International Journal of Control, Automation, and Systems 24() (2026) 726-753
http://dx.doi.org/10.1007/s12555-026-00082-x ISSN:1598-6446 eISSN:2005-4092
http://www.springer.com/12555

Physics-Informed Causal Reasoning in Physical AI: A Review on Modeling Non-Stationary Environments for Safety-Critical Control

Sung Kwon On, Sukchul Jeong, Jungjin Lee and Insoon Yang*

Abstract: Non-stationary physical environments—arising from changes in friction, payload, contact conditions, human intervention, and system degradation—remain a central obstacle to deploying learning-enabled controllers in safety-critical settings. While substantial progress has been made in learning-based control and reinforcement learning, many failures under distribution shift can be traced back to a modeling gap: the inability to infer and track the *physical causes* that deform system dynamics and safety constraints over time. This review consolidates recent advances in *physics-informed causal reasoning and modeling* for non-stationary environments, with an emphasis on (*i*) inferring safety-relevant physical context from observations and interaction (action–reaction), (*ii*) building context-modulated dynamics models with physics priors injected via losses, architectures, or explicit constraints, and (*iii*) quantifying and calibrating uncertainty in *actionable* forms that support downstream safety mechanisms, including constraint tightening, risk-sensitive formulations, and distributionally robust reasoning. We propose a unified taxonomy that disentangles *task-relevant state* from *safety-critical physical context* and highlights failure modes induced by appearance–physics mismatch and non-informative cues. We organize the literature along axes of non-stationarity types, context identification paradigms (passive vs. active), and physics-prior injection mechanisms. Beyond summarizing methods, we provide an evaluation checklist and reporting protocol to improve comparability across studies, and conclude with open problems and a research agenda toward deployment-ready *physical AI systems grounded in real-world physics*.

Keywords: Physical AI, physics-informed modeling, causal reasoning, non-stationary environments, safety-critical control.

1. INTRODUCTION

Learning-enabled control systems have recently demonstrated impressive performance in robotics and autonomous systems [1]; however, their deployment in real-world, safety-critical settings remains fundamentally limited by non-stationary physical environments [2, 3]. In safety-critical control, the objective is not only to achieve task performance but also to guarantee that system trajectories remain within admissible safety constraints, such as collision avoidance and force limits.

Changes in friction, payload, contact conditions, human intervention, and long-term system degradation can continuously deform both the underlying dynamics and the effective geometry of safety constraints, rendering models trained under nominal conditions unreliable. In such settings, failures are often not caused by poor task-level performance, but by the inability to anticipate how safety-relevant constraints evolve as the physical environment changes. This gap highlights a critical limitation of many existing approaches: while they improve prediction accuracy or policy optimization under nominal assumptions, they often overlook the need to identify and track the *physical causes* that drive non-stationarity and reshape safety constraints over time.

A promising route to closing this gap is to treat non-stationarity not as unstructured noise, but as a consequence of *latent, safety-critical physical context* that evolves and can be inferred. From this perspective, the central modeling objective is not merely to predict future states, but to disentangle (*i*) task- and control-relevant state variables from (*ii*) physical context variables that govern how constraints tighten or relax (e.g., reduced friction increasing braking distance, or increased payload shrinking admissible accelerations). Importantly, such context cannot always be recovered from static appearance alone [4]. In many safety-critical scenarios, identical visual scenes can induce markedly different dynamical responses due to hidden physical differences, leading to *spurious correlations driven by appearance–physics*

Manuscript received December 30, 2025; accepted March 4, 2026. Recommended by Editor-in-Chief Hyo-Sung Ahn.

The first two authors contributed equally. S. K. On is with the Interdisciplinary Program in Artificial Intelligence, Seoul National University, Seoul 08826, South Korea. S. Jeong, J. Lee and I. Yang are with the Department of Electrical and Computer Engineering, and ASRI, Seoul National University, Seoul 08826, South Korea (e-mail: insoonyang@snu.ac.kr).
* Corresponding author.

mismatch. Reliable context inference therefore often relies on *action–reaction* signatures in interaction sequences, where applied inputs reveal latent physical properties through their dynamical consequences [5–8].

These considerations motivate a class of *physics-informed causal reasoning and modeling* pipelines that (a) encode physical context in interpretable, physically aligned representations, (b) learn context-modulated dynamics that preserve physical consistency, and (c) quantify uncertainty in actionable forms that can be propagated to downstream safety mechanisms. Throughout this review, we use the term *Physical AI* to refer to embodied physical systems in which learning-enabled perception, modeling, and control are explicitly grounded in real-world physics, rather than treated as purely statistical function approximation problems [9, 10].

Despite substantial progress across related areas—including system identification, physics-informed learning, safe control, and robust learning—the literature remains fragmented along methodological and disciplinary boundaries [11–14]. Existing works often address context inference, dynamics adaptation, uncertainty quantification, and safety enforcement in isolation, with limited guidance on how these components should be coherently integrated to handle non-stationary physical environments. As a result, it remains unclear which modeling choices are essential for anticipating safety-constraint deformation, how uncertainty should be represented and calibrated under distribution shift, and how different forms of non-stationarity ought to be evaluated and compared. In particular, many failures arise from shortcut features or non-informative visual cues that obscure the true physical causes of unsafe behavior.

Scope and positioning. This review adopts a modeling-centric perspective. Rather than surveying safe control or safe reinforcement learning algorithms themselves, we focus on the upstream questions that often determine whether such algorithms succeed or fail under non-stationarity: how safety-relevant physical context is inferred, how dynamics models adapt coherently as context evolves, and how uncertainty is represented in forms that are actionable for safety mechanisms.[1] Detailed treatments of control barrier functions, Hamilton–Jacobi reachability, and predictive safety filters are therefore outside the scope of this review and are covered extensively in dedicated surveys and tutorial articles [15–18]. Our goal is to complement these works by clarifying the causal modeling and uncertainty issues that underpin reliable safety enforcement in evolving and deceptive physical environments.

Causal modeling viewpoint. To make the role of "causal" precise in this review, we use *causal reasoning* in a structural sense: non-stationarity is modeled via (possibly latent) physical context variables that *causally* modulate (i) the state-transition mechanism and (ii) the induced geometry of safety constraints. This viewpoint highlights *interventions*—exogenous changes of physical conditions (e.g., surface, payload, wear) and endogenous action-induced probes—as the key to identifiability, and motivates representations that remain stable across environments rather than relying on appearance-driven shortcuts.

Crucially, this notion of causality goes beyond exploiting statistical correlations: a latent context variable is treated as causal only insofar as it captures a mechanism that remains consistent under interventions and therefore predicts how dynamics and safety constraints deform as physical conditions evolve. This contrasts with purely predictive or task-conditioned representations, which may match observed trajectories under nominal operating conditions yet rely on incidental cues (spurious correlations) that fail when system properties or environmental regimes change.

Accordingly, our use of "causal" is not an exhaustive survey of causal discovery, particularly the topological perspective rooted in structure learning and Structural Causal Models (SCMs) [19–21]; rather, it is a modeling lens that connects identifiability, mechanism invariance, and actionable uncertainty to safety enforcement under distribution shift.

Contributions. In this paper, we present a comprehensive review of physics-informed causal reasoning and modeling for non-stationary environments, with safety-critical control as the motivating application. Specifically, we (i) propose a unified taxonomy that separates *task-relevant state* from *safety-critical physical context* and highlights failure modes induced by appearance–physics mismatch, (ii) organize prior work along axes of non-stationarity type, context identification paradigm (passive versus action–reaction), and physics-prior injection mechanism (loss, architecture, or constraint), and (iii) synthesize uncertainty representations—including distributionally robust formulations—that can be directly consumed by downstream safety mechanisms such as constraint tightening, risk-sensitive evaluation, and conservative fallback. We further provide an evaluation checklist and reporting protocol to improve comparability across studies, and conclude with open problems and a research agenda toward deployment-ready physical AI in non-stationary environments. Fig. 1 provides a system-level overview of how context inference, context-modulated dynamics, and actionable uncertainty connect to downstream safety mechanisms under non-stationarity.

Organization. The remainder of this review is organized as follows. Section 2 characterizes non-stationarity in

[1] While many robotic systems leverage force/torque sensing, motor-current feedback, or tactile measurements to infer interaction properties, this review does not aim to survey sensor-specific pipelines or hardware-centric solutions; instead, we emphasize modality-agnostic modeling principles for inferring and propagating safety-relevant physical context to downstream safety layers.

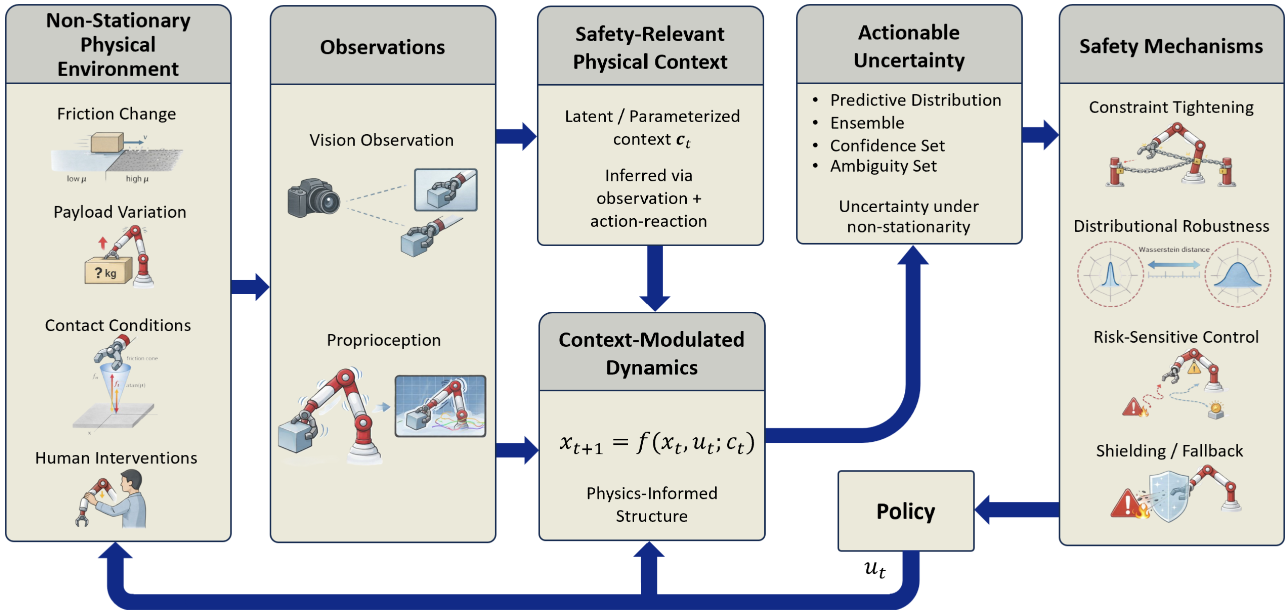

Fig. 1. Overview of the modeling-centric pipeline reviewed in this paper. Non-stationary physical environments induce time-varying safety-relevant physical context, which is inferred from observations and interaction (action–reaction) and incorporated into context-modulated, physics-informed dynamics models. The resulting uncertainty is represented in actionable forms (e.g., distributions, sets, ambiguity sets) and translated into downstream safety mechanisms such as constraint tightening, risk-sensitive evaluation, distributional robustness, and fallback strategies. In this causal framing, non-stationarity is interpreted as interventions on latent physical context that induce mechanism changes in both dynamics and safety constraints.

physical systems and formalizes how physical context deforms dynamics and safety constraints. Section 3 surveys approaches to identifying safety-relevant physical context from observations and interactions. Section 4 reviews context-modulated dynamics models and physics-informed structure. Section 5 focuses on uncertainty quantification under distribution shift, emphasizing representations that are actionable for safety mechanisms. Section 6 proposes an evaluation checklist and reporting protocol, and Section 7 concludes with open problems and directions for future research.

2. NON-STATIONARITY IN PHYSICAL SYSTEMS

Non-stationarity in physical environments refers to systematic changes in the underlying dynamics and constraints that cannot be explained by stochastic noise alone [22–24]. In safety-critical settings, such changes are particularly consequential because they directly deform the geometry of safety constraints over time. For example, reduced friction increases braking distance, additional payload tightens admissible acceleration bounds, and human intervention alters contact modes and feasible trajectories. Unlike stationary uncertainty, which can often be absorbed by robust margins, non-stationary physics introduces evolving regimes in which previously safe actions may become unsafe without explicit changes in task state. This section provides a taxonomy of non-stationarity types and highlights why they challenge safety guarantees, motivating context-aware modeling in the subsequent sections.

2.1. Types of non-stationarity in physical environments

We categorize non-stationarity in physical environments according to *what changes* and *how it impacts* safety-critical behavior. While the boundaries between categories are not always sharp, this taxonomy is useful for organizing modeling assumptions, data generation protocols, and evaluation criteria.

Parametric non-stationarity (changing physical parameters). In many robotic tasks, the dominant source of non-stationarity is a drift or shift in a low-dimensional set of physical parameters [25–29]. Examples include friction coefficients, payload mass and inertia, damping and compliance, and actuator/sensor biases. Such changes can be gradual (e.g., wear and temperature effects) or abrupt (e.g., switching surfaces or tools). Parametric shifts often preserve the functional form of the dynamics but alter its coefficients, making them well-suited for context-conditioned models in which a latent physical context variable modulates the system parameters [30, 31].

Structural non-stationarity (changing regimes and constraints). Beyond parameter drift, physical systems frequently exhibit regime changes that modify the structure of the dynamics or constraints [32–37]. Common instances are contact mode switches (sticking vs. slipping), changes in kinematic constraints (free motion vs. constrained motion), and intermittent actuation limits or saturation. Structural non-stationarity is particularly challenging because it cannot be captured by smoothly varying parameters alone; instead, it typically requires hybrid system representations (e.g., discrete modes with continuous state), explicit modeling of contact/constraint transitions, or latent variables that encode regime identity [38–40].

Interactive non-stationarity (exogenous interventions and multi-agent effects). In human-in-the-loop settings, the environment evolves through interaction [41–44]. Human intervention, crowding, occlusions, and unmodeled external forces can introduce non-stationarity that is neither purely parametric nor purely structural. The key difficulty is that the dynamics become influenced by another agent whose intentions are only partially observable. From a causal perspective, these effects correspond to interventions on the system through forces, contact events, or constraint modifications, motivating models that can attribute changes in trajectories to exogenous causes rather than to task state alone [45, 46].

Deceptive and confounded non-stationarity (appearance–physics mismatch). A safety-critical failure mode arises when visual appearance is weakly informative of physical properties [47–49]. Objects or environments that look identical may have drastically different internal mass distributions, compliance, or frictional behavior, leading to "deceptive" regimes in which purely appearance-based inference fails. In such cases, reliable context inference often requires action–reaction evidence, counterfactual or interventional data, and representations that explicitly separate nuisance visual factors from latent physical causes [5, 50].

Degradation-driven non-stationarity (aging and performance loss). Long-horizon deployment introduces persistent changes due to wear, mechanical degradation, and calibration drift [51–54]. Examples include increasing joint friction, reduced torque capability, sensor drift, and intermittent faults. Unlike one-off shifts, degradation typically accumulates over time and can interact with other sources of non-stationarity (e.g., increased friction amplifying slip under low surface friction). Modeling degradation benefits from explicit uncertainty tracking and mechanisms that distinguish temporary disturbances from persistent parameter shifts [55–57].

Implications for modeling and evaluation. These categories suggest that non-stationarity is not a monolithic phenomenon; different mechanisms call for different modeling choices and different evaluation protocols. Parametric shifts emphasize physically aligned context inference and context-modulated dynamics, whereas structural and interactive shifts highlight the need for regime-aware representations and causal attribution of exogenous interventions. Deceptive settings stress action–reaction identification and counterfactual data generation. In the next section, we review how safety-relevant physical context can be inferred under these forms of non-stationarity, and how such context can be encoded in representations that are actionable for downstream safety mechanisms.

2.2. Why non-stationarity breaks safety guarantees

Non-stationarity poses a fundamental challenge to safety-critical control because it alters the geometry of safety constraints in ways that are not captured by fixed models or static robustness margins. Safety constraints—such as collision avoidance, braking distance, force limits, or admissible accelerations—are typically derived under nominal assumptions about physical parameters. When these parameters shift, actions that were previously safe may violate constraints without any apparent change in task state. As a result, safety failures often occur not due to incorrect task execution, but due to unmodeled changes in the physical context that *invalidate* the assumed constraint feasibility [58–60].

More specifically, in non-stationary settings, the effective safety constraints become *context-dependent*: changes in physical parameters such as road friction, surface grade, or regulatory limits alter the feasible action set without necessarily manifesting as explicit state changes. Consequently, standard learning-based safety mechanisms often guarantee constraint satisfaction only in expectation over historical data distributions, rather than *conditionally* on the current physical context. This leads to delayed adaptation to regime shifts, biased safety value estimates due to distributional mismatch, and policies that implicitly average across heterogeneous safety regimes, resulting in violations even when the policy appears optimal under nominal conditions.

A common response to this challenge is to increase conservatism, for example by enlarging safety margins or relying on worst-case bounds [61–63]. While effective in preventing immediate violations, such approaches can severely degrade performance and may still fail under structured non-stationarity, where parameter shifts induce qualitative changes in dynamics (e.g., contact modes) rather than simple parametric drift. In contrast, safety under non-stationary physics requires models that can anticipate how constraints deform as context evolves, rather than treating uncertainty as unstructured noise. This observation motivates a shift from purely state-based safety reasoning toward causal modeling of safety-relevant physical context, which we address in the next section.

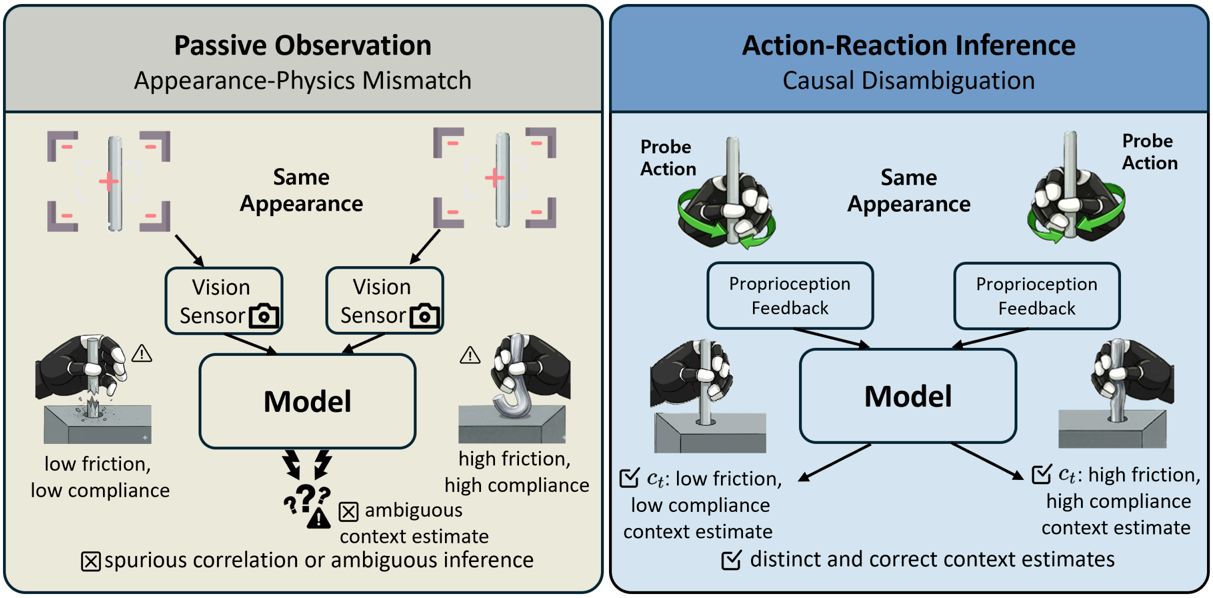

Fig. 2. Appearance–physics mismatch and action–reaction for safety-relevant context inference in a peg-in-hole scenario. Left: passive, appearance-based inference fails when visually similar environments exhibit different physical regimes, leading to spurious correlations and unsafe predictions. Right: interaction-induced action–reaction responses reveal latent physical context (e.g., friction, mass), enabling reliable inference of safety-critical conditions under non-stationary physics.

3. INFERRING SAFETY-RELEVANT PHYSICAL CONTEXT

The preceding sections highlighted that safety violations under non-stationary physics are often driven by latent physical context rather than observable task state. We now review methods for inferring such safety-relevant physical context from data. Existing approaches can be broadly categorized by whether context is inferred *passively* from observations, or *actively* through interaction-induced action–reaction signals. From a causal perspective, passive inference corresponds to an observational setting, whereas action–reaction inference exploits interventional information in which actions serve as informative probes; this distinction clarifies when safety-relevant context is identifiable and when appearance-driven shortcuts can fail under distribution shift. This section surveys both paradigms, emphasizing their underlying assumptions, identifiability properties, and implications for safety-critical modeling under non-stationarity. Related ideas have long been studied in systems and control through multiple-model filtering and experiment design [64–67], which provide complementary perspectives on context (mode) inference and informative excitation.

Fig. 2 illustrates why passive, appearance-based inference can fail under non-stationary and deceptive environments and how action–reaction cues can reveal safety-relevant physical context beyond observation alone; Table 1 summarizes the corresponding inference paradigms reviewed in the following subsections.

3.1. Passive context inference from observations

A large body of prior work seeks to infer physical properties or latent context directly from observations, without explicit interaction designed for identification. Typical examples include vision-based estimation of mass, friction, stiffness, and contact properties, as well as latent-variable models that infer hidden system parameters from state trajectories [68–76]. These approaches are attractive due to their simplicity and compatibility with passive data collection, and have been successfully applied in settings where appearance strongly correlates with physical properties.

However, passive inference faces fundamental limitations in safety-critical non-stationary environments. When multiple physical regimes induce similar observational signatures, purely observational methods may conflate nuisance factors with safety-relevant context. This issue is particularly pronounced in deceptive settings, where visual appearance is weakly informative of underlying physics, and in interactive scenarios where exogenous forces or human intervention alter dynamics without consistent visual cues [110–112]. As a result, passive context inference can lead to overconfident but miscalibrated estimates, undermining downstream safety guarantees.

Several works attempt to mitigate these issues by incorporating temporal context, uncertainty estimation, or physics-inspired priors into passive models [99, 100, 113, 114]. While such extensions improve robustness to noise,

Table 1. Taxonomy of physical-context inference under non-stationary environments (Section 3).

Paradigm	Core idea	Strengths	Failure modes / caveats	References
Passive Inference	Infer physical properties or latent context directly from observations / trajectories (no designed interaction).	Simple deployment; compatible with offline logs; effective when appearance correlates with physics.	Identifiability gaps under appearance–physics mismatch; confounding by nuisance factors; overconfident/miscalibrated estimates in deceptive or interactive regimes.	[68–76]
Active Inference	Use interaction-induced signatures (action → reaction) to disambiguate safety-relevant physical causes (e.g., friction, payload, contact regimes).	Improves identifiability; reveals hidden physics in deceptive settings; enables early detection of regime changes.	Requires safe-yet-informative probing; exploration vs safety trade-off; action constraints may limit identifiability.	[5, 67, 77–98]
Hybrid Inference & Identifiability	Combine passive observation with selective interaction; explicitly reason about identifiability under sensing/action constraints and track uncertainty over context.	Practical for deployment; reduces probing burden; supports multi-rate architectures (fast low-level cues + slow perception).	Some contexts only locally identifiable; uncertainty must be represented in actionable form; partial observability complicates attribution.	[11, 46, 65, 67, 81, 85, 93, 99–110]

they do not fundamentally resolve the identifiability gap inherent to observation-only inference. These limitations motivate approaches that leverage interaction and intervention to reveal causal physical context, which we review next.

3.2. Action–reaction and active context inference

To overcome the identifiability limitations of passive inference, a growing line of work leverages interaction to actively reveal latent physical context. From a causal perspective, physical context often acts as a hidden confounder that influences system response but cannot be reliably inferred from observation alone. By deliberately exciting the system through actions and observing the resulting reactions, active inference exploits discrepancies in action–reaction signatures to disambiguate safety-relevant physical causes [5, 81–83, 85, 89]. In systems and control, this viewpoint is closely related to experiment/input design for identification, where informative excitation is engineered to maximize identifiability under constraints [67, 77, 78, 82, 115].

In robotic systems, action–reaction signals manifest through differences in acceleration, slip, contact forces, compliance, or response delays under identical commanded actions. Even when visual appearance remains unchanged, such interaction-induced cues can expose variations in friction, mass distribution, stiffness, or contact regimes that directly affect safety constraints. This principle underlies a range of interactive perception and active system identification methods, including manipulation-based probing of liquids and other challenging physical phenomena [5, 86, 87, 90].

Active context inference is particularly well-suited to safety-critical non-stationary environments. First, it enables early detection of regime changes before constraint violations occur, allowing safety mechanisms to adapt proactively rather than reactively [92, 96]. Second, it supports causal attribution by distinguishing whether deviations in behavior arise from changes in task state, exogenous interventions, or underlying physical context [84]. Third, when combined with physics-informed priors, action–reaction inference can be performed with minimal excitation, mitigating the risk of unsafe probing [93, 94].

Nevertheless, active inference introduces its own challenges. Designing informative yet safe probing actions requires balancing exploration with constraint satisfaction, especially in human-in-the-loop settings [79, 80, 89, 95, 97]. Moreover, not all physical context variables are equally identifiable through interaction, and the degree of excitation required for reliable inference depends on both system dynamics and sensing modalities [88, 98]. These considerations motivate hybrid approaches that combine passive observation with selective, physics-aware interaction, as well as representations that explicitly encode uncertainty over inferred context.

3.3. Hybrid inference and identifiability considerations

In practice, reliable inference of safety-relevant physical context often requires a hybrid strategy that combines passive observation with selective action–reaction cues. Purely passive inference may be insufficient under deceptive or confounded regimes, while fully active probing can be impractical or unsafe in real-world deployments. Hybrid inference approaches aim to exploit readily available observational data while invoking interaction only when ambiguity remains high or safety margins begin to erode [81, 85, 102].

From a modeling standpoint, hybrid inference raises fundamental questions of identifiability. A physical context variable is identifiable if distinct values of the context induce distinguishable distributions over observations or interaction outcomes under admissible actions. In non-stationary environments, identifiability depends not only on sensing modalities and model expressiveness, but also on the set of actions that can be safely executed [93, 103, 104]. This connects directly to classical identifiability and experiment design considerations in system identification [11, 67]. As a result, some safety-critical context variables may be locally identifiable only within certain operating regimes, while remaining indistinguishable elsewhere.

These considerations highlight the importance of representations that explicitly encode uncertainty over inferred context. Rather than committing to point estimates, hybrid models benefit from maintaining distributions or sets over latent physical context, which can guide both selective interaction and downstream safety mechanisms [65, 99, 100, 110]. When uncertainty is high, conservative behavior or additional information-gathering actions may be warranted; when uncertainty collapses, models can exploit more aggressive control strategies without compromising safety.

Hybrid inference also interacts closely with system architecture and timing. In many robotic systems, fast low-level signals (e.g., joint encoders, motor currents) provide early indicators of physical regime changes, while slower perceptual streams (e.g., vision) refine context estimates over longer horizons. Effective hybrid approaches therefore integrate multi-rate inference pipelines and leverage physics-informed priors to reconcile information across time scales [85, 91]. Taken together, hybrid inference and identifiability considerations suggest that safety under non-stationary physics is fundamentally a problem of causal model design, rather than algorithm selection alone [46, 101, 105].

Box 1: Practical deep learning tools for causal reasoning under non-stationarity

While this review adopts a modeling-centric notion of causal reasoning rather than a full survey of causal discovery, it is useful to summarize concrete deep learning patterns that operationalize causal ideas in non-stationary physical systems.

Invariant representation learning. A central principle is that predictive mechanisms should remain stable across environments, while shortcut correlations vary. Recent work clarifies practical objectives and trade-offs for enforcing invariance and improving out-of-distribution generalization [101, 116–119].

Counterfactual and paired-data training. Causal reasoning is strengthened when training data contains controlled contrasts—e.g., paired samples where appearance is fixed but physical properties differ (and vice versa). In practice, this is realized via counterfactual data generation and structured augmentation schemes that break shortcut features and expose causal factors [27, 120–123].

Latent causal variable modeling. Many approaches introduce latent variables intended to capture safety-relevant physical context and constrain how these variables modulate dynamics. This aligns with recent causal representation learning results establishing identifiability of latent causal variables under interventional, multi-domain, or multi-view settings [81, 124–127].

Intervention through action–reaction. Actions can be interpreted as interventions that reveal causal structure. Recent work demonstrates this principle through active system identification and interactive physical reasoning, where exploration policies are designed to collect informative action–reaction data and disambiguate latent physical regimes [5, 93, 128, 129].

Uncertainty-aware causal reasoning. Under non-stationarity, causal assumptions may be violated or only locally valid; recent work therefore combines causal estimation with calibrated uncertainty quantification and robustness under distribution shift [99, 130–132].

In line with this view, recent work increasingly emphasizes *contextual* and *physics-informed* modeling pipelines—including embodied world models and context-aware learning under observed contexts—to improve robustness under non-stationarity and enable deployment-oriented inference architectures [106–109].

To conclude this section, Box 1 summarizes practical deep learning patterns that operationalize causal reasoning in non-stationary physical systems, bridging the conceptual discussion above with the modeling frameworks reviewed in the next section.

4. CONTEXT-MODULATED DYNAMICS MODELING

Having discussed how safety-relevant physical context can be inferred, we now review how such context is incorporated into dynamics models to handle non-stationary environments. The central question addressed in this section is how to move beyond fixed dynamics representations and construct models whose predictions adapt coherently as physical context evolves, while remaining compatible with known physical principles. We primarily adopt a discrete-time notation for clarity and alignment with sampled robotic data streams; continuous-time counterparts are discussed when they provide the most natural expression of physical structure.

Table 2 summarizes the main modeling paradigms reviewed in this section, organizing them by how context enters the dynamics and how physics priors are injected (via losses, architectures, or explicit constraints), together with their typical strengths and failure modes. We begin by contrasting fixed and context-modulated dynamics and then discuss parameter-conditioned and latent-context formulations, physics-informed structure, and common failure modes under regime shifts.

4.1. Fixed versus context-modulated dynamics

Many learning-based dynamics models assume a fixed mapping from state and control input to next state, implicitly treating non-stationarity as noise or unmodeled disturbance [68, 72, 133–136]. In its simplest form, this corresponds to learning a single transition function

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

where $x_t \in \mathcal{X}$ and $u_t \in \mathcal{U}$ denote the state and control input, and ε_t captures stochasticity and modeling error. Throughout this review, we use $f(\cdot)$ to denote a dynamics model with *fixed, shared parameters* learned offline, and reserve the time-varying variable c_t to represent non-stationary physical context.

While fixed dynamics models can perform well in stationary or narrowly varying regimes, they tend to fail under structured non-stationarity, where identical (x_t, u_t) pairs can yield qualitatively different outcomes due to changes in physical parameters or operating conditions [81, 85, 137, 160, 161]. A common mitigation strategy is to increase robustness through conservative modeling, for example by enlarging safety margins or training over randomized environments [27, 28, 63]. Although effective against unstructured uncertainty, these approaches do not explicitly represent the causes of non-stationarity and therefore struggle to extrapolate to unseen regimes. In particular, when safety constraints depend sensitively on physical parameters—such as friction, payload, or compliance—a fixed dynamics model cannot anticipate how constraint geometry deforms as context shifts.

Context-modulated dynamics models address this limitation by conditioning the dynamics on an explicit or implicit representation of physical context [124, 162, 163]. Instead of learning a single $f(x_t, u_t)$, these models represent a family of dynamics

$$x_{t+1} = f(x_t, u_t, c_t) + \varepsilon_t, \tag{2}$$

where $c_t \in \mathcal{C}$ captures the aspects of the environment responsible for non-stationary behavior. This formulation provides a natural interface between context inference (Section 3) and downstream safety reasoning, as changes in c_t directly induce predictable changes in system response [5, 81, 164–166].

4.2. Parameter-conditioned and latent-context dynamics

Context-modulated dynamics models can be broadly divided into *parameter-conditioned* and *latent-context* approaches. In parameter-conditioned models, context corresponds to interpretable physical parameters

$$\theta \in \Theta \subset \mathbb{R}^d, \qquad \theta = (\mu_f, m, \mathrm{load}, \dots), \tag{3}$$

such as friction coefficients μ_f, payload mass m, or inertial/compliance parameters, which are either estimated online or provided as inputs. Dynamics are then conditioned on θ,

$$x_{t+1} = f(x_t, u_t; \theta) + \varepsilon_t, \tag{4}$$

which yields physical interpretability and can directly leverage known structure from mechanics, but requires that the relevant parameters be identifiable and measurable (or estimable) at runtime [81, 137, 138].

Latent-context dynamics models relax this requirement by representing context as a low-dimensional latent variable inferred from interaction history [30, 85, 124]. A common abstraction is to infer c_t from a window of recent observations/inputs,

$$c_t = g(h_t), \qquad h_t := (x_{t-K:t}, u_{t-K:t-1}), \tag{5}$$

and condition either the dynamics model or the control policy on c_t. This paradigm underlies online adaptation and system identification approaches, where the latent context is updated in real time and used to condition the dynamics model or the policy [81, 85, 93, 124, 139, 140].

Table 2. Taxonomy of context-modulated dynamics modeling under non-stationary environments (Section 4).

Model class	Core idea	Physics prior injection	Strengths / failure modes	References
Fixed Dynamics	Single transition model trained to explain all regimes.	None (or implicit via data); sometimes randomized training.	+ Simple; stable in stationary/narrow regimes. – Breaks under structured parameter/regime shifts.	[68, 72, 133–136]
Parameter-Conditioned Dynamics	Condition dynamics on interpretable physical parameters (e.g., friction, mass).	Parameter bounds; physics-residual losses; scheduled/estimated variables.	+ Interpretable; supports error attribution. – Requires identifiable parameters; risk of misspecification.	[81, 137, 138]
Latent-Context Dynamics	Infer latent context from recent interaction history and condition dynamics or policy.	Often none by default; optional stability regularizers; downstream constraints.	+ Flexible; fast adaptation to heterogeneous shifts. – Latent misalignment; brittle under deceptive regimes.	[30, 81, 85, 93, 111, 124, 139, 140]
Physics-Informed Losses	Regularize dynamics by penalizing violations of governing relations.	Soft physics constraints via residual losses.	+ Improves plausibility and generalization. – Sensitive to residual design and weighting.	[12, 141–145]
Structure-Preserving Arch.	Impose physical structure by construction (e.g., energy-based).	Architecture-level physics encoding.	+ Preserves invariants; improves extrapolation. – Limited expressiveness for contact/deformation.	[146–151]
Context-Modulated World Models	Combine context modulation with physics priors for non-stationary regimes.	Loss/architecture priors; often multi-rate deployment.	+ Robust across regimes; coherent-driven modulation. – Context misidentification; latency/discretization transients.	[100, 106, 108, 152–159]

From a modeling perspective, latent-context approaches offer flexibility and scalability, particularly in complex environments where explicit parameterization is difficult. They also naturally accommodate partial observability and heterogeneous sources of non-stationarity. However, their expressiveness comes at the cost of reduced interpretability, and improper alignment of latent dimensions can lead to brittle adaptation or spurious correlations [111].

Recent work has explored hybrid formulations that combine latent context embeddings with physics-informed structure, for example by using inferred context to modulate physically meaningful components of the dynamics rather than arbitrarily altering the transition function [154, 156]. One representative pattern is to use the context variable to parameterize families of structured models,

$$x_{t+1} = f(x_t, u_t; \alpha(c_t)) + \varepsilon_t, \tag{6}$$

where $\alpha(c_t)$ denotes context-dependent parameters (e.g., frictional or inertial terms) or feature-wise modulation coefficients [146, 167, 168]. Such approaches aim to retain the adaptability of latent representations while constraining them to evolve in ways consistent with known mechanics [100, 108, 152, 153].

To illustrate this mechanism in practice, consider a standard rigid-body dynamics model where c_t is mapped to the inertial matrix $M(c_t)$ or the Coriolis term $C(x_t, c_t)$ rather than influencing the acceleration directly as a black-box input. By restricting context modulation to these specific terms, the model enforces key physical invariants—such as energy conservation or passivity—regardless of the value of the inferred context. This design pattern ensures that even if the latent context estimation is noisy or uncertain, the resulting trajectory predictions remain physically plausible, thereby enhancing the reliability of downstream safety mechanisms.

4.3. Physics-informed structure and consistency

While context-modulated dynamics provide a mechanism to adapt predictions under non-stationarity, they raise a critical modeling question: *what constraints should be* *imposed on how context influences the dynamics?* Without additional structure, highly expressive models risk fitting spurious correlations, producing physically implausible trajectories, or extrapolating poorly to unseen regimes. This issue is particularly acute in safety-critical settings, where unphysical predictions can directly translate into incorrect constraint evaluation and unsafe behavior.

Many physics-informed learning approaches are naturally formulated in continuous time, reflecting the fact that governing physical laws are expressed as differential equations [141, 142, 147, 149]. In this view, system evolution is modeled as

$$\dot{x}(t) = f_c\left(x(t), u(t); c(t)\right), \tag{7}$$

where $f_c(\cdot)$ encodes physically meaningful structure such as conservation laws, contact dynamics, or force–acceleration relationships, and $c(t)$ represents time-varying physical context. For data-driven learning and downstream decision layers, the continuous-time model is typically discretized with sampling interval Δt, yielding a discrete-time transition of the form

$$\begin{aligned} x_{t+1} &\approx x_t + \int_{t\Delta t}^{(t+1)\Delta t} f_c\left(x(\tau), u(\tau); c(\tau)\right) d\tau \\ &=: f(x_t, u_t; c_t) + \varepsilon_t, \end{aligned} \tag{8}$$

which recovers the context-modulated discrete-time formulation used throughout this section.

From a practical perspective, physics priors can be injected into learning-based dynamics through three main mechanisms: *(i) physics-informed losses* that penalize violations of governing relations [142, 144], *(ii) physics-structured architectures* that enforce qualitative behavior by construction (e.g., energy-based parameterizations) [147–149], and *(iii) explicit constraints* that restrict trajectories or parameters to physically admissible sets [142,169]. These mechanisms are complementary and often combined in modern pipelines.

A simple example of (i) is Newtonian consistency between force and acceleration [145]. Let $\hat{a}_t$ and $\hat{F}_t$ denote the acceleration and net force implied by the predicted trajectory and control input under context c_t (or parameter estimate θ). One may use a physics-informed loss function:

$$\mathcal{L}_{\text{Newton}} = \sum_t \left\| \hat{m}(c_t)\, \hat{a}_t - \hat{F}_t \right\|_2^2, \tag{9}$$

where $\hat{m}(c_t)$ denotes a context-dependent mass estimate. More generally, physics-informed losses can penalize residuals of governing equations and constraints, including kinematic relations, contact-mode consistency, and energy balance relations [12, 142, 143].

For (ii), many physics-informed models impose structure at the level of the governing equations themselves [150]. In Hamiltonian neural networks, for example, one learns an energy function $\mathcal{H}(q, p; c)$ and derives dynamics via Hamilton's equations

$$\dot{q} = \frac{\partial \mathcal{H}}{\partial p}, \qquad \dot{p} = -\frac{\partial \mathcal{H}}{\partial q}, \tag{10}$$

while Lagrangian formulations parameterize a learned $\mathcal{L}(q, \dot{q}; c)$ and obtain dynamics through Euler–Lagrange equations [146, 147, 149]. Such parameterizations encode conservation laws and qualitative properties (e.g., symmetries, invariants) by construction, and provide a principled mechanism for incorporating physical context through modulation of physically meaningful terms rather than arbitrarily altering a black-box transition function [146, 147, 151].

Finally, regarding (iii), explicit constraints are essential for modeling systems where energy-based priors alone are insufficient. While Hamiltonian or Lagrangian networks effectively regularize conservative systems, they often fail to capture *dissipative* and *nonsmooth* phenomena—such as dry friction, viscous damping, or discontinuous collisions. To address this, dynamics modeling can be cast as a constrained optimization problem or a differential variational inequality. In this view, physics priors are imposed as geometric or algebraic constraints: contact non-penetration is modeled via complementarity conditions or barrier functions, while friction is enforced through the maximum dissipation principle or friction cone limits [170, 171]. Recent approaches have integrated these complementarity constraints directly into network architectures, utilizing differentiable optimization layers to accurately capture hybrid contact dynamics [172, 173]. By restricting the learned model to satisfy these inequalities—whether through hard projections onto admissible sets or downstream shielding—we ensure that predicted trajectories respect thermodynamic irreversibility and contact geometry, preventing the nonphysical oscillations or drift common in purely conservative formulations [174, 175].

In non-stationary environments, physics-informed structure plays a dual role. First, it improves generalization by biasing the model toward physically meaningful extrapolations when encountering novel contexts [157, 158]. Second, it provides a coherent interface for context modulation: latent or estimated physical context can be used to modulate interpretable components such as inertial, frictional, damping, or contact terms, enabling the model to adapt while remaining physically plausible [155, 159]. Recent work on physics-informed world models and structured latent dynamics exemplifies this trend, demonstrating improved robustness under distribution shift and regime changes [108].

Importantly, physics-consistency should be viewed as an enabler of safety rather than a purely aesthetic modeling choice. By constraining how dynamics can evolve under context changes, such models reduce the risk of unphysical predictions that would otherwise undermine downstream safety mechanisms. Moreover, structured models facilitate error attribution by clarifying whether prediction errors arise primarily from context misestimation, unmodeled physical effects, or data scarcity. This clarity becomes essential when uncertainty must be propagated to safety-critical decision layers, motivating the uncertainty-aware modeling tools reviewed next.

4.4. Adaptation, generalization, and failure modes

Despite the benefits of context modulation and physics-informed structure, accurate dynamics modeling alone does not guarantee safe behavior under non-stationarity. In practice, adaptation mechanisms introduce new failure modes that arise from misidentification, limited excitation, and mismatches between modeling and execution timescales. Understanding these failure modes is essential for assessing when adaptive models can be trusted and when additional safeguards are required [176].

A primary source of failure stems from *context misidentification* [111, 177]. When inferred context variables are poorly aligned with the true physical causes of non-stationarity, models may adapt in the wrong direction, amplifying prediction errors rather than correcting them. This issue is exacerbated under partial observability, where multiple physical regimes induce indistinguishable responses within the admissible operating envelope. Even physics-informed models can exhibit such failures if the assumed structure does not capture the dominant modes of variation or if context is inferred from insufficiently informative interaction data [11, 64, 93, 177].

A second challenge concerns *generalization beyond the training distribution* [178, 179]. While physics-informed priors improve extrapolation relative to unconstrained models, they do not guarantee correctness under qualitatively novel regimes, such as unmodeled contact modes, deformation, wear, or human intervention. Domain randomization and context conditioning expand the range of covered scenarios, but rare or extreme combinations of physical parameters can still violate implicit assumptions embedded in the model [27, 28]. In safety-critical settings, such failures are particularly concerning because small modeling errors can translate into large violations of safety constraints.

Temporal and architectural factors further complicate adaptation [91, 180, 181]. In many robotic systems, low-level control loops operate at high frequency, while context inference and model updates occur on slower timescales. This multi-rate structure can lead to transient inconsistencies, where outdated context estimates are used for safety evaluation or where rapid control actions are taken based on stale dynamics predictions. Without careful coordination, such mismatches can induce unsafe behavior even when long-term adaptation converges [80, 85, 106, 182].

Taken together, these failure modes highlight a fundamental limitation of adaptation-centric approaches: committing to a single context estimate or a single adapted dynamics model is often unsafe during periods of ambiguity or transition. What is required instead is a principled representation of *model uncertainty* that captures both epistemic uncertainty due to limited data and structural uncertainty arising from model mismatch. Such representations enable systems to reason conservatively when confidence is low and to relax constraints only as uncertainty collapses.

This observation motivates the next section, which surveys uncertainty quantification methods tailored to non-stationary and physics-informed settings. We focus on uncertainty representations that are not only statistically meaningful, but also *actionable* for safety-critical decision making, allowing adaptive systems to balance performance and safety under evolving physical conditions.

5. UNCERTAINTY QUANTIFICATION FOR NON-STATIONARY, SAFETY-CRITICAL MODELING

The failure modes discussed in Section 4.4 suggest that committing to a single context estimate or a single adapted dynamics model is often unsafe under non-stationarity. Safety-critical systems therefore require principled uncertainty representations that capture not only stochastic disturbances, but also epistemic uncertainty due to limited data and structural uncertainty induced by regime shifts, unmodeled contacts, and human interventions. In this section, we review uncertainty quantification methods for physics-informed and context-modulated dynamics models, with an emphasis on uncertainty *representations* that are actionable for downstream safety mechanisms. We first introduce a taxonomy of uncertainty sources under non-stationarity, then survey practical estimation methods, discuss calibration and shift detection, and finally outline how uncertainty outputs can be translated into conservative safety behavior without incurring unnecessary performance loss.

Throughout this section, we use *non-stationarity* to describe time-varying physical dynamics and constraints, *distribution shift* to refer to changes in the data/transition distributions that can invalidate calibrated uncertainty, and *regime change* to emphasize discrete or structural shifts (e.g., contact-mode switches) that induce model mismatch; *shift detection* refers to mechanisms that detect such changes in real time.

5.1. Sources and taxonomy of uncertainty under non-stationarity

Uncertainty plays a fundamentally different role in safety-critical systems operating under non-stationary physics than in stationary or purely stochastic settings. In particular, uncertainty is not merely a measure of prediction variance, but a key indicator of when learned models and inferred context should no longer be trusted for safety evaluation. To clarify this distinction, we first categorize the sources of uncertainty that arise in non-stationary environments and discuss their implications for safety-critical modeling.

A classical distinction separates uncertainty into *aleatoric* and *epistemic* components. Aleatoric uncertainty captures irreducible randomness arising from sensor noise, process disturbances, or inherently stochastic phenomena [183]. In robotic systems, this includes measurement noise, unmodeled micro-perturbations, and contact variability that persists even with perfect knowledge of system dynamics. Aleatoric uncertainty can often be modeled as additive noise and is typically handled through probabilistic filtering or stochastic control formulations [184–186].

Epistemic uncertainty, by contrast, reflects a lack of knowledge about the system model itself [187]. This includes uncertainty due to limited data coverage, insufficient excitation, or ambiguity in inferred physical context. Under non-stationarity, epistemic uncertainty becomes particularly salient because learned dynamics models and context estimates may be valid only within a restricted regime of operation. Under *distribution shift*—for example, due to changes in friction, payload, wear, or human intervention—epistemic uncertainty can increase sharply, signaling that model predictions may no longer be reliable for safety assessment [99, 178, 187–189].

Beyond this classical dichotomy, non-stationary environments introduce an additional and often underappreciated source of uncertainty: *structural or regime-shift uncertainty* [190–192]. This arises when the true system dynamics no longer belong to the assumed model class, even after adaptation. Examples include unmodeled contact modes, changes in actuation limits, or qualitative alterations of interaction geometry. Structural uncertainty is not well captured by variance estimates alone and often manifests as systematic prediction errors or abrupt degradation in model performance. From a safety perspective, this form of uncertainty is especially dangerous, as it can invalidate previously learned safety margins without obvious warning.

Importantly, these sources of uncertainty interact with context inference. When physical context is inferred from data, uncertainty arises not only in state prediction but also in the context estimate itself. Ambiguity in inferred context can propagate nonlinearly through context-modulated dynamics, leading to large uncertainty in downstream safety-relevant quantities such as stopping distance, collision probability, or constraint violation risk. As discussed in Section 4.4, committing to a single context estimate during such periods of ambiguity can be unsafe.

From the standpoint of safety-critical modeling, the key question is therefore not only *how much* uncertainty is present, but *what kind* of uncertainty is being represented and *how* it should be used. Aleatoric uncertainty may be accommodated through probabilistic constraints, while epistemic and structural uncertainty often call for conservative behavior, constraint tightening, or explicit reasoning over sets of plausible models [63]. This motivates uncertainty representations that go beyond scalar confidence measures and instead provide structured outputs—such as predictive distributions, sample ensembles, or confidence sets—that can be directly consumed by downstream safety mechanisms.

In the following subsections, we review practical uncertainty quantification methods through this lens. Rather than categorizing methods solely by their statistical formulation, we emphasize the *form of uncertainty output* they produce and its suitability for safety-critical decision making under non-stationarity.

5.2. Uncertainty representations and practical estimation methods

Having identified the major sources of uncertainty under non-stationary physics, we now review practical methods for uncertainty quantification (UQ). Rather than organizing methods purely by their probabilistic formulation, we categorize them by the *form of uncertainty output* they produce. This perspective is particularly relevant for safety-critical systems, where uncertainty must ultimately be consumed by downstream safety mechanisms such as constraint tightening, risk assessment, or conservative control.

Predictive distributions. A common approach to UQ is to represent uncertainty through a predictive distribution over future states or trajectories,

$$p(x_{t+1} \mid x_t, u_t), \tag{11}$$

often parameterized by a mean and covariance $(\mu_{t+1}, \Sigma_{t+1})$. Such representations arise naturally in Bayesian neural networks, variational state-space models, and Gaussian process-based dynamics models [68, 72, 136, 193]. Predictive distributions are attractive due to their compactness and compatibility with stochastic control and risk-constrained formulations [186, 194, 195]. However, in practice they rely on strong assumptions about distributional form and can underestimate uncertainty under distribution shift or model mismatch.

Ensemble-based uncertainty. Ensemble methods approximate epistemic uncertainty by maintaining a collection of models $\{f^{(k)}\}_{k=1}^{K}$ trained with different initializations, data subsets, or stochastic regularization [99, 196]. Uncertainty is then represented implicitly through the dispersion of predictions,

$$\{x_{t+1}^{(k)} = f^{(k)}(x_t, u_t)\}_{k=1}^{K}. \tag{12}$$

Ensembles are widely used in model-based reinforcement learning and robotics due to their simplicity and empirical robustness [99, 100, 180]. They provide a nonparametric approximation of epistemic uncertainty and are less sensitive to model misspecification than single-model predictors. Nonetheless, ensembles can be computationally expensive and may still fail to capture structural uncertainty when all models share the same inductive bias.

Confidence sets and set-valued prediction. An alternative representation is to output a set $\mathcal{X}_{t+1}$ that is guaranteed (with high probability) to contain the true next state,

$$x_{t+1} \in \mathcal{X}_{t+1}, \tag{13}$$

rather than a full distribution [197–199]. Such set-valued predictions arise in conformal prediction, interval methods, and robust identification frameworks [200–205]. Confidence sets are particularly appealing for safety-critical applications because they can be directly mapped to constraint tightening or worst-case analysis without relying on distributional assumptions. Their main limitation lies in potential conservatism, especially when uncertainty is high or poorly calibrated.

Residual- and error-based uncertainty indicators. In many deployed systems, uncertainty is inferred indirectly from prediction errors, residual statistics, or inconsistency measures between model outputs and observed behavior [12, 206–208]. Examples include monitoring normalized prediction residuals, innovation sequences in filters, or violation rates of physics-consistency constraints. While such indicators do not constitute uncertainty estimates in a strict probabilistic sense, they often serve as effective triggers for conservative behavior, fallback strategies, or model adaptation. Their effectiveness depends critically on appropriate thresholds and calibration.

Uncertainty over context and model parameters. In context-modulated dynamics, uncertainty arises not only in state prediction but also in inferred context or physical parameters [124, 209–211]. Representations such as

$$p(c_t \mid h_t) \quad \text{or} \quad c_t \in \mathcal{C}_t \tag{14}$$

capture ambiguity over physical regimes and can propagate nonlinearly into safety-relevant quantities. Explicitly representing uncertainty over context is essential in non-stationary environments, as errors in context inference can dominate downstream risk estimation. Methods that collapse context uncertainty into a point estimate risk brittle behavior during regime transitions.

Discussion. Each uncertainty representation offers distinct trade-offs between expressiveness, computational cost, and suitability for safety-critical use. Predictive distributions are compact but assumption-laden; ensembles are flexible but resource-intensive; set-valued predictions provide strong safety guarantees at the cost of conservatism; and residual-based indicators offer pragmatic but heuristic signals. Importantly, the choice of UQ method should be guided not only by estimation accuracy, but by how naturally its outputs can be translated into downstream safety mechanisms. This translation, and the associated challenges of calibration and reliability under distribution shift, are the focus of the next subsection.

5.3. Calibration and shift detection under non-stationarity

Uncertainty quantification is only meaningful for safety-critical decision making if it is *reliable*. In particular, uncertainty estimates must be calibrated such that stated confidence levels accurately reflect empirical error frequencies. Under non-stationary physics, however, maintaining calibration is challenging: models trained or adapted under one regime may become systematically miscalibrated as physical context shifts.

Conformal prediction provides distribution-free coverage guarantees under minimal assumptions, and thus offers an attractive calibration tool when parametric uncertainty models are unreliable [201, 202, 212, 213]. Classical conformal methods rely on exchangeability, which is violated under non-stationarity, motivating adaptive and online variants that update calibration sets over time and aim to maintain long-run coverage under distribution shift [214–217]. These methods are particularly relevant when uncertainty estimates are fed directly into safety margins or constraint tightening rules.

Distribution shift is closely tied to calibration: safety-critical systems must detect when they are operating outside the distributional regime in which their uncertainty estimates remain trustworthy. Practical shift-detection indicators include abrupt changes in input statistics, persistent prediction residuals, physics-consistency violations, or disagreement across ensemble members [178, 218, 219]. From a systems perspective, these indicators can be interpreted as evidence that the current dynamics no longer match the assumed model class—often due to regime changes or model mismatch—and can be used to trigger conservative behavior, re-identification, or fallback control [11, 205].

5.4. Actionable uncertainty for safety mechanisms

The ultimate value of uncertainty quantification in safety-critical systems lies not in uncertainty estimation itself, but in how uncertainty is translated into concrete safety behavior. For example, in peg-in-hole insertion(Fig. 2), non-stationary physical conditions such as changes in friction, hole clearance, surface wear, slight deformation, or misalignment can make contact transitions difficult to predict using a single nominal model. Under such uncertainty, a controller that commits too aggressively to one predicted contact outcome may cause excessive insertion force, jamming, or unstable contact behavior. This motivates uncertainty-aware control, where uncertainty estimates are used to induce safer behavior, such as reducing insertion speed, increasing compliance, tightening force-related constraints, or switching to a more conservative search or recovery strategy when the predicted contact dynamics are unreliable.

In this subsection, we discuss how different forms of uncertainty representations can be systematically mapped to safety mechanisms, enabling conservative decision making under ambiguity without unnecessary loss of performance. Table 3 summarizes common uncertainty representations and illustrates how they are operationalized in downstream safety mechanisms, ranging from constraint tightening and risk-sensitive formulations to distributionally robust reasoning and conservative fallback strategies.

From uncertainty to constraint tightening. One of the most direct uses of uncertainty is constraint tightening [220, 221]. When uncertainty is represented through confidence sets, safety constraints can be conservatively adjusted to account for worst-case deviations. For example, if a safety constraint is expressed as $h(x_{t+1}) \geq 0$ and uncertainty is captured by a set $\mathcal{X}_{t+1}$, a sufficient condition for safety is

$$\min_{x \in \mathcal{X}_{t+1}} h(x) \geq 0. \tag{15}$$

This approach is particularly attractive in systems with hard safety constraints, where violation is unacceptable [63, 222, 223].

Risk-sensitive and chance-constrained formulations. When uncertainty is represented probabilistically, safety requirements can be expressed via risk measures [224, 225]. A common example is a chance constraint

$$\mathbb{P}\left(h(x_{t+1}) \geq 0 \mid x_t, u_t\right) \geq 1 - \delta, \tag{16}$$

where δ specifies an acceptable violation probability [186]. Alternative risk measures, such as Conditional Value-at-Risk (CVaR), provide a continuous trade-off between conservatism and performance by penalizing tail events more heavily than mean behavior [226]. These formulations are most effective when predictive uncertainty is well calibrated, as emphasized in Section 5.3 [227–231].

Distributionally robust safety under ambiguity sets. Under non-stationary physics, a particularly compelling viewpoint is *distributional robustness* [232], where uncertainty is represented as an ambiguity set over *transition* or *disturbance* distributions rather than a single predictive model [233–240]. Let $\widehat{P}$ denote a nominal (data-driven) conditional distribution of the next state given the current state and input, and consider a Wasserstein ambiguity set

$$\mathcal{D}(\widehat{P}, \rho) := \{P : W(P, \widehat{P}) \leq \rho\}, \tag{17}$$

where ρ controls the admissible amount of distribution mismatch. A distributionally robust safety requirement can then be expressed by enforcing safety under all transitions in $\mathcal{D}$, [241] e.g., a worst-case bound on a safety risk functional $\mathcal{R}$ defined on the one-step transition [242],

$$\sup_{P \in \mathcal{D}(\widehat{P}, \rho)} \mathbb{E}_P[\mathcal{R}(x_t, u_t, x_{t+1})] \leq \delta. \tag{18}$$

This formulation is well aligned with non-stationarity: rather than committing to a single probabilistic model, it explicitly accounts for distribution shift, and the radius ρ provides a transparent knob to trade off conservatism and performance [189, 243–246].

Crucially, the ambiguity radius ρ can be made *adaptive* using uncertainty signals from modeling and context inference. For example, increases in epistemic or shift uncertainty can trigger larger ρ (more conservative robustness), while improved calibration and reduced uncertainty allow ρ to shrink, yielding a principled "uncertainty-to-robustness" translation mechanism.

Scenario-based and sampling-based robustness. A complementary line of work enforces safety and robustness via randomized constraints built from sampled uncertainties (scenario approaches) [247, 248]. Such methods provide finite-sample guarantees for constraint satisfaction under convexity assumptions and have been widely used in systems and control to handle uncertainty without explicit distributional models [249–251]. From the present viewpoint, scenario approaches can be interpreted as another form of uncertainty-to-safety translation, where data directly induces conservative feasible sets.

Uncertainty-triggered mode switching and fallback. In deployed systems, uncertainty often serves as a trigger rather than a direct optimization variable. Sudden increases in prediction dispersion, physics residuals, or disagreement across ensemble members can signal loss of model validity, prompting a switch to conservative fallback controllers, reduced operating envelopes, or information-gathering actions [60, 252, 253]. This mechanism is especially effective in multi-rate architectures, where fast safety layers can react immediately to uncertainty spikes while slower adaptation processes update the underlying model.

Discussion. Across these mechanisms, a common theme emerges: uncertainty must be represented in a form that downstream safety layers can consume *directly*, without requiring prohibitive online inference. Whether through tightened constraints, risk bounds, ambiguity sets, or mode-switching triggers, actionable uncertainty provides a principled interface to balance safety and performance under non-stationary physics. Importantly, no single uncertainty representation is universally optimal. Effective safety-critical systems often combine multiple representations, leveraging coarse but reliable signals (e.g., residual thresholds) to guard against rare, catastrophic failures, while exploiting more refined probabilistic estimates when confidence is high.

A critical practical consideration for these methods is their computational cost, which often conflicts with the

Table 3. Actionable uncertainty: translating uncertainty representations into safety mechanisms under non-stationary physics (Section 5).

Uncertainty output	How it is obtained	Safety translation	References
Predictive distribution (μ, Σ)	Probabilistic dynamics; Bayesian state-space models; Gaussian approximations	Chance constraints; risk-sensitive bounds (e.g., CVaR)	[68, 72, 136, 186, 193–195]
Ensemble dispersion	Deep ensembles; bootstrap; stochastic regularization	Empirical worst-case risk; uncertainty-triggered conservatism	[99, 100, 180, 196]
Confidence set / interval	Conformal prediction; set-valued identification; residual bounds	Constraint tightening; robust feasibility checks	[197–205]
Ambiguity set of distributions	Wasserstein or moment-based ambiguity sets	Distributionally robust safety constraints	[189, 232–243, 245, 246]
Scenario-based uncertainty	Randomized samples of uncertainty; scenario programs	Finite-sample safety guarantees via sampled constraints	[247–251]
Residual / inconsistency indicators	Prediction errors; innovation statistics; physics residuals	Mode switching; fallback control; re-identification triggers	[11, 12, 206–208]
Context uncertainty $p(c_t \mid h_t)$ or $\mathcal{C}_t$	Posterior inference over physical context; context ensembles	Conservative constraint deformation during regime transitions	[81, 85, 124, 209–211]

high-frequency control loops required in safety-critical systems. In practice, deployments frequently navigate this challenge via *multi-rate architectures*, where computationally intensive uncertainty quantification is performed at a lower frequency than the inner control loop. We further address these real-time implementation trade-offs in Section 7.5.

6. EVALUATION CHECKLIST AND REPORTING PROTOCOL

As research on physical AI for non-stationary environments continues to expand, comparing methods across papers has become increasingly difficult. Differences in task setups, non-stationarity sources, safety definitions, and evaluation metrics often obscure the true strengths and limitations of proposed approaches. In this section, we propose an evaluation checklist and reporting protocol aimed at improving transparency, comparability, and reproducibility for safety-critical modeling and control under non-stationary physics.

The goal of this checklist is not to prescribe a single benchmark or experimental setup, but to identify a minimal set of questions that authors should address to clarify what their method handles well, where it may fail, and under what assumptions its safety claims hold.

6.1. Characterization of non-stationarity

A first requirement for meaningful evaluation is a clear characterization of the non-stationarity being considered. Authors should specify:

- **Source of non-stationarity**: e.g., parametric drift (friction, mass, wear), regime switching (contact modes, payload changes), external interventions (human interaction), or structural changes.
- **Temporal profile**: whether changes are abrupt, gradual, periodic, or adversarial.
- **Observability**: which aspects of the non-stationarity are directly observed, indirectly inferred, or entirely latent.

Explicitly stating these factors helps distinguish between methods designed for mild distribution shift and those intended for genuinely deceptive or safety-critical regime changes.

6.2. Context inference and identifiability

When a method relies on inferred physical context or latent variables, evaluation should explicitly address identifiability and inference reliability. In safety-critical settings, failures often arise not from poor nominal prediction, but from incorrect or ambiguous attribution of physical context. We therefore recommend that authors report the following items:

- **Inference paradigm:** whether physical context is inferred passively from observations, actively through action–reaction signals, or via a hybrid strategy.
- **Excitation and safety compatibility:** what actions or excitations are required for reliable context inference, and whether these actions respect safety constraints throughout deployment.
- **Identifiability limits:** failure cases in which distinct physical regimes are observationally indistinguishable within the admissible operating envelope, including whether identifiability is global or only local to specific regimes.
- **Intervention specification:** a clear description of what physical factors are intervened upon to induce non-stationarity (e.g., friction, payload, contact mode), and whether these changes are exogenous or action-induced.
- **Appearance–physics mismatch stress tests:** controlled contrasts in which visual appearance varies while physical properties are fixed (and vice versa), to expose shortcut features and spurious correlations in context inference.
- **Sensitivity analysis:** quantitative or qualitative analysis of how downstream safety metrics (e.g., constraint violation rate, fallback frequency) degrade under context misidentification.

Whenever possible, ablation studies that isolate the effect of context inference accuracy on safety outcomes are strongly encouraged, as they provide critical insight into whether safety improvements stem from correct physical attribution or incidental robustness. Taken together, these items emphasize that context inference should be evaluated not only by prediction accuracy, but by causal identifiability under safety constraints.

6.3. Dynamics modeling and physics consistency

For context-modulated or physics-informed dynamics models, evaluation should clarify how physical structure influences performance and safety:

- Which physical priors are enforced (loss-based, architectural, or constraint-based).
- How violations of physics consistency are measured and monitored.
- Whether improved prediction accuracy translates into improved safety-relevant quantities (e.g., stopping distance, collision risk).
- Robustness of learned dynamics under extrapolation to unseen physical regimes.

Reporting physics residuals alongside standard prediction errors can provide valuable insight into model behavior under shift.

6.4. Uncertainty quantification and calibration

Because safety under non-stationarity hinges on uncertainty awareness, evaluation should include:

- The form of uncertainty output (distribution, ensemble, set, ambiguity set).
- Calibration diagnostics appropriate to the chosen representation (e.g., coverage, reliability).
- Behavior of uncertainty estimates before, during, and after regime changes.
- Interaction between uncertainty signals and safety mechanisms (e.g., constraint tightening, fallback triggering).

Methods that report uncertainty without demonstrating its calibration or downstream use should be interpreted with caution in safety-critical settings.

6.5. Safety metrics and failure analysis

Safety evaluation should go beyond average performance metrics. We recommend reporting:

- Explicit safety constraints and violation definitions.
- Worst-case or tail-risk statistics (e.g., maximum violation, CVaR).
- Frequency and severity of safety interventions or fallback activations.
- Qualitative failure analyses highlighting when and why safety mechanisms fail.

Such analyses help distinguish between methods that fail gracefully and those that fail catastrophically.

6.6. Deployment considerations

Finally, for methods targeting real-world deployment, evaluation should consider:

- Computational overhead and real-time feasibility.
- Multi-rate architectures and latency between inference, adaptation, and control.
- Behavior under partial system failure or degraded sensing.
- Assumptions required for safe operation in practice.

Explicit discussion of these factors helps bridge the gap between laboratory demonstrations and deployment-ready physical AI systems.

Summary. Together, these checklist items provide a structured lens for evaluating safety-critical learning systems under non-stationary physics. While not all items will be relevant for every contribution, addressing them transparently can significantly improve the interpretability and comparability of experimental results, and accelerate progress toward reliable deployment.

7. OPEN PROBLEMS AND RESEARCH AGENDA

Despite significant recent progress, building safety-critical physical AI systems that operate reliably under non-stationary environments remains an open and fundamentally challenging problem. In this section, we highlight key open issues that emerge from the preceding review and outline a research agenda toward deployment-ready systems.

7.1. Causal identifiability under safety constraints

A central challenge lies in identifying safety-relevant physical context when informative excitation is itself constrained by safety requirements. While active and hybrid inference approaches improve identifiability relative to passive observation, they often rely on probing actions whose informativeness competes with safety [60, 254–256]. Developing principled frameworks that characterize *what can be identified safely*, and under which operating envelopes, remains largely open. This includes understanding identifiability under partial observability, limited excitation, and human-in-the-loop interaction, as well as designing actions that are maximally informative subject to hard safety constraints.

7.2. Disentanglement and verification of uncertainty sources

While the theoretical distinction between aleatoric, epistemic, and structural uncertainty is clear (as discussed in Section 5.1), developing algorithmic mechanisms to automatically *disentangle* and *verify* these sources from raw data remains an open challenge [257–260]. Current methods often conflate these sources; for instance, high ensemble variance can result equally from a novel physical regime (structural uncertainty) or simply from noisy data in a known regime (aleatoric uncertainty). This ambiguity leads to suboptimal or unsafe downstream responses: while aleatoric noise should be handled via probabilistic chance constraints, epistemic or structural gaps require model adaptation or conservative fallback. Developing rigorous verification tools that can distinguish "I don't know" (epistemic) from "The world has changed" (structural) is a critical step toward *competency-aware* physical AI.

7.3. Uncertainty representations beyond point estimates

Most existing systems ultimately commit to a single adapted model or context estimate, even when uncertainty remains high. As discussed in Sections 5.1–5.4, such commitment can be unsafe under regime transitions or ambiguous observations [197, 261]. An open research direction is the development of uncertainty representations that persist through the modeling and control stack, enabling decision making over sets or distributions of plausible dynamics rather than point estimates. This includes scalable methods for propagating uncertainty through context-modulated dynamics and translating it into conservative yet nontrivial safety behavior.

7.4. Calibration and robustness under evolving regimes

Ensuring calibration of uncertainty estimates over long horizons remains an unresolved problem in non-stationary environments. While adaptive and online calibration methods offer partial solutions, their interaction with regime changes, structural shifts, and delayed feedback is poorly understood [217, 262, 263]. Future work should investigate principled mechanisms for maintaining reliability guarantees under continual distribution shift, as well as detecting when calibration itself has failed. This is especially critical when uncertainty estimates directly influence safety margins or robustness parameters.

7.5. Bridging modeling fidelity and real-time deployment

Physics-informed and context-modulated models often increase computational complexity, raising concerns about real-time feasibility and robustness to latency [12, 264]. Current mitigation strategies face fundamental limitations in safety-critical settings. For instance, reduced ensemble methods approximate epistemic uncertainty by evaluating only a small subset of models, significantly lowering inference costs; however, they often under-represent uncertainty in out-of-distribution regimes where ensemble diversity is most needed [100, 265]. Similarly, amortized or distilled uncertainty surrogates compress expensive predictors into lightweight models, yet their estimates are tied to the training distribution and risk becoming overconfident under novel contexts or sensing degradation [266]. While multi-rate architectures partially address latency issues by decoupling inference from control, they introduce distinct failure modes arising from timing mismatches and stale context estimates [85, 264, 267]. A key open question is how to systematically co-design modeling, inference, and execution layers such that safety guarantees are preserved despite asynchronous updates, partial failures, or degraded sensing.

7.6. Evaluation standards and failure transparency

As emphasized in Section 6, the lack of standardized evaluation protocols showing how and why methods fail remains a major obstacle to progress [268, 269]. Future benchmarks should go beyond average performance and explicitly stress-test systems under deceptive, adversarial, or rare physical regimes. Equally important is transparent reporting of failure cases and near-miss events, which can provide valuable insight into the limits of current approaches and guide safer system design.

7.7. Toward deployment-ready physical AI

Ultimately, the goal of research in this area is not merely improved modeling accuracy, but the deployment of systems that fail gracefully and predictably under uncertainty. Achieving this requires a shift from optimizing nominal performance toward designing architectures that explicitly reason about model validity, uncertainty, and safety throughout their lifetime. We believe that integrating physics-informed causal modeling, uncertainty-aware reasoning, and principled robustness mechanisms provides a promising foundation for this vision, but substantial theoretical and practical challenges remain.

Outlook. Addressing these open problems will require closer integration between systems and control theory, machine learning, and robotics. By grounding learning-based methods in physical structure, causal reasoning, and actionable uncertainty, future work can move toward physical AI systems that operate safely not only when conditions are favorable, but also when they are deceptive, evolving, and uncertain.

8. CONCLUSION

Non-stationary physical environments remain a fundamental barrier to the reliable deployment of learning-enabled controllers in safety-critical systems. Across the literature reviewed in this paper, a recurring theme is that safety failures under shift are often driven not by poor nominal performance, but by an inability to infer and track the physical causes that deform dynamics and safety constraints over time. This observation motivates a modeling-centric perspective: rather than treating non-stationarity as unstructured noise, we should represent it through safety-relevant physical context and incorporate it coherently into dynamics modeling.

This review organized recent progress through four connected lenses. We first characterized types of non-stationarity in physical systems and emphasized their implications for safety-constraint geometry. We then surveyed methods for inferring safety-relevant physical context, highlighting the identifiability gap of purely passive inference and the role of action–reaction evidence and hybrid strategies. Next, we reviewed context-modulated dynamics modeling and physics-informed structure, clarifying how physics priors can be injected via losses, architectures, or explicit constraints. Finally, we synthesized uncertainty quantification under non-stationarity, focusing on uncertainty representations that are actionable for safety mechanisms, including constraint tightening, risk-sensitive formulations, and distributionally robust reasoning.

Beyond summarizing methods, we proposed an evaluation checklist and reporting protocol to improve comparability and transparency across studies, and we outlined open problems and research directions toward deployment-ready physical AI. We hope that this perspective encourages tighter integration between causal reasoning, physics-informed modeling, and uncertainty-aware safety enforcement, ultimately enabling systems that not only perform well under nominal conditions but also remain reliable when environments are evolving, deceptive, and uncertain.

DECLARATIONS

Conflict of Interest

The authors declare that there is no competing financial interest or personal relationship that could have appeared to influence the work reported in this paper.

Authors' Contributions

Sung Kwon On and Sukchul Jeong drafted the manuscript and prepared the figures. Jungjin Lee contributed to the writing. Insoon Yang designed the review and supervised the writing.

Funding

This work was supported in part by the Information and Communications Technology Planning and Evaluation (IITP) grants funded by MSIT No. 2022-0-00124, No. 2022-0-00480 and No. RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University).

REFERENCES

[1] I. Radosavovic, T. Xiao, B. Zhang, T. Darrell, J. Malik, and K. Sreenath, "Real-world humanoid locomotion with reinforcement learning," *Science Robotics*, vol. 9, no. 89, p. eadi9579, 2024.

[2] R. Sinha, A. Sharma, S. Banerjee, T. Lew, R. Luo, S. M. Richards, Y. Sun, E. Schmerling, and M. Pavone, "A system-level view on out-of-distribution data in robotics," *arXiv preprint arXiv:2212.14020*, 2023.

[3] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1749–1767, 2023.

[4] H.-Y. Tung, M. Ding, Z. Chen, D. Bear, C. Gan, J. Tenenbaum, D. Yamins, J. Fan, and K. Smith, "Physion++: Evaluating physical scene understanding that requires online inference of different physical properties," in *Advances in Neural Information Processing Systems*, 2023, pp. 67 048–67 068.

[5] J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, and G. S. Sukhatme, "Interactive perception: Leveraging action in perception and perception in action," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1273–1291, 2017.

[6] D. Zheng, V. Luo, J. Wu, and J. B. Tenenbaum, "Unsupervised learning of latent physical properties using perception-prediction networks," in *Conference on Uncertainty in Artificial Intelligence*, 2018, p. 497–507.

[7] Z. Xu, J. Wu, A. Zeng, J. B. Tenenbaum, and S. Song, "Densephysnet: Learning dense physical object representations via multi-step dynamic interactions," in *Robotics: Science and Systems*, 2019, pp. 1–10.

[8] Y. Li, "Learning structured world models from and for physical interactions," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025, pp. 28 717–28 718.

[9] T. X. Nghiem, J. Drgoňa, C. Jones, Z. Nagy, R. Schwan, B. Dey, A. Chakrabarty, S. Di Cairano, J. A. Paulson, A. Carron, M. N. Zeilinger, W. Shaw Cortez, and D. L. Vrabie, "Physics-informed machine learning for modeling and control of dynamical systems," in *2023 American Control Conference*, 2023, pp. 3735–3750.

[10] N. Agarwal, A. Ali, M. Bala, Y. Balaji, E. Barker, T. Cai, P. Chattopadhyay, Y. Chen, Y. Cui, Y. Ding *et al.*, "Cosmos world foundation model platform for physical AI," *arXiv preprint arXiv:2501.03575*, 2025.

[11] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Prentice Hall, 1999.

[12] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Reviews Physics*, vol. 3, no. 6, pp. 422–440, 2021.

[13] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review Control, Robotics, and Autonomous Systems.*, vol. 5, pp. 411–444, 2022.

[14] G. Pillonetto, A. Aravkin, D. Gedon, L. Ljung, A. H. Ribeiro, and T. B. Schön, "Deep networks for system identification: A survey," *Automatica*, vol. 171, p. 111907, 2025.

[15] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-jacobi reachability: A brief overview and recent advances," in *2017 IEEE 56th Annual Conference on Decision and Control*, 2017, pp. 2242–2253.

[16] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference*, 2019, pp. 3420–3431.

[17] K. P. Wabersich, A. J. Taylor, J. J. Choi, K. Sreenath, C. J. Tomlin, A. D. Ames, and M. N. Zeilinger, "Data-driven safety filters: Hamilton-jacobi reachability, control barrier functions, and predictive methods for uncertain systems," *IEEE Control Systems Magazine*, vol. 43, no. 5, pp. 137–177, 2023.

[18] S. Hwang, I. Jang, D. Kim, and H. J. Kim, "Safe motion planning and control for mobile robots: A survey," *International Journal of Control, Automation and Systems*, vol. 22, p. 2955–2969, 2024.

[19] J. Pearl, *Causality: models, reasoning, and inference*. Cambridge University Press, 2000.

[20] X. Zheng, B. Aragam, P. K. Ravikumar, and E. P. Xing, "Dags with no tears: Continuous optimization for structure learning," in *Advances in Neural Information Processing Systems*, 2018.

[21] A. Zanga, E. Ozkirimli, and F. Stella, "A survey on causal discovery: theory and practice," *International Journal of Approximate Reasoning*, vol. 151, pp. 101–129, 2022.

[22] J. Gama, I. Žliobaitundefined, A. Bifet, M. Pechenizkiy, and A. Bouchachia, "A survey on concept drift adaptation," *ACM Computing Surveys*, vol. 46, no. 4, 2014.

[23] S. Padakandla, "A survey of reinforcement learning algorithms for dynamically varying environments," *ACM Computing Surveys*, vol. 54, no. 6, 2021.

[24] Z. Liu, J. Lu, J. Xuan, and G. Zhang, "Learning latent and changing dynamics in real non-stationary environments," *IEEE Transactions on Knowledge and Data Engineering*, vol. 37, no. 4, pp. 1930–1942, 2025.

[25] P. A. Ioannou and J. Sun, *Robust Adaptive Control*. Prentice-Hall, Inc., 1995.

[26] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[27] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017.

[28] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2018, pp. 3803–3810.

[29] Z. Wu and Y. Bai, "Adaptive super-twisting sliding mode control for marine electric propulsion system based on load observer," *International Journal of Control, Automation and Systems*, 2026.

[30] K. Lee, Y. Seo, S. Lee, H. Lee, and J. Shin, "Context-aware dynamics model for generalization in model-based reinforcement learning," in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 5757–5766.

[31] J. Kwon, Y. Efroni, C. Caramanis, and S. Mannor, "Rl for latent mdps: Regret guarantees and a lower bound," in *Advances in Neural Information Processing Systems*, 2021, pp. 24 523–24 534.

[32] C. Tomlin, J. Lygeros, and S. Shankar Sastry, "A game theoretic approach to controller design for hybrid systems," *Proceedings of the IEEE*, vol. 88, no. 7, pp. 949–970, 2000.

[33] D. Liberzon, *Switching in Systems and Control*. Birkhäuser Boston, MA, 2003.

[34] R. Goebel, R. G. Sanfelice, and A. R. Teel, "Hybrid dynamical systems," *IEEE Control Systems Magazine*, vol. 29, no. 2, pp. 28–93, 2009.

[35] T. Marcucci, R. Deits, M. Gabiccini, A. Bicchi, and R. Tedrake, "Approximate hybrid model predictive control for multi-contact push recovery in complex environments," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics*, 2017, pp. 31–38.

[36] L. Lindemann, H. Hu, A. Robey, H. Zhang, D. Dimarogonas, S. Tu, and N. Matni, "Learning hybrid control barrier functions from data," in *Proceedings of the 2020 Conference on Robot Learning*, 2021, pp. 1351–1370.

[37] P. M. Wensing, M. Posa, Y. Hu, A. Escande, N. Mansard, and A. D. Prete, "Optimization-based control for dynamic legged robots," *IEEE Transactions on Robotics*, vol. 40, pp. 43–63, 2024.

[38] S. Paoletti, A. L. Juloski, G. Ferrari-Trecate, and R. Vidal, "Identification of hybrid systems a tutorial," *European Journal of Control*, vol. 13, no. 2, pp. 242–260, 2007.

[39] M. Poli, S. Massaroli, L. Scimeca, S. Chun, S. J. Oh, A. Yamashita, H. Asama, J. Park, and A. Garg, "Neural hybrid automata: Learning dynamics with multiple modes and stochastic transitions," in *Advances in Neural Information Processing Systems*, 2021.

[40] C. N. Mavridis and K. H. Johansson, "Real-time switched system identification with online deterministic annealing," *IEEE Transactions on Automatic Control*, pp. 1–12, 2025.

[41] P. Hernandez-Leal, M. Kaisers, T. Baarslag, and E. M. De Cote, "A survey of learning in multiagent environments: Dealing with non-stationarity," *arXiv preprint arXiv:1707.09183*, 2017.

[42] S. Ahlberg and D. V. Dimarogonas, "Human-in-the-loop control synthesis for multi-agent systems under hard and soft metric interval temporal logic specifications," in *2019 IEEE 15th International Conference on Automation Science and Engineering*, 2019, pp. 788–793.

[43] D. Hughes, A. Agarwal, Y. Guo, and K. Sycara, "Inferring non-stationary human preferences for human-agent teams," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication*, 2020, pp. 1178–1185.

[44] Y. Chandak, S. Shankar, N. Bastian, B. da Silva, E. Brunskill, and P. S. Thomas, "Off-policy evaluation for action-dependent non-stationary environments," in *Advances in Neural Information Processing Systems*, 2022, pp. 9217–9232.

[45] Y. Bengio, T. Deleu, N. Rahaman, N. R. Ke, S. Lachapelle, O. Bilaniuk, A. Goyal, and C. Pal, "A meta-transfer objective for learning to disentangle causal mechanisms," in *International Conference on Learning Representations*, 2020.

[46] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, and Y. Bengio, "Toward causal representation learning," *Proceedings of the IEEE*, vol. 109, no. 5, pp. 612–634, 2021.

[47] J. Wu, J. J. Lim, H. Zhang, J. B. Tenenbaum, and W. T. Freeman, "Physics 101: Learning physical object properties from unlabeled videos," in *27th British Machine Vision Conference*, 2016, p. 39.1 – 39.12.

[48] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, "Shortcut learning in deep neural networks," *Nature Machine Intelligence*, vol. 2, pp. 665–673, 2020.

[49] Z. Zhao, Y. Li, W. Li, Z. Qi, L. Ruan, Y. Zhu, and K. Althoefer, "Tac-man: Tactile-informed prior-free manipulation of articulated objects," *IEEE Transactions on Robotics*, vol. 41, pp. 538–557, 2025.

[50] P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine, "Learning to poke by poking: Experiential learning of intuitive physics," vol. 29, 2016.

[51] M. Blanke, M. Kinnaert, J. Lunze, and M. Staroswiecki, *Diagnosis and Fault-Tolerant Control*. Springer Berlin, Heidelberg, 2003.

[52] L. Balzano and R. Nowak, "Blind calibration of sensor networks," in *Proceedings of the 6th International Conference on Information Processing in Sensor Networks*, 2007, p. 79–88.

[53] A. Vergara, S. Vembu, T. Ayhan, M. A. Ryan, M. L. Homer, and R. Huerta, "Chemical gas sensor drift compensation using classifier ensembles," *Sensors and Actuators B: Chemical*, vol. 166-167, pp. 320–329, 2012.

[54] P. Scholl, M. Iskandar, S. Wolf, J. Lee, A. Bacho, A. Dietrich, A. Albu-Schäffer, and G. Kutyniok, "Learning-based adaption of robotic friction models," *Robotics and Computer-Integrated Manufacturing*, vol. 89, p. 102780, 2024.

[55] K. J. Astrom and B. Wittenmark, *Adaptive Control*. Addison-Wesley Longman Publishing Co., Inc., 1994.

[56] R. Isermann, "Supervision, fault-detection and fault-diagnosis methods — an introduction," *Control Engineering Practice*, vol. 5, no. 5, pp. 639–652, 1997.

[57] D. Ordoñez Apraez, V. Kostic, G. Turrisi, P. Novelli, C. Mastalli, C. Semini, and M. Pontil, "Dynamics harmonic analysis of robotic systems: Application in data-driven Koopman modelling," in *Proceedings of the 6th Annual Learning for Dynamics amp; Control Conference*, vol. 242, 2024, pp. 1318–1329.

[58] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, "Examples when nonlinear model predictive control is nonrobust," *Automatica*, vol. 40, no. 10, pp. 1729–1738, 2004.

[59] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.

[60] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2018.

[61] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice-Hall, Inc., 1996.

[62] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*, 1999, pp. 207–226.

[63] D. Q. Mayne, M. M. Seron, and S. V. Raković, "Robust model predictive control of constrained linear systems with bounded disturbances," *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.

[64] H. A. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with markovian switching coefficients," *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, 2002.

[65] K. Narendra and J. Balakrishnan, "Adaptive control using multiple models," *IEEE Transactions on Automatic Control*, vol. 42, no. 2, pp. 171–187, 1997.

[66] Y. Bar-Shalom, T. Kirubarajan, and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., 2002.

[67] H. Hjalmarsson, "From experiment design to closed-loop control," *Automatica*, vol. 41, no. 3, pp. 393–438, 2005.

[68] R. G. Krishnan, U. Shalit, and D. Sontag, "Deep Kalman filters," *arXiv preprint: arXiv:1511.05121*, 2015.

[69] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum, "Galileo: Perceiving physical object properties by integrating a physics engine with deep learning," in *Advances in Neural Information Processing Systems*, 2015.

[70] P. Battaglia, R. Pascanu, M. Lai, D. Jimenez Rezende, and k. kavukcuoglu, "Interaction networks for learning about objects, relations and physics," in *Advances in Neural Information Processing Systems*, 2016.

[71] M. Brandão, K. Hashimoto, and A. Takanishi, "Friction from vision: A study of algorithmic and human performance with consequences for robot perception and teleoperation," in *IEEE-RAS International Conference on Humanoid Robots*, 2016, pp. 428–435.

[72] M. Karl, M. Soelch, J. Bayer, and P. van der Smagt, "Deep variational bayes filters: Unsupervised learning of state space models from raw data," *arXiv preprint, arXiv: 1605.06432*, 2016.

[73] R. Riochet, M. Y. Castro, M. Bernard, A. Lerer, R. Fergus, V. Izard, and E. Dupoux, "Intphys: A framework and benchmark for visual intuitive physics reasoning," *arXiv preprint arXiv:1803.07616*, 2018.

[74] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 8459–8468.

[75] S. Aguilera, M. A. Murtaza, Y. Zhao, and S. Hutchinson, "Mass estimation of a moving object through minimal manipulation interaction," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 6461–6467.

[76] Q. Garrido, N. Ballas, M. Assran, A. Bardes, L. Najman, M. Rabbat, E. Dupoux, and Y. LeCun, "Intuitive physics understanding emerges from self-supervised pretraining on natural videos," *arXiv preprint arXiv:2502.11831*, 2025.

[77] J. Swevers, C. Ganseman, D. B. Tukel, J. De Schutter, and H. Van Brussel, "Optimal robot excitation and identification," *IEEE Transactions on Robotics and Automation*, vol. 13, no. 5, pp. 730–740, 2002.

[78] H. Jansson and H. Hjalmarsson, "Input design via lmis admitting frequency-wise model specifications in confidence regions," *IEEE Transactions on Automatic Control*, vol. 50, no. 10, pp. 1534–1549, 2005.

[79] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Information gathering actions over human internal state," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 66–73.

[80] F. Berkenkamp, M. Turchetta, A. Schoellig, and A. Krause, "Safe model-based reinforcement learning with stability guarantees," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[81] W. Yu, J. Tan, C. K. Liu, and G. Turk, "Preparing for the unknown: Learning a universal policy with online system identification," in *Robotics: Science and Systems*, 2017.

[82] S. Dean, S. Tu, N. Matni, and B. Recht, "Safely learning to control the constrained linear quadratic regulator," in *2019 American Control Conference*, 2019, pp. 5582–5588.

[83] S. Huh and I. Yang, "Safe reinforcement learning for probabilistic reachability and safety specifications: A lyapunov-based approach," *arXiv preprint arXiv:2002.10126*, 2020.

[84] O. Ahmed, F. Träuble, A. Goyal, A. Neitz, Y. Bengio, B. Schölkopf, M. Wüthrich, and S. Bauer, "CausalWorld: A robotic manipulation benchmark for causal structure and transfer learning," *arXiv preprint arXiv:2010.04296*, 2020.

[85] A. Kumar, Z. Fu, D. Pathak, and J. Malik, "RMA: Rapid motor adaptation for legged robots," in *Robotics: Science and Systems*, 2021.

[86] T. Lee, B. D. Lee, and F. C. Park, "Optimal excitation trajectories for mechanical systems identification," *Automatica*, vol. 131, p. 109773, 2021.

[87] C. C. Matl, "Interactive perception for robotic manipulation of liquids, grains, and doughs," Ph.D. dissertation, University of California, Berkeley, 2021.

[88] Y. Yu, S. Talebi, H. J. Van Waarde, U. Topcu, M. Mesbahi, and B. Açıkmeşe, "On controllability and persistency of excitation in data-driven control: Extensions of willems' fundamental lemma," in *Proceedings of the IEEE Conference on Decision and Control*, 2021, pp. 6485–6490.

[89] T. Lew, A. Sharma, J. Harrison, A. Bylard, and M. Pavone, "Safe active dynamics learning and control: A sequential exploration–exploitation framework," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 2888–2907, 2022.

[90] J. H. Park, G. P. Dalwankar, A. Bartsch, A. George, and A. B. Farimani, "Fluid viscosity prediction leveraging computer vision and robot interaction," *arXiv preprint arXiv:2308.02715*, 2023.

[91] P. Wu, A. Escontrela, D. Hafner, P. Abbeel, and K. Goldberg, "Daydreamer: World models for physical robot learning," in *Proceedings of the Conference on Robot Learning*, 2023, pp. 2226–2240.

[92] M. Lange-Hegermann and C. Zimmer, "Future-aware safe active learning of time-varying systems using gaussian processes," *arXiv preprint arXiv:2405.10581*, 2024.

[93] M. Memmel, A. Wagenmaker, C. Zhu, P. Yin, D. Fox, and A. Gupta, "ASID: Active exploration for system identification in robotic manipulation," in *International Conference on Learning Representations*, 2024.

[94] B. Zhang, D. Haugk, and R. Vasudevan, "System identification for constrained robots," *arXiv preprint arXiv:2408.08830*, 2024.

[95] J. Wang, Y. Yuan, H. Che, H. Qi, Y. Ma, J. Malik, and X. Wang, "Lessons from learning to spin pens," *arXiv preprint arXiv:2407.18902*, 2024.

[96] B. Zhang, Z. Zhou, and R. Vasudevan, "Provably-safe, online system identification," *arXiv preprint arXiv:2504.21486*, 2025.

[97] Y. Wei, Z. Yi, H. Li, S. Soedarmadji, and Y. Sui, "Safe bayesian optimization for the control of high-dimensional embodied systems," in *Conference on Robot Learning*, 2025, pp. 4771–4792.

[98] S. Khorshidi, M. Dawood, B. Nederkorn, M. Bennewitz, and M. Khadiv, "Physically-consistent parameter identification of robots in contact," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2025, pp. 677–683.

[99] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, 2017.

[100] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," in *Advances in Neural Information Processing Systems*, 2018.

[101] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, "Invariant risk minimization," *arXiv preprint, arXiv: 1907.02893*, 2019.

[102] K. Kim, M. Sano, J. De Freitas, N. Haber, and D. Yamins, "Active world model learning with progress curiosity," in *Proceedings of the International Conference on Machine Learning*, 2020, pp. 5306–5315.

[103] T. Wang, S. S. Du, A. Torralba, P. Isola, A. Zhang, and Y. Tian, "Denoised MDPs: Learning world models better than the world itself," *arXiv preprint arXiv:2206.15477*, 2022.

[104] N. Dirkx, M. Bosselaar, and T. Oomen, "Peak amplitude-constrained experiment design for FRF identification of MIMO motion systems," in *Proceedings of the IEEE International Conference on Advanced Motion Control*, 2022, pp. 256–261.

[105] T. Gupta, W. Gong, C. Ma, N. Pawlowski, A. Hilmkil, M. Scetbon, M. Rigter, A. Famoti, A. J. Llorens, J. Gao *et al.*, "The essential role of causality in foundation world models for embodied AI," *arXiv preprint arXiv:2402.06665*, 2024.

[106] Y. Shang, X. Zhang, Y. Tang, L. Jin, C. Gao, W. Wu, and Y. Li, "Roboscape: Physics-informed embodied world model," in *arXiv preprint arXiv:2506.23135*, 2025.

[107] S. Prasanna, K. Farid, R. Rajan, and A. Biedenkapp, "Learning contextual world models aids zero-shot generalization under observed contexts," in *Reinforcement Learning Conference and Reinforcement Learning Journal*, 2024.

[108] W. Li, H. Zhao, Z. Yu, Y. Du, Q. Zou, R. Hu, and K. Xu, "PIN-wm: Learning physics-informed world models for non-prehensile manipulation," *arXiv preprint arXiv:2504.16693*, 2025.

[109] G. Zollicoffer, T. Chopra, M. Yan, X. Ma, K. Eaton, and M. Riedl, "World model robustness via surprise recognition," *arXiv preprint arXiv:2512.01119*, 2025.

[110] Z. Mao, M. E. Umasudhan, and I. Ruchkin, "How safe will I be given what I saw? calibrated prediction of safety chances for image-controlled autonomy," *arXiv preprint arXiv:2508.09346*, 2025.

[111] P. De Haan, D. Jayaraman, and S. Levine, "Causal confusion in imitation learning," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[112] M. A. Alcorn, Q. Li, Z. Gong, C. Wang, L. Mai, W.-S. Ku, and A. Nguyen, "Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4845–4854.

[113] S. Pfrommer, M. Halm, and M. Posa, "Contactnets: Learning discontinuous contact dynamics with smooth, implicit representations," in *Conference on Robot Learning*, 2021, pp. 2279–2291.

[114] D. Liu, J. Zhang, A.-D. Dinh, E. Park, S. Zhang, A. Mian, M. Shah, and C. Xu, "Generative physical AI in vision: A survey," *arXiv preprint arXiv:2501.10928*, 2025.

[115] A. Wagenmaker, G. Shi, and K. G. Jamieson, "Optimal exploration for model-based rl in nonlinear systems," in *Advances in Neural Information Processing Systems*, 2023, pp. 15 406–15 455.

[116] C. Lu, Y. Wu, J. M. Hernández-Lobato, and B. Schölkopf, "Invariant causal representation learning for out-of-distribution generalization," in *International Conference on Learning Representations*, 2021.

[117] B. Sadeghi, S. Dehdashtian, and V. Boddeti, "On characterizing the trade-off in invariant representation learning," in *International Conference on Learning Representations*, 2024.

[118] V. Y. Nastl and M. Hardt, "Do causal predictors generalize better to new domains?" in *Advances in Neural Information Processing Systems*, 2024, pp. 31 202–31 315.

[119] Y. Yang, B. Huang, F. Feng, X. Wang, S. Tu, and L. Xu, "Towards generalizable reinforcement learning via causality-guided self-adaptive representations," in *Proceeding of the 30th International Conference on Learning Representations*, 2025.

[120] S. Pitis, E. Creager, A. Mandlekar, and A. Garg, "Mocoda: Model-based counterfactual data augmentation," in *Advances in Neural Information Processing Systems*, 2022, pp. 18 143–18 156.

[121] Z. Sun, B. He, J. Liu, X. Chen, C. Ma, and S. Zhang, “Offline imitation learning with variational counterfactual reasoning,” in *Advances in Neural Information Processing Systems*, 2023, pp. 43 729–43 741.

[122] A. Feder, Y. Wald, C. Shi, S. Saria, and D. Blei, “Data augmentations for improved (large) language model generalization,” in *Advances in Neural Information Processing Systems*, 2023, pp. 70 638–70 653.

[123] N. Armengol Urpí, M. Bagatella, M. Vlastelica, and G. Martius, “Causal action influence aware counterfactual data augmentation,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 1709–1729.

[124] K. Rakelly, A. Zhou, C. Finn, S. Levine, and D. Quillen, “Efficient off-policy meta-reinforcement learning via probabilistic context variables,” in *Proceedings of the International Conference on Machine Learning*, 2019, pp. 5331–5340.

[125] S. Buchholz, G. Rajendran, E. Rosenfeld, B. Aragam, B. Schölkopf, and P. Ravikumar, “Learning linear causal representations from interventions under general nonlinear mixing,” in *Advances in Neural Information Processing Systems*, 2023, pp. 45 419–45 462.

[126] D. Yao, D. Xu, S. Lachapelle, S. Magliacane, P. Taslakian, G. Martius, J. von Kügelgen, and F. Locatello, “Multi-view causal representation learning with partial observability,” in *The 12th International Conference on Learning Representations*, 2024.

[127] I. Ng, X. Dong, H. Dai, B. Huang, P. Spirtes, and K. Zhang, “Score-based causal discovery of latent variable causal models,” in *Forty-first International Conference on Machine Learning*, 2024.

[128] S. Li, K. Wu, C. Zhang, and Y. Zhu, “I-phyre: Interactive physical reasoning,” in *International Conference on Learning Representations*, 2024.

[129] N. Sobanbabu, G. He, T. He, Y. Yang, and G. Shi, “Sampling-based system identification with active exploration for legged sim2real learning,” in *Proceedings of The 9th Conference on Robot Learning*, 2025, pp. 578–598.

[130] A. M. Alaa, Z. Ahmad, and M. van der Laan, “Conformal meta-learners for predictive inference of individual treatment effects,” in *Advances in Neural Information Processing Systems*, 2023, pp. 47 682–47 703.

[131] J. Ai and Z. Ren, “Not all distributional shifts are equal: Fine-grained robust conformal inference,” in *Proceedings of the 41st International Conference on Machine Learning*, 2024, pp. 641–665.

[132] Z. Zhou, M. Q. Elahi, and M. Kocaoglu, “Sample efficient bayesian learning of causal graphs from interventions,” in *Advances in Neural Information Processing Systems*, 2024, pp. 9178–9204.

[133] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, “Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2018, pp. 7559–7566.

[134] T. Salzmann, E. Kaufmann, J. Arrizabalaga, M. Pavone, D. Scaramuzza, and M. Ryll, “Real-time neural MPC: Deep learning model predictive control for quadrotors and agile robotic platforms,” *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2397–2404, 2023.

[135] T. M. Moerland, J. Broekens, A. Plaat, C. M. Jonker *et al.*, “Model-based reinforcement learning: A survey,” *Foundations and Trends® in Machine Learning*, vol. 16, no. 1, pp. 1–118, 2023.

[136] M. Deisenroth and C. E. Rasmussen, “Pilco: A model-based and data-efficient approach to policy search,” in *Proceedings of the 28th International Conference on machine learning*, 2011, pp. 465–472.

[137] V. Verdult and M. Verhaegen, “Subspace identification of multivariable linear parameter-varying systems,” *Automatica*, vol. 38, no. 5, pp. 805–814, 2002.

[138] T. Lee, J. Kwon, P. M. Wensing, and F. C. Park, “Robot model identification and learning: A modern perspective,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2023.

[139] J. Shin, A. Hakobyan, M. Park, Y. Kim, G. Kim, and I. Yang, “Infusing model predictive control into meta-reinforcement learning for mobile robots in dynamic environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 065–10 072, 2022.

[140] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, “In-hand object rotation via rapid motor adaptation,” in *Proceedings of the Conference on Robot Learning*, 2023, pp. 1722–1732.

[141] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, “Neural ordinary differential equations,” in *Advances in Neural Information Processing Systems*, 2018.

[142] M. Raissi, P. Perdikaris, and G. E. Karniadakis, “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations,” *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.

[143] G. Pizzuto and M. Mistry, “Physics-penalised regularisation for learning dynamics models with contact,” in *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, vol. 144, 2021, pp. 611–622.

[144] J. Yu, L. Lu, X. Meng, and G. E. Karniadakis, “Gradient-enhanced physics-informed neural networks for forward and inverse pde problems,” *Computer Methods in Applied Mechanics and Engineering*, vol. 393, p. 114823, 2022.

[145] M. Haghighatlari, J. Li, X. Guan, O. Zhang, A. Das, C. J. Stein, F. Heidar-Zadeh, M. Liu, M. Head-Gordon, L. Bertels *et al.*, “NewtonNet: A newtonian message passing network for deep learning of interatomic potentials and forces,” *Digital Discovery*, vol. 1, no. 3, pp. 333–343, 2022.

[146] M. Lutter, C. Ritter, and J. Peters, “Deep lagrangian networks: Using physics as model prior for deep learning,” in *International Conference on Learning Representations*, 2019.

[147] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[148] Y. D. Zhong, B. Dey, and A. Chakraborty, "Symplectic ODE-net: Learning Hamiltonian dynamics with control," *arXiv preprint arXiv:1909.12077*, 2019.

[149] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, "Lagrangian neural networks," *arXiv preprint arXiv:2003.04630*, 2020.

[150] C. Neary and U. Topcu, "Compositional learning of dynamical system models using port-Hamiltonian neural networks," in *Proceedings of the Learning for Dynamics and Control Conference*, 2023, pp. 679–691.

[151] E. Dierkes, C. Offen, S. Ober-Blöbaum, and K. Flaßkamp, "Hamiltonian neural networks with automatic symmetry detection," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 33, no. 6, 2023.

[152] F. Fuchs, D. Worrall, V. Fischer, and M. Welling, "SE(3)-transformers: 3D roto-translation equivariant attention networks," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 1970–1981.

[153] F. Djeumou, C. Neary, E. Goubault, S. Putot, and U. Topcu, "Neural networks with physics-informed architectures and constraints for dynamical systems modeling," in *Proceedings of the Learning for Dynamics and Control Conference*, 2022, pp. 263–277.

[154] R. Dai, G. Evangelisti, and S. Hirche, "Physically consistent modeling and identification of nonlinear friction with dissipative gaussian processes," in *Proceedings of the Learning for Dynamics and Control Conference*, 2024, pp. 1415–1426.

[155] W. Cho, M. Jo, H. Lim, K. Lee, D. Lee, S. Hong, and N. Park, "Parameterized physics-informed neural networks for parameterized PDEs," *arXiv preprint arXiv:2408.09446*, 2024.

[156] L. Schulze, J. Peters, and O. Arenz, "Context-aware deep lagrangian networks for model predictive control," *arXiv preprint arXiv:2506.15249*, 2025.

[157] T. Koshizuka and I. Sato, "Understanding generalization in physics-informed models through affine variety dimensions," *arXiv preprint arXiv:2501.18879*, 2025.

[158] J. Drgoňa, T. X. Nghiem, T. Beckers, M. Fazlyab, E. Mallada, C. Jones, D. Vrabie, S. L. Brunton, and R. Findeisen, "Safe physics-informed machine learning for dynamics and control," in *Proceedings of the American Control Conference*, 2025, pp. 591–606.

[159] M. H. Taufik and T. Alkhalifah, "Latentpinns: Generative physics-informed neural networks via a latent representation learning," *Artificial Intelligence in Geosciences*, p. 100115, 2025.

[160] A. H. Chang, C. M. Hubicki, J. J. Aguilar, D. I. Goldman, A. D. Ames, and P. A. Vela, "Learning terrain dynamics: A gaussian process modeling and optimal control adaptation framework applied to robotic jumping," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 4, pp. 1581–1596, 2020.

[161] A. Popov, A. Degirmenci, D. Wehr, S. Hegde, R. Oldja, A. Kamenev, B. Douillard, D. Nistér, U. Muller, R. Bhargava *et al.*, "Mitigating covariate shift in imitation learning for autonomous vehicles using latent space generative world models," *arXiv preprint arXiv:2409.16663*, 2024.

[162] R. Tóth, *Modeling and Identification of Linear Parameter-Varying Systems*. Springer, 2010.

[163] J. Mohammadpour and C. W. Scherer, *Control of Linear Parameter Varying Systems with Applications*. Springer Science & Business Media, 2012.

[164] S. Di Cairano, "Indirect adaptive model predictive control for linear systems with polytopic uncertainty," in *Proceedings of the American Control Conference*, 2016, pp. 3570–3575.

[165] M. M. Morato, J. E. Normey-Rico, and O. Sename, "Model predictive control design for linear parameter varying systems: A survey," *Annual Reviews in Control*, vol. 49, pp. 64–80, 2020.

[166] Z. Mao and I. Ruchkin, "Towards physically interpretable world models: Meaningful weakly supervised representations for visual trajectory prediction," *arXiv preprint arXiv:2412.12870*, 2024.

[167] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, "FiLM: Visual reasoning with a general conditioning layer," in *AAAI Conference on Artificial Intelligence*, 2018.

[168] M. O'Connell, G. Shi, X. Shi, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, "Neural-fly enables rapid learning for agile flight in strong winds," *Science Robotics*, vol. 7, no. 66, p. eabm6597, 2022.

[169] A. Tuor, J. Drgona, and D. Vrabie, "Constrained neural ordinary differential equations with stability guarantees," *arXiv preprint arXiv:2004.10883*, 2020.

[170] D. E. Stewart and J. C. Trinkle, "An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction," *International Journal for Numerical Methods in Engineering*, vol. 39, no. 15, pp. 2673–2691, 1996.

[171] E. Todorov, "Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in mujoco," in *IEEE International Conference on Robotics and Automation*, 2014, pp. 6054–6061.

[172] A. Hochlehnert, A. Terenin, S. Sæmundsson, and M. Deisenroth, "Learning contact dynamics using physically structured neural networks," in *International Conference on Artificial Intelligence and Statistics*, 2021, pp. 2152–2160.

[173] W. Jin, A. Aydinoglu, M. Halm, and M. Posa, "Learning linear complementarity systems," in *Proceedings of the Learning for Dynamics and Control Conference*, 2022, pp. 1137–1149.

[174] J. Urain, M. Ginesi, D. Tateo, and J. Peters, "Imitationflow: Learning deep stable stochastic dynamic systems by normalizing flows," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.

[175] M. Valente, T. C. Dias, V. Guerra, and R. Ventura, “Physics-consistent machine learning with output projection onto physical manifolds,” *Communications Physics*, vol. 8, no. 1, p. 433, 2025.

[176] S. Jeen, T. Bewley, and J. Cullen, “Conservative world models,” in *Neural Information Processing Systems Workshop on Generalization in Planning*, 2023.

[177] I. Mayer, J. Josse, F. Raimundo, and J.-P. Vert, “MissDeepCausal: Causal inference from incomplete data using deep latent variable models,” *arXiv preprint arXiv:2002.10837*, 2020.

[178] Y. Ovadia, E. Fertig, J. Ren, Z. Nado, D. Sculley, S. Nowozin, J. Dillon, B. Lakshminarayanan, and J. Snoek, “Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[179] D. Krueger, E. Caballero, J.-H. Jacobsen, A. Zhang, J. Binas, D. Zhang, R. Le Priol, and A. Courville, “Out-of-distribution generalization via risk extrapolation,” *arXiv preprint arXiv:2003.00688*, 2020.

[180] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.

[181] D. M. Cherenson and D. Panagou, “Adaptive control allocation for underactuated time-scale separated non-affine systems,” *arXiv preprint arXiv:2510.07507*, 2025.

[182] K. P. Wabersich and M. N. Zeilinger, “A predictive safety filter for learning-based control of constrained nonlinear dynamical systems,” *Automatica*, vol. 129, p. 109597, 2021.

[183] E. Hüllermeier and W. Waegeman, “Aleatoric and epistemic uncertainty in machine learning: an introduction to concepts and methods,” *Machine Learning*, no. 110, p. 457–506, 2021.

[184] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2000.

[185] P. R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Society for Industrial and Applied Mathematics, 2015.

[186] A. Mesbah, “Stochastic model predictive control: An overview and perspectives for future research,” *IEEE Control Systems Magazine*, vol. 36, no. 6, pp. 30–44, 2016.

[187] L. P. Swiler, T. L. Paez, and R. L. Mayes, “Epistemic uncertainty quantification tutorial,” in *Proceedings of the 27th International Modal Analysis Conference*, vol. 2, no. 9, 2009.

[188] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *Proceedings of the 33rd International Conference on Machine Learning*, 2016, pp. 1050–1059.

[189] A. Hakobyan and I. Yang, “Distributionally robust risk map for learning-based motion planning and control: A semidefinite programming approach,” *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 718–737, 2023.

[190] J. Doyle, “Analysis of feedback systems with structured uncertainties,” *IEE Proceedings D (Control Theory and Applications)*, vol. 129, no. 6, pp. 242–250, 1982.

[191] H. Blom and Y. Bar-Shalom, “The interacting multiple model algorithm for systems with markovian switching coefficients,” *IEEE Transactions on Automatic Control*, vol. 33, no. 8, pp. 780–783, 1988.

[192] I. Franović, S. Eydam, and D. Eroglu, “Regime switching in coupled nonlinear systems: Sources, prediction, and control—minireview and perspective on the focus issue,” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 34, no. 12, 2024.

[193] V. Dhiman, M. J. Khojasteh, M. Franceschetti, and N. Atanasov, “Control barriers in bayesian learning of system dynamics,” *IEEE Transactions on Automatic Control*, vol. 68, no. 1, pp. 214–229, 2023.

[194] S. Samuelson and I. Yang, “Safety-aware optimal control of stochastic systems using conditional value-at-risk,” in *2018 Annual American Control Conference*, 2018, pp. 6285–6290.

[195] A. Tsiamis, D. S. Kalogerias, L. F. O. Chamon, A. Ribeiro, and G. J. Pappas, “Risk-constrained linear-quadratic regulators,” in *2020 59th IEEE Conference on Decision and Control*, 2020, pp. 3040–3047.

[196] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Y. Zou, S. Levine, C. Finn, and T. Ma, “Mopo: Model-based offline policy optimization,” in *Advances in Neural Information Processing Systems*, 2020, pp. 14 129–14 142.

[197] M. Milanese and A. Vicino, “Optimal estimation theory for dynamic systems with set membership uncertainty: An overview,” *Automatica*, vol. 27, no. 6, pp. 997–1009, 1991.

[198] S. Jafarpour, A. Harapanahalli, and S. Coogan, “Interval reachability of nonlinear dynamical systems with neural network controllers,” in *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, 2023, pp. 12–25.

[199] N. Risso, B. Altin, R. G. Sanfelice, and J. Sprinkle, *Set-Valued Model Predictive Control*. Springer International Publishing, 2023.

[200] M. Milanese and C. Novara, “Set membership identification of nonlinear systems,” *Automatica*, vol. 40, no. 6, pp. 957–975, 2004.

[201] V. Vovk, A. Gammerman, and G. Shafer, *Algorithmic Learning in a Random World*. Springer, 2005.

[202] G. Shafer and V. Vovk, “A tutorial on conformal prediction,” *Journal of Machine Learning Research*, vol. 9, pp. 371–421, 2008.

[203] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, “Safe planning in dynamic environments using conformal prediction,” *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 5116–5123, 2023.

[204] K. Y. Chee, T. C. Silva, M. A. Hsieh, and G. J. Pappas, “Uncertainty quantification and robustification of model-based controllers using conformal prediction,” in *Proceedings of the 6th Annual Learning for Dynamics amp; Control Conference*, vol. 242, 2024, pp. 528–540.

[205] A. Aboudonia and J. Lygeros, "Adaptive learning-based model predictive control for uncertain interconnected systems: A set membership identification approach," *Automatica*, vol. 171, p. 111943, 2025.

[206] R. Mehra, "On the identification of variances and adaptive kalman filtering," *IEEE Transactions on Automatic Control*, vol. 15, no. 2, pp. 175–184, 1970.

[207] A. S. Willsky, "A survey of design methods for failure detection in dynamic systems," *Automatica*, vol. 12, no. 6, pp. 601–611, 1976.

[208] M. Basseville and I. V. Nikiforov, *Detection of Abrupt Changes: Theory and Application*. Prentice-Hall, Inc., 1993.

[209] N. Kantas, A. Doucet, S. Singh, and J. Maciejowski, "An overview of sequential monte carlo methods for parameter estimation in general state-space models," *IFAC Proceedings Volumes*, vol. 42, no. 10, pp. 774–785, 2009.

[210] S. Sæmundsson, K. Hofmann, and M. Deisenroth, "Meta reinforcement learning with latent variable gaussian processes," in *34th Conference on Uncertainty in Artificial Intelligence 2018, UAI 2018*, vol. 34, 2018, pp. 642–652.

[211] M. Lauricella and L. Fagiano, "Set membership identification of linear systems with guaranteed simulation accuracy," *IEEE Transactions on Automatic Control*, vol. 65, no. 12, pp. 5189–5204, 2020.

[212] Y. Romano, E. Patterson, and E. J. Candès, "Conformalized quantile regression," in *Advances in Neural Information Processing Systems*, 2019.

[213] R. J. Tibshirani, R. Foygel Barber, E. Candes, and A. Ramdas, "Conformal prediction under covariate shift," in *Advances in Neural Information Processing Systems*, 2019.

[214] I. Gibbs and E. J. Candès, "Adaptive conformal inference under distribution shift," in *Advances in Neural Information Processing Systems*, 2021.

[215] R. Foygel Barber, E. J. Candes, A. Ramdas, and R. J. Tibshirani, "Conformal prediction beyond exchangeability," *The Annals of Statistics*, vol. 51, no. 2, 2022.

[216] A. Dixit, L. Lindemann, S. X. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, "Adaptive conformal prediction for motion planning among dynamic agents," in *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, vol. 211, 2023, pp. 300–314.

[217] I. Gibbs and E. J. Candès, "Conformal inference for online prediction with arbitrary distribution shifts," *Journal of Machine Learning Research*, vol. 25, no. 162, pp. 1–36, 2024.

[218] V. Lin, K. J. Jang, S. Dutta, M. Caprio, O. Sokolsky, and I. Lee, "Dc4l: Distribution shift recovery via data-driven control for deep learning models," in *6th Annual Learning for Dynamics & Control Conference*, 2024, pp. 1526–1538.

[219] A. Farid, S. Veer, D. Pachisia, and A. Majumdar, "Task-driven detection of distribution shifts with statistical guarantees for robot learning," *IEEE Transactions on Robotics*, vol. 41, pp. 926–945, 2024.

[220] A. Ben-Tal and A. Nemirovski, "Robust convex optimization," *Mathematics of Operations Research*, vol. 23, no. 4, pp. 769–805, 1998.

[221] W. Langson, I. Chryssochoos, S. Raković, and D. Mayne, "Robust model predictive control using tubes," *Automatica*, vol. 40, no. 1, pp. 125–133, 2004.

[222] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.

[223] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[224] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.

[225] M. P. Chapman, R. Bonalli, K. M. Smith, I. Yang, M. Pavone, and C. J. Tomlin, "Risk-sensitive safety analysis using conditional value-at-risk," *IEEE Transactions on Automatic Control*, vol. 67, no. 12, pp. 6521–6536, 2021.

[226] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[227] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, "Risk-sensitive and robust decision-making: a CVaR optimization approach," in *Advances in Neural Information Processing Systems*, vol. 28, 2015.

[228] C. W. Miller and I. Yang, "Optimal control of conditional value-at-risk in continuous time," *SIAM Journal on Control and Optimization*, vol. 55, no. 2, pp. 856–884, 2017.

[229] A. Hakobyan, G. C. Kim, and I. Yang, "Risk-aware motion planning and control using CVaR-constrained optimization," *IEEE Robotics and Automation letters*, vol. 4, no. 4, pp. 3924–3931, 2019.

[230] M. Ahmadi, X. Xiong, and A. D. Ames, "Risk-averse control via CVaR barrier functions: Application to bipedal robot locomotion," *IEEE Control Systems Letters*, vol. 6, pp. 878–883, 2021.

[231] P. Akella, A. Dixit, M. Ahmadi, L. Lindemann, M. P. Chapman, G. J. Pappas, A. D. Ames, and J. W. Burdick, "Risk-aware robotics: Tail risk measures in planning, control, and verification [focus on education]," *IEEE Control Systems*, vol. 45, no. 4, pp. 46–78, 2025.

[232] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafieezadeh-Abadeh, "Wasserstein distributionally robust optimization: Theory and applications in machine learning," in *Operations Research & Management Science in the Age of Analytics*. INFORMS, 2019, pp. 130–166.

[233] B. P. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 430–442, 2015.

[234] I. Yang, "Wasserstein distributionally robust stochastic control: A data-driven approach," *IEEE Transactions on Automatic Control*, vol. 66, no. 8, pp. 3863–3870, 2021.

[235] M. Schuurmans and P. Patrinos, "A general framework for learning-based distributionally robust MPC of Markov jump systems," *IEEE Transactions on Automatic Control*, 2023.

[236] K. Kim and I. Yang, "Distributional robustness in minimax linear quadratic control with Wasserstein distance," *SIAM Journal on Control and Optimization*, vol. 61, no. 2, pp. 458–483, 2023.

[237] B. Li, T. Guan, L. Dai, and G.-R. Duan, "Distributionally robust model predictive control with output feedback," *IEEE Transactions on Automatic Control*, vol. 69, no. 5, pp. 3270–3277, 2024.

[238] A. Hakobyan and I. Yang, "Wasserstein distributionally robust control of partially observable linear stochastic systems," *IEEE Transactions on Automatic Control*, vol. 69, no. 9, pp. 6121–6136, 2024.

[239] R. D. McAllister and P. M. Esfahani, "Distributionally robust model predictive control: Closed-loop guarantees and scalable algorithms," *IEEE Transactions on Automatic Control*, vol. 70, no. 5, pp. 2963–2978, 2025.

[240] J.-S. Brouillon, A. Martin, J. Lygeros, F. Dörfler, and G. F. Trecate, "Distributionally robust infinite-horizon control: from a pool of samples to the design of dependable controllers," *IEEE Transactions on Automatic Control*, vol. 70, no. 10, pp. 6465–6480, 2025.

[241] I. Yang, "A dynamic game approach to distributionally robust safety specifications for stochastic systems," *Automatica*, vol. 94, pp. 94–101, 2018.

[242] A. Hakobyan and I. Yang, "Wasserstein distributionally robust motion control for collision avoidance using conditional value-at-risk," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 939–957, 2022.

[243] J. Queeney and M. Benosman, "Risk-averse model uncertainty for distributionally robust safe reinforcement learning," in *Advances in Neural Information Processing Systems*, vol. 36, 2023, pp. 1659–1680.

[244] A. Hakobyan and I. Yang, "Distributionally robust optimization with unscented transform for learning-based motion control in dynamic environments," in *2023 IEEE International Conference on Robotics and Automation*, 2023, pp. 3225–3232.

[245] J. M. Nadales, A. Hakobyan, D. M. De La Pena, D. Limon, and I. Yang, "Risk-aware wasserstein distributionally robust control of vessels in natural waterways," *IEEE Transactions on Control Systems Technology*, vol. 32, no. 4, pp. 1471–1478, 2024.

[246] S. Safaoui and T. H. Summers, "Distributionally robust cvar-based safety filtering for motion planning in uncertain environments," in *IEEE International Conference on Robotics and Automation*, 2024, pp. 103–109.

[247] M. C. Campi, S. Garatti, and M. Prandini, "The scenario approach for systems and control design," *Annual Reviews in Control*, vol. 33, no. 2, pp. 149–157, 2009.

[248] M. C. Campi, S. Garatti, and F. A. Ramponi, "A general scenario theory for nonconvex optimization and decision making," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4067–4078, 2018.

[249] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Transactions on Automatic Control*, vol. 51, no. 5, pp. 742–753, 2006.

[250] G. C. Calafiore and L. Fagiano, "Robust model predictive control via scenario optimization," *IEEE Transactions on Automatic Control*, vol. 58, no. 1, pp. 219–224, 2012.

[251] L. Romao, K. Margellos, and A. Papachristodoulou, "Tight generalization guarantees for the sampling and discarding approach to scenario optimization," in *IEEE Conference on Decision and Control*, 2020, pp. 2228–2233.

[252] J. P. Hespanha, D. Liberzon, and A. S. Morse, "Overcoming the limitations of adaptive control by means of logic-based switching," *Systems & Control Letters*, vol. 49, no. 1, pp. 49–65, 2003.

[253] Y. Zhang and J. Jiang, "Bibliographical review on reconfigurable fault-tolerant control systems," *Annual Reviews in Control*, vol. 32, no. 2, pp. 229–252, 2008.

[254] T. M. Moldovan and P. Abbeel, "Safe exploration in markov decision processes," in *Proceedings of the 29th International Conference on Machine Learning*, 2012, pp. 1451–1458.

[255] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based model predictive control for safe exploration," in *IEEE Conference on Decision and Control*, 2018, pp. 6059–6066.

[256] C. Zimmer, M. Meister, and D. Nguyen-Tuong, "Safe active learning for time-series modeling with gaussian processes," vol. 31, 2018.

[257] B. Charpentier, R. Senanayake, M. Kochenderfer, and S. Günnemann, "Disentangling epistemic and aleatoric uncertainty in reinforcement learning." *arXiv preprint arXiv:2206.01558*, 2022.

[258] J. Ferré, A. Rokem, E. A. Buffalo, J. N. Kutz, and A. Fairhall, "Non-stationary dynamic mode decomposition," *IEEE Access*, vol. 11, pp. 117 159–117 176, 2023.

[259] M. Chan, M. Molina, and C. Metzler, "Estimating epistemic and aleatoric uncertainty with a single model," in *Advances in Neural Information Processing Systems*, 2024, pp. 109 845–109 870.

[260] Z. An, Z. Hou, and W. Du, "Disentangling uncertainties by learning compressed data representation," in *Proceedings of the Learning for Dynamics and Control Conference*, 2025.

[261] M. Lorenzen, M. Cannon, and F. Allgöwer, "Robust mpc with recursive model update," *Automatica*, vol. 103, pp. 461–471, 2019.

[262] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *International Conference on Machine Learning*, 2017, pp. 1321–1330.

[263] M. Zaffran, O. Feron, Y. Goude, J. Josse, and A. Dieuleveut, "Adaptive conformal predictions for time series," in *Proceedings of the 39th International Conference on Machine Learning*, vol. 162, 2022, pp. 25 834–25 866.

[264] A. Cervin, "Integrated control and real-time scheduling," Ph.D. dissertation, Department of Automatic Control, Lund Institute of Technology, 2003.

[265] T. Kurutach, I. Clavera, Y. Duan, A. Tamar, and P. Abbeel, "Model-ensemble trust-region policy optimization," in *International Conference on Learning Representations*, 2018.

[266] C. Xu, Q. Li, J. Luo, and S. Levine, "Rldg: Robotic generalist policy distillation via reinforcement learning," *arXiv preprint arXiv:2412.09858*, 2024.

[267] J. Nilsson, B. Bernhardsson, and B. Wittenmark, "Stochastic analysis and control of real-time systems with random time delays," *Automatica*, vol. 34, no. 1, pp. 57–64, 1998.

[268] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger, "Deep reinforcement learning that matters," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018.

[269] R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare, "Deep reinforcement learning at the edge of the statistical precipice," in *Advances in Neural Information Processing Systems*, vol. 34, 2021, pp. 29 304–29 320.

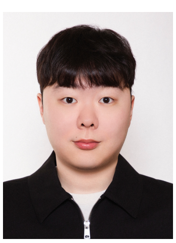

Sung Kwon On received the M.Eng. degree in Engineering Science from the University of Oxford, Oxford, U.K., in 2020. From 2021 to 2024, he was an Assistant Professor with the Department of Electronic and Communication Engineering, Republic of Korea Air Force Academy. In 2024, he was a Visiting Scholar with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. He is currently working toward the Ph.D. degree with the Interdisciplinary Program in Artificial Intelligence, Seoul National University, Seoul, South Korea. His research interests include reinforcement learning and safety control in robotics.

Sukchul Jeong received the B.S. in Mechanical Engineering from Seoul National University, Seoul, South Korea, in 2025. He is currently working toward the Ph.D. degree with Department of Electrical and Computer Engineering, Seoul National University. His research intersets include variational inference, stochasting optimal control, and reinforcement learning.

Jungjin Lee received the B.S. in Mechanical Engineering from Seoul National University, Seoul, South Korea, in 2024. He is currently working toward the Ph.D. degree with Department of Electrical and Computer Engineering, Seoul National University. His research interests include nonlinear dynamical systems, data-driven modeling, and stochastic optimal control.

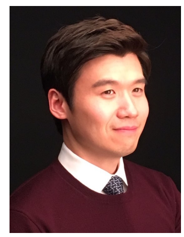

Insoon Yang is currently a Professor in the Department of Electrical and Computer Engineering (ECE) at Seoul National University (SNU). Previously, he was an Assistant Professor of ECE at USC (2016–2018) and a Postdoctoral Associate at MIT (2015–2016). He received B.S. degrees in Mathematics and Mechanical and Aerospace Engineering from SNU (2009). He obtained his M.S. (2012) and Ph.D. (2015) in EECS, as well as an M.A. in Mathematics (2013), from UC Berkeley. His research interests lie at the intersection of control theory, optimization, and machine learning, with a focus on decision-making under uncertainty and physical AI. He received the 2015 Eli Jury Award and was a finalist for the Best Student Paper Award at the 55th IEEE Conference on Decision and Control (CDC).