sensory fusion for robotic manipulation. In: *Conference on Robot Learning*. PMLR, pp. 1368–1378.

Li R, Jabri A, Darrell T and Agrawal P (2020) Towards practical multi-object manipulation using relational reinforcement learning. In: *2020 ieee international conference on robotics and automation (icra)*. IEEE, pp. 4051–4058.

Li X, Vasile CI and Belta C (2017) Reinforcement learning with temporal logic rewards. In: *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 3834–3839.

Liang G, Ijspeert AJ, Yim M and Lam TL (2026) Modular reconfigurable robots: Toward on-demand multifunctional applications. *Science Robotics* 11(111): eadz1999.

Liconti D, Zhou Y, Toshimitsu Y, Hinchet R and Katzschmann RK (2026) A benchmark of dexterity for anthropomorphic robotic hands. *arXiv preprint arXiv:2604.09294* .

Lin Y, Wang AS, Undersander E and Rai A (2022) Efficient and interpretable robot manipulation with graph neural networks. *IEEE Robotics and Automation Letters* 7(2): 2740–2747.

Liu H, Dass S, Martín-Martín R and Zhu Y (2024) Model-based runtime monitoring with interactive imitation learning. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 4154–4161.

Liu H, Nasiriany S, Zhang L, Bao Z and Zhu Y (2022) Robot learning on the job: Human-in-the-loop autonomy and learning during deployment. *The International Journal of Robotics Research* : 02783649241273901.

Lowe DG (2004) Distinctive image features from scale-invariant keypoints. *International journal of computer vision* 60(2): 91–110.

Lozano-Perez T, Mason MT and Taylor RH (1984) Automatic synthesis of fine-motion strategies for robots. *The International Journal of Robotics Research* 3(1): 3–24.

Lu H, Dong Y, Weng Z, Pokorny F, Lundell J and Kragic D (2025) Grasping a handful: Sequential multi-object dexterous grasp generation. *IEEE Robotics and Automation Letters* .

Luo J, Hu Z, Xu C, Tan YL, Berg J, Sharma A, Schaal S, Finn C, Gupta A and Levine S (2024) Serl: A software suite for sample-efficient robotic reinforcement learning. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 16961–16969.

Luo J, Xu C, Wu J and Levine S (2025) Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *Science Robotics* 10(105): eads5033.

Madry A, Makelov A, Schmidt L, Tsipras D and Vladu A (2018) Towards deep learning models resistant to adversarial attacks. *International Conference on Learning Representations* .

Mahler J, Liang J, Niyaz S, Laskey M, Doan R, Liu X, Aparicio J and Goldberg K (2017) Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. *Robotics: Science and Systems XIII* .

Mahler J, Pokorny FT, Niyaz S and Goldberg K (2018) Synthesis of energy-bounded planar caging grasps using persistent homology. *IEEE Transactions on Automation Science and Engineering* 15(3): 908–918.

Maler O and Nickovic D (2004) Monitoring temporal properties of continuous signals. In: *International symposium on formal techniques in real-time and fault-tolerant systems*. Springer, pp. 152–166.

Mandlekar A, Nasiriany S, Wen B, Akinola I, Narang Y, Fan L, Zhu Y and Fox D (2023) Mimicgen: A data generation system for scalable robot learning using human demonstrations. In: *Conference on Robot Learning*. PMLR, pp. 1820–1864.

Mandlekar A, Xu D, Martín-Martín R, Zhu Y, Fei-Fei L and Savarese S (2020) Human-in-the-loop imitation learning using remote teleoperation. *arXiv preprint arXiv:2012.06733* .

Mann J, Sargeant BL, Watson-Capps JJ, Gibson QA, Heithaus MR, Connor RC and Patterson E (2008) Why do dolphins carry sponges? *PloS one* 3(12): e3868.

Manuelli L, Gao W, Florence P and Tedrake R (2019) kpam: Keypoint affordances for category-level robotic manipulation. In: *The International Symposium of Robotics Research*. Springer, pp. 132–157.

Martín-Martín R (2018) *Leveraging problem structure in interactive perception for robot manipulation of constrained mechanisms*. Technische Universitaet Berlin (Germany).

Mason M (1985) The mechanics of manipulation. In: *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, volume 2. IEEE, pp. 544–548.

Mason MT (2007) Compliance and force control for computer controlled manipulators. *IEEE Transactions on Systems, Man, and Cybernetics* 11(6): 418–432.

Mason MT (2012) Creation myths: The beginnings of robotics research. *IEEE robotics & automation magazine* 19(2): 72–77.

Mason MT (2018) Toward robotic manipulation. *Annual Review of Control, Robotics, and Autonomous Systems* 1(1): 1–28.

Mayne DQ, Seron MM and Raković SV (2005) Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica* 41(2): 219–224.

Meng Y and Fan C (2023) Signal temporal logic neural predictive control. *IEEE Robotics and Automation Letters* 8(11): 7719–7726.

Mishra UA, Xue S, Chen Y and Xu D (2023) Generative skill chaining: Long-horizon skill planning with diffusion models. In: *Conference on Robot Learning*. PMLR, pp. 2905–2925.

Moos J, Hansel K, Abdulsamad H, Stark S, Clever D and Peters J (2022) Robust reinforcement learning: A review of foundations and recent advances. *Machine Learning and Knowledge Extraction* 4(1): 276–315.

Morrison D, Corke P and Leitner J (2019) Multi-view picking: Next-best-view reaching for improved grasping in clutter. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 8762–8768.

Nachum O, Gu SS, Lee H and Levine S (2018) Data-efficient hierarchical reinforcement learning. *Advances in neural information processing systems* 31.

Nakamura K, Peters L and Bajcsy A (2025) Generalizing safety beyond collision-avoidance via latent-space reachability analysis. *arXiv preprint arXiv:2502.00935* .

Nubert J, Köhler J, Berenz V, Allgöwer F and Trimpe S (2020) Safe and fast tracking on a robot manipulator: Robust mpc and neural network control. *IEEE Robotics and Automation Letters* 5(2): 3050–3057.

OnlineDeliveryin (2024) Mixed fruits with basket (2 kg). https://www.onlinedelivery.in/mixed-fruits-with-basket-2-kg. Accessed: 2026-03-13.

Open X-Embodiment Collaboration (2024) Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 6892–6903.

OpenAI, Akkaya I, Andrychowicz M, Chociej M, Litwin M, McGrew B, Petron A, Paino A, Plappert M, Powell G, Ribas R, Schneider J, Tezak N, Tworek J, Welinder P, Weng L, Yuan Q, Zaremba W and Zhang L (2019) Solving rubik's cube with a robot hand. *arXiv preprint arXiv:1910.07113* .

Oquab M, Darcet T, Moutakanni T, Vo H, Szafraniec M, Khalidov V, Fernandez P, Haziza D, Massa F, El-Nouby A, Assran M, Ballas N, Galuba W, Howes R, Huang PY, Li SW, Misra I, Rabbat M, Sharma V, Synnaeve G, Xu H, J'egou H, Mairal J, Labatut P, Joulin A and Bojanowski P (2024) Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal* .

Pagoli A, Chapelle F, Corrales JA, Mezouar Y and Lapusta Y (2021) A soft robotic gripper with an active palm and reconfigurable fingers for fully dexterous in-hand manipulation. *IEEE Robotics and Automation Letters* 6(4): 7706–7713.

Pan C, Anantharaman G, Huang NC, Jin C, Pfrommer D, Yuan C, Permenter F, Qu G, Boffi N, Shi G and Simchowitz M (2026) Much ado about noising: Dispelling the myths of generative robotic control. *International Conference on Learning Representations* .

Paul C (2006) Morphological computation: A basis for the analysis of morphology and control requirements. *Robotics and Autonomous Systems* 54(8): 619–630.

Pereira GA, Campos MF and Kumar V (2004) Decentralized algorithms for multi-robot manipulation via caging. *The International Journal of Robotics Research* 23(7-8): 783–795.

Peshkin MA and Sanderson AC (2002) Planning robotic manipulation strategies for workpieces that slide. *IEEE Journal on Robotics and Automation* 4(5): 524–531.

Pinto L, Davidson J and Gupta A (2017a) Supervision via competition: Robot adversaries for learning tasks. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1601–1608.

Pinto L, Davidson J, Sukthankar R and Gupta A (2017b) Robust adversarial reinforcement learning. In: *International conference on machine learning*. PMLR, pp. 2817–2826.

Platt R, Tedrake R, Kaelbling L and Lozano-Perez T (2010) Belief space planning assuming maximum likelihood observations. *Robotics: Science and Systems VI* .

Pokorny FT and Kragic D (2015) Data-driven topological motion planning with persistent cohomology. *2015 Robotics: Science and Systems Conference* 11.

Polverini MP, Nicolis D, Zanchettin AM and Rocco P (2017) Implicit robot force control based on set invariance. *IEEE Robotics and Automation Letters* 2(3): 1288–1295.

Qin Z, Fang K, Zhu Y, Fei-Fei L and Savarese S (2020) Keto: Learning keypoint representations for tool manipulation. In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 7278–7285.

Rajeswaran A, Kumar V, Gupta A, Vezzani G, Schulman J, Todorov E and Levine S (2018) Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *Robotics: Science and Systems XIV* .

Roa MA and Suárez R (2015) Grasp quality measures: review and performance. *Autonomous robots* 38(1): 65–88.

Rodriguez A (2021) The unstable queen: Uncertainty, mechanics, and tactile feedback. *Science Robotics* 6(54): eabi4667.

Rodriguez A, Mason MT and Ferry S (2012) From caging to grasping. *The International Journal of Robotics Research* 31(7): 886–900.

Ross S, Gordon G and Bagnell D (2011) A reduction of imitation learning and structured prediction to no-regret online learning. In: *Proceedings of the fourteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings, pp. 627–635.

Rublee E, Rabaud V, Konolige K and Bradski G (2011) Orb: An efficient alternative to sift or surf. In: *2011 International conference on computer vision*. Ieee, pp. 2564–2571.

Ryu H, Lee Hi, Lee JH and Choi J (2022) Equivariant descriptor fields: Se (3)-equivariant energy-based models for end-to-end visual robotic manipulation learning. *arXiv preprint arXiv:2206.08321* .

Saha M and Isto P (2007) Manipulation planning for deformable linear objects. *IEEE Transactions on Robotics* 23(6): 1141–1150.

Sankar S, Cheng WY, Zhang J, Slepyan A, Iskarous MM, Greene RJ, DeBrabander R, Chen J, Gupta A and Thakor NV (2025) A natural biomimetic prosthetic hand with neuromorphic tactile sensing for precise and compliant grasping. *Science Advances* 11(10): eadr9300.

Shi F, Zhang C, Miki T, Lee J, Hutter M and Coros S (2024) Rethinking robustness assessment: Adversarial attacks on learning-based quadrupedal locomotion controllers. *Robotics: Science and System XX* .

Shintake J, Cacucciolo V, Floreano D and Shea H (2018) Soft robotic grippers. *Advanced materials* 30(29): 1707035.

Shirai Y, Jha DK, Raghunathan A and Romeres D (2022) Chance-constrained optimization in contact-rich systems for robust manipulation. *arXiv preprint arXiv:2203.02616* .

Simchowitz M, Pfrommer D and Jadbabaie A (2025) The pitfalls of imitation learning when actions are continuous. *arXiv preprint arXiv:2503.09722* .

Siméon T, Laumond JP, Cortés J and Sahbani A (2004) Manipulation planning with probabilistic roadmaps. *The International Journal of Robotics Research* 23(7-8): 729–746.

Simeonov A, Du Y, Tagliasacchi A, Tenenbaum JB, Rodriguez A, Agrawal P and Sitzmann V (2022) Neural descriptor fields: Se (3)-equivariant object representations for manipulation. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 6394–6400.

Skogestad S and Postlethwaite I (2005) *Multivariable feedback control: analysis and design*. john Wiley & sons.

Song S, Drotlef DM, Majidi C and Sitti M (2017) Controllable load sharing for soft adhesive interfaces on three-dimensional surfaces. *Proceedings of the National Academy of Sciences* 114(22): E4344–E4353.

Sternad D and Hasson CJ (2016) Predictability and robustness in the manipulation of dynamically complex objects. *Progress in Motor Control: Theories and Translations* : 55–77.

Suh HT, Pang T, Zhao T and Tedrake R (2025) Dexterous contact-rich manipulation via the contact trust region. *The International Journal of Robotics Research* : 02783649251398875.

Sun Z and Song S (2026) Latent policy barrier: Learning robust visuomotor policies by staying in-distribution. *Advances in Neural Information Processing Systems* 38: 174280–174305.

Sutton RS, Precup D and Singh S (1998) Intra-option learning about temporally abstract actions. In: *ICML*, volume 98. pp. 556–564.

Takano R, Oyama H and Yamakita M (2021) Continuous optimization-based task and motion planning with signal temporal logic specifications for sequential manipulation. In: *2021 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 8409–8415.

Tao S, Xiang F, Shukla A, Qin Y, Hinrichsen X, Yuan X, Bao C, Lin X, Liu Y, Chan Tk, Gao Y, Li X, Mu T, Xiao N, Gurha A, Huang Z, Calandra R, Chen R, Luo S and Su H (2024) Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai. *arXiv preprint arXiv:2410.00425* .

Tedrake R (2009) Lqr-trees: Feedback motion planning on sparse randomized trees. *Robotics: Science and Systems V* .

To K, Lai PY and Pak H (2001) Jamming of granular flow in a two-dimensional hopper. *Physical review letters* 86(1): 71.

Tobin J, Fong R, Ray A, Schneider J, Zaremba W and Abbeel P (2017) Domain randomization for transferring deep neural networks from simulation to the real world. In: *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, pp. 23–30.

TRI LBM Team, Barreiros J, Beaulieu A, Bhat A, Cory R, Cousineau E, Dai H, Fang CH, Hashimoto K, Irshad MZ, Itkina M, Kuppuswamy N, Lee KH, Liu K, McConachie D, McMahon I, Nishimura H, Phillips-Grafflin C, Richter C, Shah P, Srinivasan K, Wulfe B, Xu C, Zhang M, Alspach A, Angeles M, Arora K, Guizilini VC, Castro A, Chen D, Chu TS, Creasey S, Curtis S, Denitto R, Dixon E, Dusel E, Ferreira M, Goncalves A, Gould G, Guoy D, Gupta S, Han X, Hatch K, Hathaway B, Henry A, Hochsztein H, Horgan P, Iwase S, Jackson D, Karamcheti S, Keh S, Masterjohn J, Mercat J, Miller P, Mitiguy P, Nguyen T, Nimmer J, Noguchi Y, Ong R, Onol A, Pfannenstiehl O, Poyner R, Rocha LPM, Richardson G, Rodriguez C, Seale D, Sherman M, Smith-Jones M, Tago D, Tokmakov P, Tran M, Hoorick BV, Vasiljevic I, Zakharov S, Zolotas M, Ambrus R, Fetzer-Borelli K, Burchfiel B, Kress-Gazit H, Feng S, Ford S and Tedrake R (2026) A careful examination of large behavior models for multitask dexterous manipulation. *Science Robotics* 11(113): eaea6201.

Valassakis E, Papagiannis G, Di Palo N and Johns E (2022) Demonstrate once, imitate immediately (dome): Learning visual servoing for one-shot imitation learning. In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 8614–8621.

Varnai P and Dimarogonas DV (2020) On robustness metrics for learning stl tasks. In: *2020 American Control Conference (ACC)*. IEEE, pp. 5394–5399.

Wang D, Hart S, Surovik D, Kelestemur T, Huang H, Zhao H, Yeatman M, Wang J, Walters R and Platt R (2025a) Equivariant diffusion policy. In: *Conference on Robot Learning*. PMLR, pp. 48–69.

Wang G, Ren K, Morgan AS and Hang K (2025b) Caging in time: A framework for robust object manipulation under uncertainties and limited robot perception. *The International Journal of Robotics Research* : 02783649251343926.

Xie H, Sun M, Fan X, Lin Z, Chen W, Wang L, Dong L and He Q (2019) Reconfigurable magnetic microrobot swarm: Multimode transformation, locomotion, and manipulation. *Science robotics* 4(28): eaav8006.

Xiong H, Xu X, Wu J, Hou Y, Bohg J and Song S (2025) Vision in action: Learning active perception from human demonstrations. In: *Conference on Robot Learning*. PMLR, pp. 5450–5463.

Xu K, Hu Z, Doshi R, Rovinsky A, Kumar V, Gupta A and Levine S (2023) Dexterous manipulation from images: Autonomous real-world rl via substep guidance. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 5938–5945.

Xu K, Huang H, Shi Y, Li H, Long P, Caichen J, Sun W and Chen B (2015) Autoscanning for coupled scene reconstruction and proactive object analysis. *ACM Transactions on Graphics (TOG)* 34(6): 1–14.

Xu X, Hou Y, Liu Z and Song S (2026) Compliant residual dagger: Improving real-world contact-rich manipulation with human corrections. *The Thirty-ninth Annual Conference on Neural Information Processing Systems* .

Xue H, Ren J, Chen W, Zhang G, Fang Y, Gu G, Xu H and Lu C (2025a) Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation. *Robotics: Science and Systems* .

Xue Z, Deng S, Chen Z, Wang Y, Yuan Z and Xu H (2025b) Demogen: Synthetic demonstration generation for data-efficient visuomotor policy learning. *Robotics: Science and Systems* .

Yim M, Duff DG and Roufas KD (2000) Polybot: a modular reconfigurable robot. In: *Proceedings 2000 ICRA. millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, volume 1. IEEE, pp. 514–520.

Yu T, Xiao T, Stone A, Tompson J, Brohan A, Wang S, Singh J, Tan C, M D, Peralta J, Ichter B, Hausman K and Xia F (2023) Scaling robot learning with semantically imagined experience. *Robotics: Science and Systems* .

Zahid M and Pokorny FT (2024) Cloudgripper: An open source cloud robotics testbed for robotic manipulation research, benchmarking and data collection at scale. In: *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 12076–12082.

Zeng A, Song S, Yu KT, Donlon E, Hogan FR, Bauza M, Ma D, Taylor O, Liu M, Romo E, Fazeli N, Alet F, Dafle NC, Holladay R, Morona I, Nair PQ, Green D, Taylor I, Liu W, Funkhouser T and Rodriguez A (2022) Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. *The International Journal of Robotics Research* 41(7): 690–705.

Zhang X, Chang M, Kumar P and Gupta S (2024) Diffusion meets dagger: Supercharging eye-in-hand imitation learning. In: *Robotics science and systems*. Robotics science and systems.

Zhou A, Kim MJ, Wang L, Florence P and Finn C (2023a) Nerf in the palm of your hand: Corrective augmentation for robotics via novel-view synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 17907–17917.

Zhou G, Ke L, Srinivasa S, Gupta A, Rajeswaran A and Kumar V (2023b) Real world offline reinforcement learning with realistic data source. In: *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 7176–7183.

Zhu Y, Stone P and Zhu Y (2022) Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation. *IEEE Robotics and Automation Letters* 7(2): 4126–4133.

Zhu Y, Wong J, Mandlekar A, Martín-Martín R, Joshi A, Nasiriany S and Zhu Y (2020) robosuite: A modular simulation framework and benchmark for robot learning. *arXiv preprint arXiv:2009.12293* .

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