

A Comprehensive Survey of Cross-Domain Policy Transfer for Embodied Agents

Haoyi Niu¹, Jianming Hu^{1*}, Guyue Zhou^{1*} and Xianyuan Zhan^{1,2*}

¹Tsinghua University

²Shanghai Artificial Intelligence Laboratory

nhy22@mails.tsinghua.edu.cn, zhanxianyuan@air.tsinghua.edu.cn

Abstract

The burgeoning fields of robot learning and embodied AI have triggered an increasing demand for large quantities of data. However, collecting sufficient unbiased data from the target domain remains a challenge due to costly data collection processes and stringent safety requirements. Consequently, researchers often resort to data from easily accessible source domains, such as simulation and laboratory environments, for cost-effective data acquisition and rapid model iteration. Nevertheless, the environments and embodiments of these source domains can be quite different from their target domain counterparts, underscoring the need for effective cross-domain policy transfer approaches. In this paper, we conduct a systematic review of existing cross-domain policy transfer methods. Through a nuanced categorization of domain gaps, we encapsulate the overarching insights and design considerations of each problem setting. We also provide a high-level discussion about the key methodologies used in cross-domain policy transfer problems. Lastly, we summarize the open challenges that lie beyond the capabilities of current paradigms and discuss potential future directions in this field.

1 Introduction

The past few years have seen rapid progress in the fields of robot learning and embodied AI, integrating advances in computer vision, decision-making, and even language processing to build capable embodied agents [Duan *et al.*, 2022; Open X-Embodiment and others, 2023]. This naturally leads to a surge in demand for high-quality and large-scale training data. However, collecting large amounts of data in the **target domain** (where the policy is deployed, i.e., the real world/task environment) at will can be prohibitively costly due to efficiency issues and safety concerns, e.g., in autonomous driving and industrial robot control [Chen *et al.*, 2023; Nguyen-Tuong and Peters, 2011]. Instead, a popular practice is to uti-

lize data from easily accessible **source domains** (e.g., simulation or laboratory environments) that allow safe exploration and cheap data collection. Despite the advances in simulation modeling technologies, high-fidelity simulators still struggle to capture nuanced physical properties, as well as delicate environmental and embodiment details of target domains [Tobin *et al.*, 2017; Peng *et al.*, 2018]. In some settings, although human demonstration videos can be easily recorded in a controllable manner in the target environment, the distinct embodiment from the target robot agents hinders their direct use in policy learning [Yu *et al.*, 2018]. Such intricate environment and embodiment discrepancies, also referred to as *domain gaps*, negatively impact policies trained on source domain data and inevitably lead to their deployment failures in the target domains. The data bottlenecks in real-world tasks and the wide existence of domain gaps naturally stimulated cross-domain policy transfer studies, aiming to fully exploit existing off-domain data to learn transferable policies.

Cross-domain policy transfer has emerged as a crucial class of methods for erasing the influences of domain gaps on policies, serving as a bedrock for large-scale real-world deployment of embodied agents [Chebotar *et al.*, 2019; Bewley *et al.*, 2019]. However, existing approaches in this direction are highly fragmented, primarily due to diverse types of domain gaps, various learning paradigms, as well as distinct data constraints and setting assumptions. Such fragmentation seriously shadows our understanding of the underlying connections and differentiations among existing policy transfer strategies, making it difficult for researchers to gain a holistic view of this field and embark on new research endeavors.

In light of this, we present the first comprehensive review of cross-domain policy transfer methods for embodied agents. We begin by unifying the related notations and definitions in cross-domain settings, based on which we provide clear categorization of different types of domain gaps, with discussion on their distinctions and connections. Then, we consolidate a vast diversity of existing methods, organizing them into four commonly encountered domain gap categories in the literature, i.e., appearance, viewpoint, dynamics, and morphology gaps. Furthermore, we provide overarching methodological insights developed in existing approaches, and discuss open challenges and promising future research directions. A detailed architecture of the survey, including domain gap taxonomy, methodology classification, and fu-

*Corresponding authors. We will keep updating a collection of research papers on this topic in <https://github.com/t6-thu/awesome-cross-domain-policy-transfer-for-embodied-agents>.

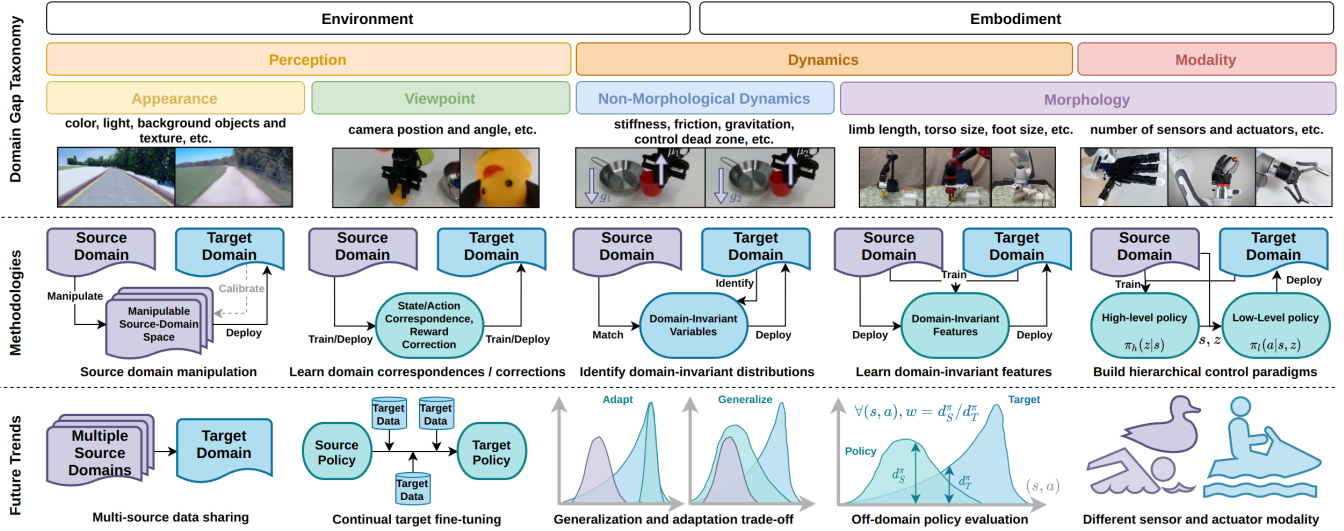


Figure 1: The main architecture of the survey: domain gap taxonomy, overarching insights on methodologies, and future trends.

ture trends are illustrated in Figure 1. We hope our survey can bring new insights and expedite future research in cross-domain policy transfer for embodied agents.

2 Overviews

2.1 Notations and Definitions

Definition 1 (Domain). *With environment \mathcal{E}_{env} , embodiment \mathcal{E}_{emb} , and the embodiment dynamics influenced by both properties $T(\mathcal{E}_{env}, \mathcal{E}_{emb})$, we can define a domain $\Omega = \{\mathcal{E}_{env}, \mathcal{E}_{emb}, T(\mathcal{E}_{env}, \mathcal{E}_{emb})\}$.*

Definition 2 (Domain-Dependent Markov Decision Process (D-MDP)). *Adapting standard MDP with domain-dependent information as $\mathcal{M}(\Omega) := \langle \mathcal{S}_\Omega, \mathcal{A}_\Omega, T_\Omega, R_\Omega, \gamma \rangle$ where $\mathcal{S}_\Omega, \mathcal{A}_\Omega, T_\Omega$ are domain-dependent state, action spaces, and dynamics; R_Ω is the task-relevant reward function also influenced by \mathcal{S}_Ω and \mathcal{A}_Ω ; γ denotes the discount factor.*

To keep the notations uncluttered, we denote $\Omega_{src} = \{\Omega_{src}^i\}_{i=1}^N, N \in \mathbb{N}^+$ as the source domain(s), Ω_{tgt} as the target domain, and use subscript *src* and *tgt* to indicate elements from source and target domains in the rest of the paper. Next, we introduce the definition for cross-domain policy transfer:

Definition 3 (Cross-Domain Policy Transfer). *Given a policy class Π , one or multiple policies $\pi_{src} \in \Pi$ obtained in Ω_{src} , and an explicit or implicit policy transfer method $h: \Pi \rightarrow \Pi$, an effective cross-domain policy transfer is achieved when $J_{tgt}(\pi_{src}) \ll J_{tgt}(h(\pi_{src})) \leq J_{tgt}(\pi_{tgt}^*)$, where $\pi_{tgt}^* \in \Pi$ is the optimal target policy and $J_{tgt}(\cdot)$ is some policy performance evaluation metric based on the criteria (e.g., expected cumulative rewards, success rate, etc.) of target domain Ω_{tgt} .*

2.2 Categorization of Domain Gaps

As indicated in the above definitions, domain gaps arise from both the inconsistencies in environments [Tobin *et al.*, 2017; Peng *et al.*, 2018; Chebotar *et al.*, 2019] and embodiments [Gupta *et al.*, 2021b]. In the following, we establish

fine-grained classifications of domain gaps and discuss their connections as well as distinctions.

In terms of environment inconsistencies, **appearance gaps** arise when observations in the source domain (e.g., simulations) exhibit differences in colors, background objects, illumination conditions, and rendering textures as compared to the target domain (e.g., reality), such as variations in coarse and fine rendering or high and low resolutions [Tobin *et al.*, 2017; Andrychowicz *et al.*, 2020]. Additionally, the configuration of sensor setups (e.g., camera position and angles, etc.) can significantly influence the downstream policy learning of embodied agents, we refer to these as **viewpoint gaps** [Sermanet *et al.*, 2018]. Appearance and viewpoint gaps are sometimes jointly termed **visual gaps** that systematically distort the state space \mathcal{S}_{src} in the source domain, or more generally, **perception gaps** [Wang *et al.*, 2022a] in cases where observations come from other perception sensors besides visual cameras, such as lidars.

At the intersection of embodiment and environment variations, **dynamics gaps** [Peng *et al.*, 2018; Eysenbach *et al.*, 2020; Niu *et al.*, 2022] occur when interactions between embodiments and their deploying environments, or interactions among different parts of the embodiment itself, follow different transitional dynamics ($T_{src} \neq T_{tgt}$), such as stiffness, gear dead zones of embodiments, body mass, and friction. Focusing on the embodiment aspect, **morphology gaps** [Hejna *et al.*, 2020b; Gupta *et al.*, 2021a] arise when target embodiments exhibit different morphological designs compared to the source domain agents, e.g., variations in joint types, module shapes, and lengths, which may ultimately lead to a dynamics mismatch. Morphological differences may also encompass variations in the dimensions and semantic meanings of state and action spaces $\mathcal{S}_{src}, \mathcal{A}_{src}$, such as the number of observational sensors, limbs, and end effectors. In some literature, these are referred to as **modality gaps** [Wang *et al.*, 2022b; Salhotra *et al.*, 2023] for different sensing and actuation modalities. Occasionally, researchers merge the gaps

resulting from morphology discrepancies and parts of non-morphological dynamics disparities, which stem from internal physical properties, into a unified term: *embodiment gaps*. This term represents a more ego-centric robotic perspective, independent of external environmental variations, and should be clearly distinguished from morphology gaps.

3 Policy Transfer Across Different Gaps

3.1 Cross-Appearance Policy Transfer

To remedy the appearance gap, a class of unsupervised transfer learning techniques originating from visual domain adaptation has been proposed. These techniques map observational representations from the source domain (e.g., simulations) to the target domain (real world) while ensuring consistent data distributions. This growing collection of approaches is particularly adept at image-to-image translation across domains, making the observation information transferable for end-to-end vision-based autonomous robots and vehicles [Triess *et al.*, 2021].

Various schemes for the mapping function have been proposed, addressing the problem from different perspectives. Some studies adopt the cycle consistency philosophy from CycleGAN [Zhu *et al.*, 2017] to ensure a photo-realistic image translation process, achieving good real-world transfer performance with only a modest number of real-world observations [Rao *et al.*, 2020]. Conversely, real-to-sim translation [Zhang *et al.*, 2019] requires a pre-trained adaptation module to convert real images captured by cameras into simulation-like images at test time, which can be computationally inefficient during real-world deployment. As an alternative, training a canonical domain-invariant representation, such as semantic segmentation [Pan *et al.*, 2017; Mueller *et al.*, 2018; Wang *et al.*, 2022a], enables observations from both domains to be translated into an intermediate and lower-dimensional representation. This unifies the semantic meaning of the observation space while easing the burden of the downstream policy learning module. In contrast to explicit representations like semantic segmentation, intra-domain image reconstruction with direct and cyclic losses [Bewley *et al.*, 2019] offers another way to enhance transferability, where a bi-directional image translation strategy is introduced to form an implicit structure of representation.

Domain randomization and visual data augmentation [Tobin *et al.*, 2017; Laskin *et al.*, 2020; Kar *et al.*, 2019; Yue *et al.*, 2019] instead opt for a domain generalization approach that focuses on manipulating pixel-level physical mechanisms. These methods do not require massive data from the target domain or learning transferable embeddings. Additionally, proper model setups can also enhance the policy feasibility for real-world execution. For instance, interactive imitation learning [Lee *et al.*, 2022] is specially tailored in simulation to distill state-based experts into a “student” vision-based policy, allowing for in-domain data augmentation from randomized simulations.

3.2 Cross-Viewpoint Policy Transfer

In many cases, training data may not always be available from a first-person or ego-centric viewpoint, which is often the most convenient and desirable observational input [Pathak *et al.*, 2018]. Embodied agents often have different camera setup positions and angles, resulting in observational information with systematic bias. Policy learning that relies on robust cross-appearance visual encoders can still be vulnerable to viewpoint discrepancies, such as a wrist camera on source-domain demonstrators and a side camera in the target environment. To address this, agents need to translate (imagine) third-person observations from their own viewpoint.

To relax the assumption in “learning from demonstrations” that demonstrations come solely from an identical observational configuration, third-person imitation learning [Stadie *et al.*, 2016] constructs an architecture based on generative adversarial imitation learning, which minimizes class loss (expert vs. non-expert) while maximizing domain confusion. Minimizing class loss enables the model to accurately predict the correct class label for a given input, which is crucial for task completion; maximizing domain confusion allows the model to generalize better across different domains, adapting to real-world situations where internet-scale data comes from third-person demonstrators with different viewpoints from embodied agents designated for later deployment.

In a different vein, [Liu *et al.*, 2018] learns a context translation model that can convert a demonstration from one context (e.g., a third-person viewpoint and a human demonstrator) to another context (e.g., a first-person viewpoint and a robot). This approach directly predicts demonstrator behavior sequences from the target robot’s viewpoint, which is claimed to excel in more complex manipulation skill acquisition. Another line of work [Sadeghi *et al.*, 2018] suggests that training a deep convolutional recurrent neural network implicitly learns to identify the effects of actions in image space from the past history of observations and actions. This enables robots to understand how actions affect their motion from the current viewpoint, given a small number of labeled target image queries. Contrastive learning [Sermanet *et al.*, 2018] is also employed to discover attributes that remain consistent across viewpoints or even change throughout task progress.

However, in situations where source domain demonstrators have not only different viewpoints but also distinct morphological embodiments, learning domain-invariant features alone may not suffice for transferrable agent learning. To address these challenges, meta-learning approaches [Yu *et al.*, 2018] have been introduced, although models for each unseen task must be trained separately, necessitating more data and high-capacity models for generalization. Alternatively, a hierarchical setup has also been proposed [Sharma *et al.*, 2019], in which an embodiment-agnostic high-level module learns to generate first-person sub-goals conditioned on third-person demonstrations and an embodiment-specific low-level controller predicts actions to achieve those sub-goals.

3.3 Cross-Dynamics Policy Transfer

Embodied tasks, regardless of whether they have visual observation or not, involve complex transition dynamics that dictate interactions with the environment and constraints

within the embodiment’s components. These intricacies pose considerable challenges in building high-fidelity simulators or finding unbiased source domains. Traditional system identification methods [Ljung, 1998; Kolev and Todorov, 2015; Yu *et al.*, 2017] address domain inconsistency through dynamics model fitting and calibration, which often involve extensive target domain data collection and perform poorly in complex physical dynamics.

An alternative approach is to modify the source domain configuration directly, assuming access to manipulable source domains. Many works have attempted to randomize the physical parameter space of the source domain simulators with different configurations to improve generalization across various target domains [Rajeswaran *et al.*, 2017; Peng *et al.*, 2018; Andrychowicz *et al.*, 2020]. However, this approach often requires manually specified, sufficiently large parameter spaces for adjustment [Vuong *et al.*, 2019], which can be impractical for complex embodied systems. Active domain randomization [Mehta *et al.*, 2020] addresses the sample complexity issue by adaptively selecting parameters from the most informative configurations according to the discrepancies of policy rollouts in randomized and reference environment instances. Another perspective is to revisit classical system identification and lower its demand for target data, such as incorporating target-domain prior information to guide accurate and efficient source-domain posterior distribution calibration [Muratore *et al.*, 2021; Ramos *et al.*, 2019; Tan *et al.*, 2016; Du *et al.*, 2021]. Grounded action transformation (GAT) [Hanna and Stone, 2017] learns target-domain forward dynamics models and adjusts source-domain inverse dynamics accordingly, modifying source dynamics to better match target dynamics.

When source domains are not white-box or modifiable, many recent dynamics adaptation approaches focus on regularizing *policy learning* rather than *dynamics modeling*, assuming a fixed source domain. GARAT [Desai *et al.*, 2020] learns an adversarial imitation-from-observation policy by discriminating between generated actions and target environment actions, bypassing the need for a modifiable source domain. DARC and related methods [Eysenbach *et al.*, 2020; Liu *et al.*, 2022] solve cross-dynamics reinforcement learning (RL) via reward correction to compensate for dynamics shifts across domains in online or offline settings. H2O and H2O+ [Niu *et al.*, 2022; Niu *et al.*, 2023] introduce a dynamics-aware hybrid offline-and-online RL paradigm, integrating learning from online simulation and offline real-world data in a single-stage learning process while correcting dynamics gaps during policy learning. VGDF [Xu *et al.*, 2023a] samples source domain transitions $(s_{src}, a_{src}, s'_{src})$ with small value difference between s'_{src} and s'_{tgt} (obtained from a learned target domain dynamics model), and combines the selected source domain data and target domain counterparts for policy learning.

To address dynamics gaps more affordably, some approaches harness task-relevant, domain-agnostic information in state transitions. SAIL [Liu *et al.*, 2020] advocates for state alignment in cross-domain imitation learning (IL), as optimal policies heuristically induce similar state trajectories under different imitator and expert dynamics. SAIL en-

forces global state distribution matching based on Wasserstein distance and local state transition alignment based on β -VAE. Concurrent work [Gangwani and Peng, 2020] leverages the Wasserstein distance of state visitation distributions from both domains and an adversarial IL paradigm for policy optimization. Additionally, incorporating an inverse dynamics policy learned with target demonstrations [Christiano *et al.*, 2016] offers an alternative for matching (next-)state distributions. HIDIL [Jiang *et al.*, 2020] extends this idea with Horizon-Adaptive Inverse Dynamics, matching states from both domains in an H-step horizon and recovering feasible actions in the target domain based on the inverse dynamics policy. SOIL [Radosavovic *et al.*, 2021] and SRPO [Xue *et al.*, 2023] further develop these ideas, with the latter extending this insight to RL and providing theoretical grounding for the assumption that optimal policies under different dynamics induce similar stationary state visitation distributions.

However, the assumption of identical state reachability in source and target domains does not always hold in real-world situations. Feasibility MDP (f-MDP) [Cao *et al.*, 2021; Cao and Sadigh, 2021] addresses this issue by calculating feasibility scores to weigh the learning signal of each demonstration. Cold diffusion techniques can also be adapted for feasibility-guided trajectory planning by degrading every state in source trajectories to the nearest recorded state in the target replay buffer [Wang *et al.*, 2023].

3.4 Cross-Morphology Policy Transfer

Morphology gaps typically only affect low-level control, which naturally favors a decoupled hierarchical solution with a morphology-specific low-level policy and a transferable high-level policy [Hejna *et al.*, 2020b]. The high-level policy takes morphology-agnostic state observations and generates sub-goals for the low-level policy to follow. MAIL [Salhotra *et al.*, 2023] leverages domain-invariant features in the observation space, such as end effector positions, as optimal position trajectories should be task-relevant and morphology-independent. It performs position-based matching at the high level and uses inverse dynamics to recover morphology-specific low-level action commands. Alternatively, TAME [Hejna *et al.*, 2020a] explores joint optimization for the best morphology design that benefits the embodiment for successful task execution, and optimal policies corresponding to this morphology.

Morphology can also be considered as another modality that can be conditioned on models with the Transformer architecture [Vaswani *et al.*, 2017]. Morphology-aware Transformer [Yu *et al.*, 2023] captures meaningful patterns between robot embodiment and actions using a causally masked Transformer, allowing conditional action generation based on desired robot embodiment, past states, and past actions. However, learning a universal controller for a population of morphologies is resource-intensive and infeasible due to the exponentially increasing morphology representation space [Gupta *et al.*, 2021b]. Focusing on learning embodiment-dependent policies, MetaMorph [Gupta *et al.*, 2021a] first encodes morphology representation into a vector sequence, concatenates it with proprioceptive position embeddings, and processes it using a morphology-aware Trans-

former. The Transformer output is successively concatenated with exteroceptive observations before passing through the final action decoder. Large-scale pre-training over libraries of different morphologies is also utilized to facilitate sample-efficient transfer to new robot morphologies and tasks.

3.5 Cross-Multi-Gap Policy Transfer

In many complex tasks, we might simultaneously encounter multiple types of domain gaps due to substantially different embodiments and deployed environments. However, most previous works focus on paired and temporal-aligned data from source and target domains, and only address a certain type of domain gaps, which limits their applicability in general settings unless extra designs are introduced.

From a general perspective, correspondence learning across domains can empirically align the properties of both source and target domains by constructing direct mappings. GAMA [Kim *et al.*, 2020] learns state and action correspondence mapping $f : \mathcal{S}_{src} \rightarrow \mathcal{S}_{tgt}$, $g : \mathcal{A}_{tgt} \rightarrow \mathcal{A}_{src}$ from unpaired and unaligned demonstrations, and then adapts source policy π_{src} to feasible target policy as $\pi_{tgt} = g \circ \pi_{src} \circ f$. To jointly optimize state and action correspondence models, adopting dynamics cycle consistency [Zhang *et al.*, 2020] allows for splitting action mapping g into coupled dual mappings $p : \mathcal{S}_{src} \times \mathcal{A}_{src} \rightarrow \mathcal{A}_{tgt}$ and $q : \mathcal{S}_{tgt} \times \mathcal{A}_{tgt} \rightarrow \mathcal{A}_{src}$, which builds correlations between action and state correspondences. For various scenarios that only involve suboptimal policies as demonstrators, such as internet videos of humans performing tasks, [Raychaudhuri *et al.*, 2021] enforces cycle consistency on the state space together with a normalized position estimator function to align trajectories across domains without the need for expert actions. To handle the remaining domain misalignment issue of adopting unsupervised cycle consistency techniques, WeaSCL [Wang *et al.*, 2022b] introduces weak supervision into correspondence learning with temporal ordering and paired abstraction data.

Instead of learning direct correspondences, an emerging avenue of studies has extended learning domain-invariant features from closing appearance and viewpoint gaps to simultaneously addressing other domain gaps, e.g., dynamics and morphology gaps. From internet-scale cross-embodiment videos, XIRL [Zakka *et al.*, 2021] leverages temporal cycle consistency to ensure task-progress aware and domain-agnostic representation learning so that distance from goal state represented in that embedding space can be regarded as rewards used for policy training on novel embodiments. Motivated by abstracting task information from state space to ease the burden of downstream policy transfer, [Franzmeier *et al.*, 2022] proposes a mutual information criterion to reduce target state space with mapping $f' : \mathcal{S}_{tgt} \rightarrow \mathcal{Z}$ to a task-relevant domain-invariant embedding \mathcal{Z} , and then jointly learning source mapping $g' : \mathcal{S}_{src} \rightarrow \mathcal{Z}$ and an adversarial imitation policy that generates state transitions closely resembling the expert target state transitions. From a more robotic perspective, skill acquisition is an explicit procedure for grounding task-specific domain-agnostic features for easy transfer. With paired data from both domains, agents learn multiple skills and transfer knowledge by training in invariant feature spaces, upon which target domain agents can ac-

quire new skills mastered by source domain agents [Gupta *et al.*, 2017]. In a hierarchical scheme, STAR [Pertsch *et al.*, 2022] pre-trains a low-level policy to decode actions from learned high-level semantic skill policies that select a transferable skill in target task learning. With unpaired and unaligned cross-embodiment videos, XSkill [Xu *et al.*, 2023b] pre-trains skill discovery models for further skill identification. In XSkill, a skill alignment transformer is introduced to detect, align, and compose the learned skills to complete new tasks, and then pass the inferred skills to a skill-conditioned diffusion policy to output the robot’s actions.

Additionally, contrastive learning also offers a general and natural solution for aligning domain representation with positive and negative samples, from large amounts of in-the-wild cross-embodiment unpaired data. Polybot [Yang *et al.*, 2023] aligns observation and action spaces using an engineering approach and then aligns policy’s internal representations through contrastive learning to combat other domain discrepancies. Based on prompt-based learning, CONPE [Choi *et al.*, 2023] develops a novel contrastive prompt ensemble framework that uses the CLIP vision-language model [Radford *et al.*, 2021] as the visual encoder and facilitates dynamic adjustments of visual representations against domain changes through an ensemble of contrastively learned visual prompts. VIP [Ma *et al.*, 2022] contrastively pre-trains an (implicit) visual goal-conditioned value function that aims to capture task-agnostic goal-oriented representations, which can generalize to unseen domains and tasks. As a multi-modal extension of VIP, LIV [Ma *et al.*, 2023] learns vision-language representations from language-annotated videos. In a similar setting, DecisionNCE [Li *et al.*, 2024] learns universal embodied multimodal representations through an infoNCE-style learning objective, derived based on reward reparameterization under the preference-based learning framework. RT-X [Open X-Embodiment and others, 2023] harnesses domain alignment and scene understanding ability of large vision-language models, unifying domain representations with domain-invariant task-relevant language instructions.

4 Discussions on Methodologies

4.1 Source Domain Manipulation

Modifying the source domain modeling is undoubtedly the most straightforward solution to close the domain gaps when source domains (e.g., simulators) are manipulable. Such modifications often include: randomizing partially known or modifiable modeling configurations for better generalization in target domains [Tobin *et al.*, 2017; Rajeswaran *et al.*, 2017; Peng *et al.*, 2018; Lee *et al.*, 2022; Mehta *et al.*, 2020; Andrychowicz *et al.*, 2020], model calibration to match better with the target domains [Kolev and Todorov, 2015; Tan *et al.*, 2016; Yu *et al.*, 2017; Hanna and Stone, 2017; Ramos *et al.*, 2019; Muratore *et al.*, 2021; Du *et al.*, 2021], and adapt morphological design configurations from libraries of candidates [Hejna *et al.*, 2020a; Gupta *et al.*, 2021b; Gupta *et al.*, 2021a; Yu *et al.*, 2023]. Large-scale source-domain randomization has been widely applied to visual properties (i.e., color, texture, lighting condition, shapes and types of interactive objects, camera position, orientation, and

field of view) [Tobin *et al.*, 2017], dynamics parameters (i.e., mass, damping, friction, and control timestep) [Peng *et al.*, 2018], or oftentimes both [Lee *et al.*, 2022]. With principled morphology representation, we could effectively manipulate action space, sensory inputs, module shape, and size so that agents can generalize in the vast modular morphology design space. This promotes domain generalization ability maximally with large-scale parallel computing resources to support high-dimensional parameter randomization in simulation, which is highly recommended in industrial practice [Lee *et al.*, 2022].

In cases when access to target domain data is allowed but suffers costly and laborious collection, source-domain calibration turns out to be fundamental, direct, and effective for simple and well-modeled environments [Yu *et al.*, 2017]. However, target domain environments can be quite complex in practice, even strong calibration approaches are likely to under-model intricate target domain physical properties and procedures, such as non-rigidity, gear dead zone, wear-and-tear, and rolling friction. Such properties are hard to capture in current simulation modeling technologies. To summarize, source domain manipulation is highly dependent on manipulable source domains, nuanced and comprehensive knowledge of environment modeling, and rich computation resources, as preferred by the industry field.

4.2 Learn Domain Correspondences / Corrections

Learning mapping functions (correspondences) or correction terms is another class of commonly used techniques to handle domain gaps. Image-to-image translation [Zhu *et al.*, 2017; Pan *et al.*, 2017; Zhang *et al.*, 2019; Rao *et al.*, 2020] build mappings between visual representations across domains. The viewpoint context translation model [Liu *et al.*, 2018] converts data from the source context to ones from the target viewpoint. Unlike previous works, learning state and action correspondences for domain alignment [Kim *et al.*, 2020; Zhang *et al.*, 2020] allows for direct leverage of unpaired data. Additional designs can also be added in this general framework to reduce the need for expert action collection [Raychaudhuri *et al.*, 2021], which fits well with the common setting of learning from action-free video demonstrations. To address the accuracy issue of correspondence learning under stricter conditions, WeaSCL [Wang *et al.*, 2022b] finds a trade-off between strong supervision of strictly paired data and regularization over unpaired data. In addition to learning correspondences, learning reward correction for domain gap compensation [Eysenbach *et al.*, 2020; Liu *et al.*, 2022; Xue *et al.*, 2023] and dynamics ratio for reweighting learning signals on source domain samples [Niu *et al.*, 2022; Niu *et al.*, 2023] have also become viable practices in cross-domain RL.

4.3 Identify Domain-Invariant Distributions

Identifying domain-invariant distributions from accessible and organized data offers a seemingly simpler solution for cross-domain policy transfer without extracting detailed domain-dependent correspondences. For example, state regularization in IL [Christiano *et al.*, 2016; Liu *et al.*, 2020; Gangwani and Peng, 2020; Jiang *et al.*, 2020; Radosavovic *et*

al., 2021] and RL [Xue *et al.*, 2023] develop upon the insight that optimal policies across different domains induce similar state visitation distribution. Cold diffusion is used in diffusion-based planning to constrain the generated states inside the state distribution of the target replay buffer [Wang *et al.*, 2023]. In some cases, partial state distribution matching (e.g., position-based matching) is adopted since not all dimensions of the state space are semantically identical across domains [Salhotra *et al.*, 2023]. However, this avenue of work typically only addresses dynamics and morphology gaps without discrepancies in visual representation; otherwise, extra cross-domain correspondences/representations are required to align the state space before performing state-based matching. Additionally, extra efforts are needed to apply these works to long-horizon compositional tasks since state distribution matching lacks task-progress awareness.

4.4 Learn Domain-Invariant Features

Learning task-relevant domain-invariant representations is also a principled and popular direction to bridge domain gaps [Stadie *et al.*, 2016; Sermanet *et al.*, 2018; Mueller *et al.*, 2018; Bewley *et al.*, 2019; Zakka *et al.*, 2021; Franzmeyer *et al.*, 2022; Wang *et al.*, 2022a], which sometimes also appear in the form of skills [Gupta *et al.*, 2017; Pertsch *et al.*, 2022; Xu *et al.*, 2023b] and sub-goals [Sharma *et al.*, 2019]. As canonical representation across domains can be reused in multiple and even new target contexts, offering great flexibility and data efficiency, however, these approaches could also suffer some barriers as compared to learning correspondences. The domain-invariant representations, for instance, could require additional efforts for tackling issues like uninformative degenerated mapping [Gupta *et al.*, 2017]. Furthermore, this line of works often solely focuses on learning invariant features in observations, unlike learning state and action correspondences that could seamlessly align different MDPs temporally, which brings better task progress awareness for planning tasks. This highlights the extra need to borrow off-the-shelf temporal vision alignment techniques for pairing demonstrations, e.g. temporal contrastive network [Sermanet *et al.*, 2018] and temporal cycle consistency [Dwibedi *et al.*, 2019]. However, with self-supervision on comprehensive long-horizon multi-skill demonstrations, the learned representations sometimes could also be progression/distance-aware, holding advantages of yielding goal-directed visual reward [Sermanet *et al.*, 2018; Zakka *et al.*, 2021; Ma *et al.*, 2022; Li *et al.*, 2024] and guiding the downstream policy optimization process. It also provides possibilities for incorporating language instructions into unified vision-language cross-embodiment representations [Ma *et al.*, 2023; Li *et al.*, 2024].

4.5 Build Hierarchical Control Paradigms

The essence of hierarchical frameworks in cross-domain settings is to decouple the action output procedure into domain-independent high-level policy learning (e.g., skill acquisition [Gupta *et al.*, 2017; Pertsch *et al.*, 2022; Xu *et al.*, 2023b] and sub-goal generation [Hejna *et al.*, 2020b; Sharma *et al.*, 2019]) and domain-specific low-level policy learning. Such a treatment greatly reduces the difficulties in domain gap

modeling and has proven to excel in solving many complex tasks. In a similar philosophy, some meta-learning methods [Finn *et al.*, 2017; Yu *et al.*, 2018; Nagabandi *et al.*, 2018; Zintgraf *et al.*, 2019; Rakelly *et al.*, 2019] train context-based hidden embeddings for fast adaptation, with which the meta-learned policy can adapt to target environments by fine-tuning on a small amount of target domain data.

5 Open Challenges and Future Trends

5.1 Different Sensor and Actuator Modalities

The transferable knowledge from source domain embodiments can be quite limited when dealing with significantly different state and action modalities [Wang *et al.*, 2022b; Salhotra *et al.*, 2023]. Recently, large and expressive foundation models, combined with large-scale data collected from diverse robotic tasks, have emerged as a promising direction for cross-embodiment policy transfer [Open X-Embodiment and others, 2023]. Octo [Octo Model Team *et al.*, 2023], a transformer-based diffusion policy model, serves as a versatile policy initialization that can be effectively fine-tuned to adapt to new observation and action spaces. Its block-wise attention structure allows for adding or removing new inputs and outputs with various modalities as needed. However, it remains unclear what and how an agent can learn from data from significantly different embodiments of the same task.

Most studies in cross-modality settings heavily rely on expert demonstrations, which are costly to collect and limited in size, causing the issue that expert policies are hard to learn or transfer across modalities. Future work could focus on learning effective and transferable information from non-optimal, in-the-wild demonstrations to address these limitations.

5.2 Multi-Source Data Sharing

Current research on learning correspondences between domains typically focuses on narrow settings, where data are assumed to originate from only two domains [Stadie *et al.*, 2016; Bewley *et al.*, 2019; Niu *et al.*, 2022]. However, in practice, it is crucial to handle multiple source domains to overcome the data scarcity issue [Xue *et al.*, 2023]. Modern cross-domain methods need more flexible interfaces to incorporate multi-source data with domain gaps of varying scales.

A versatile and expressive feature space could also be developed to unify the representation of data across different domains. In addition to focusing on representation learning, another potential direction is to directly filter or edit source data according to learned criteria in the target domain. This perspective emphasizes manipulating data as a means to address the challenges of multi-source data sharing.

5.3 Continual Target Fine-Tuning

Current cross-domain policy transfer approaches often lack flexible designs to accommodate various forms of source-domain information, such as data and pre-trained policy. Sometimes, we might desire to fine-tune the source-domain policy using target data, as target data might not be readily available at the beginning of training and only obtainable sporadically, encompassing different coverages and skillsets.

This necessitates a policy model that is compatible with continual learning. Essentially, it could also help relax the long-standing assumption that the target domain remains time-invariant, while real-world systems often deviate from this due to factors such as wear and tear.

A potential approach to leverage multi-stage target data is the adoption of continual fine-tuning [Li *et al.*, 2023; Smith *et al.*, 2023]. These techniques enable the continuous integration of target data for fine-tuning, thereby facilitating continual skill acquisition and policy adaptation in time-varying target domains. As a result, the development of a generalist, multi-skill, multi-domain policy after fine-tuning becomes more achievable, paving the way for more robust and versatile cross-domain solutions in future research.

5.4 Generalization and Adaptation Trade-Off

Current cross-domain transferable policies tend to be either highly adaptive to a single accessible target domain or moderately generalizable to various random domains. Striking a balance between generalization and adaptation, akin to the trade-off between generalists and specialists in a cross-domain setting, appears to be a challenging task. To find the nuanced equilibrium, employing powerful foundation models and extensive cross-domain data could serve as a practical solution to such demanding requirements [Reed *et al.*, 2022].

Recently, we have witnessed the burgeoning emergence of large (vision-)language models, which can naturally serve as powerful domain aligners since language is fundamental, easily attainable, information-abstract, and domain-transferable [Open X-Embodiment and others, 2023; Ma *et al.*, 2023; Li *et al.*, 2024]. Moreover, language models have the potential to serve as domain generalizers due to their strong common sense reasoning abilities across a wide range of everyday scenarios. However, vision-language representation alignment remains a longstanding challenge, as we cannot expect language models to either generalize to or adapt to desired target domains without resolving the alignment of representations with diverse input modalities.

5.5 Off-Domain Policy Evaluation

Evaluating policies in target domains can sometimes be prohibitively expensive and even hazardous, while continuous access to source domains allows for faster and more controlled policy evaluation, albeit with reduced reliability due to domain gaps. However, there is a scarcity of theoretical or principled criteria to determine whether a policy model evaluated in source domains can be successfully transferred to target domains [Katdare *et al.*, 2023]. Therefore, future research should develop reliable *off-domain policy evaluation* methods (an extension of off-policy evaluation under domain gaps) together with standardized real-world benchmarks [Walke *et al.*, 2023; Open X-Embodiment and others, 2023]. These efforts are expected to offer a principled procedure and rigorous criteria for evaluating the transferability of policies using accessible source domains and limited pre-collected target data, which are the common settings in real-world scenarios.

6 Conclusion

In this survey, we provide the first comprehensive review of the rapidly evolving field of cross-domain policy transfer for embodied agents. We have unified the notations and definitions in cross-domain settings, distinguishing various types of domain gaps and clarifying their connections and differences. By categorizing the highly fragmented approaches in the literature, we shed light on the methods used to address appearance, viewpoint, dynamics, and morphology gaps. Moreover, we provide overarching insights shared among these methodological approaches and discuss open challenges as well as promising future trends. As the field of embodied AI continues to evolve, addressing these challenges and embracing emerging trends will be crucial for developing more robust and versatile solutions for real-world deployment. We hope our survey can serve as a useful tool for future research, offering a clear understanding of the current state of cross-domain policy transfer and providing a roadmap for tackling the remaining challenges in this exciting and rapidly growing field.

Acknowledgments

This work is supported by National Natural Science Foundation of China under Grant No. 62333015 and No. 62133002, National Key Research and Development Program of China under Grant (2022YFB2502904), Beijing Natural Science Foundation L231014, and funding from Haomo.AI.

References

- [Andrychowicz *et al.*, 2020] Marcin Andrychowicz, Bowen Baker, Maciek Chociej, et al. Learning dexterous in-hand manipulation. *IJRR*, 2020.
- [Bewley *et al.*, 2019] Alex Bewley, Jessica Rigley, Yuxuan Liu, et al. Learning to drive from simulation without real world labels. In *ICRA*, 2019.
- [Cao and Sadigh, 2021] Zhangjie Cao and Dorsa Sadigh. Learning from imperfect demonstrations from agents with varying dynamics. *RAL*, 2021.
- [Cao *et al.*, 2021] Zhangjie Cao, Yilun Hao, Mengxi Li, and Dorsa Sadigh. Learning feasibility to imitate demonstrators with different dynamics. In *CoRL*, 2021.
- [Chebotar *et al.*, 2019] Yevgen Chebotar, Ankur Handa, Viktor Makoviychuk, Miles Macklin, Jan Issac, Nathan Ratliff, and Dieter Fox. Closing the sim-to-real loop: Adapting simulation randomization with real world experience. In *ICRA*, 2019.
- [Chen *et al.*, 2023] Li Chen, Penghao Wu, Kashyap Chitta, Bernhard Jaeger, Andreas Geiger, and Hongyang Li. End-to-end autonomous driving: Challenges and frontiers. *arXiv*, 2023.
- [Choi *et al.*, 2023] Wonje Choi, Woo Kyung Kim, SeungHyun Kim, and Honguk Woo. Efficient policy adaptation with contrastive prompt ensemble for embodied agents. In *NeurIPS*, 2023.
- [Christiano *et al.*, 2016] Paul Christiano, Zain Shah, Igor Mordatch, et al. Transfer from simulation to real world through learning deep inverse dynamics model. *arXiv*, 2016.
- [Desai *et al.*, 2020] Siddharth Desai, Ishan Durugkar, Haresh Karnan, et al. An imitation from observation approach to transfer learning with dynamics mismatch. *NeurIPS*, 2020.
- [Du *et al.*, 2021] Yuqing Du, Olivia Watkins, Trevor Darrell, Pieter Abbeel, and Deepak Pathak. Auto-tuned sim-to-real transfer. In *ICRA*, 2021.
- [Duan *et al.*, 2022] Jiafei Duan, Samson Yu, Hui Li Tan, Hongyuan Zhu, and Cheston Tan. A survey of embodied ai: From simulators to research tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2022.
- [Dwibedi *et al.*, 2019] Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal cycle-consistency learning. In *CVPR*, 2019.
- [Eysenbach *et al.*, 2020] Benjamin Eysenbach, Shreyas Chaudhari, Swapnil Asawa, Sergey Levine, and Ruslan Salakhutdinov. Off-dynamics reinforcement learning: Training for transfer with domain classifiers. In *ICLR*, 2020.
- [Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- [Franzmeyer *et al.*, 2022] Tim Franzmeyer, Philip Torr, and João F Henriques. Learn what matters: cross-domain imitation learning with task-relevant embeddings. *NeurIPS*, 2022.
- [Gangwani and Peng, 2020] Tanmay Gangwani and Jian Peng. State-only imitation with transition dynamics mismatch. In *ICLR*, 2020.
- [Gupta *et al.*, 2017] Abhishek Gupta, Coline Devin, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Learning invariant feature spaces to transfer skills with reinforcement learning. In *ICLR*, 2017.
- [Gupta *et al.*, 2021a] Agrim Gupta, Linxi Fan, Surya Ganguli, and Li Fei-Fei. Metamorph: Learning universal controllers with transformers. In *ICLR*, 2021.
- [Gupta *et al.*, 2021b] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *Nature communications*, 2021.
- [Hanna and Stone, 2017] Josiah Hanna and Peter Stone. Grounded action transformation for robot learning in simulation. In *AAAI*, 2017.
- [Hejna *et al.*, 2020a] Joey Hejna, Pieter Abbeel, and Lerrel Pinto. Task-agnostic morphology evolution. In *ICLR*, 2020.
- [Hejna *et al.*, 2020b] Joey Hejna, Lerrel Pinto, and Pieter Abbeel. Hierarchically decoupled imitation for morphological transfer. In *ICML*, 2020.
- [Jiang *et al.*, 2020] Shengyi Jiang, Jingcheng Pang, and Yang Yu. Offline imitation learning with a misspecified simulator. *NeurIPS*, 2020.

- [Kar *et al.*, 2019] Amlan Kar, Aayush Prakash, Ming-Yu Liu, Eric Cameracci, Justin Yuan, Matt Rusiniak, David Acuna, Antonio Torralba, and Sanja Fidler. Meta-sim: Learning to generate synthetic datasets. In *ICCV*, 2019.
- [Katdare *et al.*, 2023] Pulkit Katdare, Nan Jiang, and Katherine Rose Driggs-Campbell. Marginalized importance sampling for off-environment policy evaluation. In *CoRL*, 2023.
- [Kim *et al.*, 2020] Kuno Kim, Yihong Gu, Jiaming Song, Shengjia Zhao, and Stefano Ermon. Domain adaptive imitation learning. In *ICML*, 2020.
- [Kolev and Todorov, 2015] Svetoslav Kolev and Emanuel Todorov. Physically consistent state estimation and system identification for contacts. In *Humanoids*, 2015.
- [Laskin *et al.*, 2020] Misha Laskin, Kimin Lee, Adam Stooke, et al. Reinforcement learning with augmented data. *NeurIPS*, 2020.
- [Lee *et al.*, 2022] Alex X Lee, Coline Manon Devin, Yuxiang Zhou, et al. Beyond pick-and-place: Tackling robotic stacking of diverse shapes. In *CoRL*, 2022.
- [Li *et al.*, 2023] Jianxiong Li, Xiao Hu, Haoran Xu, Jingjing Liu, Xianyuan Zhan, and Ya-Qin Zhang. Proto: Iterative policy regularized offline-to-online reinforcement learning. *arXiv*, 2023.
- [Li *et al.*, 2024] Jianxiong Li, Jinliang Zheng, Yinan Zheng, Liyuan Mao, Xiao Hu, Sijie Cheng, Haoyi Niu, et al. Decisionncc: Embodied multimodal representations via implicit preference learning. In *ICML*, 2024.
- [Liu *et al.*, 2018] YuXuan Liu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *ICRA*, 2018.
- [Liu *et al.*, 2020] Fangchen Liu, Zhan Ling, Tongzhou Mu, and Hao Su. State alignment-based imitation learning. In *ICLR*, 2020.
- [Liu *et al.*, 2022] Jinxin Liu, Zhang Hongyin, and Donglin Wang. Dara: Dynamics-aware reward augmentation in offline reinforcement learning. In *ICLR*, 2022.
- [Ljung, 1998] Lennart Ljung. System identification. In *Signal analysis and prediction*. Springer, 1998.
- [Ma *et al.*, 2022] Yecheng Jason Ma, Shagun Sodhani, Dinesh Jayaraman, Osbert Bastani, Vikash Kumar, and Amy Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *ICLR*, 2022.
- [Ma *et al.*, 2023] Yecheng Jason Ma, Vikash Kumar, Amy Zhang, Osbert Bastani, and Dinesh Jayaraman. Liv: Language-image representations and rewards for robotic control. In *ICML*, 2023.
- [Mehta *et al.*, 2020] Bhairav Mehta, Manfred Diaz, Florian Golemo, Christopher J Pal, and Liam Paull. Active domain randomization. In *CoRL*, 2020.
- [Mueller *et al.*, 2018] Matthias Mueller, Alexey Dosovitskiy, Bernard Ghanem, and Vladlen Koltun. Driving policy transfer via modularity and abstraction. In *CoRL*, 2018.
- [Muratore *et al.*, 2021] Fabio Muratore, Christian Eilers, Michael Gienger, and Jan Peters. Data-efficient domain randomization with bayesian optimization. *RAL*, 2021.
- [Nagabandi *et al.*, 2018] Anusha Nagabandi, Ignasi Clavera, Simin Liu, et al. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. In *ICLR*, 2018.
- [Nguyen-Tuong and Peters, 2011] Duy Nguyen-Tuong and Jan Peters. Model learning for robot control: a survey. *Cognitive processing*, 2011.
- [Niu *et al.*, 2022] Haoyi Niu, Shubham Sharma, Yiwen Qiu, Ming Li, Guyue Zhou, Jianming HU, and Xianyuan Zhan. When to trust your simulator: Dynamics-aware hybrid offline-and-online reinforcement learning. In *NeurIPS*, 2022.
- [Niu *et al.*, 2023] Haoyi Niu, Tianying Ji, Bingqi Liu, Haocheng Zhao, Xiangyu Zhu, Jianying Zheng, Pengfei Huang, Guyue Zhou, Jianming Hu, and Xianyuan Zhan. H2o+: An improved framework for hybrid offline-and-online rl with dynamics gaps. *arXiv*, 2023.
- [Octo Model Team *et al.*, 2023] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, et al. Octo: An open-source generalist robot policy, 2023.
- [Open X-Embodiment and others, 2023] Open X-Embodiment et al. Open x-embodiment: Robotic learning datasets and RT-x models. *arXiv*, 2023.
- [Pan *et al.*, 2017] Xinlei Pan, Yurong You, Ziyang Wang, and Cewu Lu. Virtual to real reinforcement learning for autonomous driving. *BMVC*, 2017.
- [Pathak *et al.*, 2018] Deepak Pathak, Parsa Mahmoudieh, Michael Luo, et al. Zero-shot visual imitation. In *ICLR*, 2018.
- [Peng *et al.*, 2018] Xue Bin Peng, Marcin Andrychowicz, Wojciech Zaremba, and Pieter Abbeel. Sim-to-real transfer of robotic control with dynamics randomization. In *ICRA*, 2018.
- [Pertsch *et al.*, 2022] Karl Pertsch, Ruta Desai, Vikash Kumar, et al. Cross-domain transfer via semantic skill imitation. In *CoRL*, 2022.
- [Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- [Radosavovic *et al.*, 2021] Ilija Radosavovic, Xiaolong Wang, Lerrel Pinto, and Jitendra Malik. State-only imitation learning for dexterous manipulation. In *IROS*, 2021.
- [Rajeswaran *et al.*, 2017] Aravind Rajeswaran, Sarvjeet Ghotra, Balaraman Ravindran, and Sergey Levine. EPOpt: Learning robust neural network policies using model ensembles. In *ICLR*, 2017.
- [Rakelly *et al.*, 2019] Kate Rakelly, Aurick Zhou, Chelsea Finn, Sergey Levine, and Deirdre Quillen. Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *ICML*, 2019.

- [Ramos *et al.*, 2019] Fabio Ramos, Rafael Possas, and Dieter Fox. Bayessim: adaptive domain randomization via probabilistic inference for robotics simulators. In *RSS*, 2019.
- [Rao *et al.*, 2020] Kanishka Rao, Chris Harris, Alex Irpan, Sergey Levine, Julian Ibarz, and Mohi Khansari. Rl-cyclegan: Reinforcement learning aware simulation-to-real. In *CVPR*, 2020.
- [Raychaudhuri *et al.*, 2021] Dripta S Raychaudhuri, Sujoy Paul, Jeroen Vanbaer, and Amit K Roy-Chowdhury. Cross-domain imitation from observations. In *ICML*, 2021.
- [Reed *et al.*, 2022] Scott Reed, Konrad Zolna, Emilio Parisotto, et al. A generalist agent. *TMLR*, 2022.
- [Sadeghi *et al.*, 2018] Fereshteh Sadeghi, Alexander Toshev, Eric Jang, and Sergey Levine. Sim2real viewpoint invariant visual servoing by recurrent control. In *CVPR*, 2018.
- [Salhotra *et al.*, 2023] Gautam Salhotra, I-Chun Arthur Liu, and Gaurav S. Sukhatme. Learning robot manipulation from cross-morphology demonstration. In *CoRL*, 2023.
- [Sermanet *et al.*, 2018] Pierre Sermanet, Corey Lynch, Yevgen Chebotar, et al. Time-contrastive networks: Self-supervised learning from video. *ICRA*, 2018.
- [Sharma *et al.*, 2019] Pratyusha Sharma, Deepak Pathak, and Abhinav Gupta. Third-person visual imitation learning via decoupled hierarchical controller. *NeurIPS*, 2019.
- [Smith *et al.*, 2023] Laura Smith, Yunhao Cao, and Sergey Levine. Grow your limits: Continuous improvement with real-world rl for robotic locomotion. *arXiv*, 2023.
- [Stadie *et al.*, 2016] Bradley C Stadie, Pieter Abbeel, and Ilya Sutskever. Third person imitation learning. In *ICLR*, 2016.
- [Tan *et al.*, 2016] Jie Tan, Zhaoming Xie, Byron Boots, and C Karen Liu. Simulation-based design of dynamic controllers for humanoid balancing. In *IROS*, 2016.
- [Tobin *et al.*, 2017] Josh Tobin, Rachel Fong, Alex Ray, et al. Domain randomization for transferring deep neural networks from simulation to the real world. In *IROS*, 2017.
- [Triess *et al.*, 2021] Larissa T Triess, Mariella Dreissig, Christoph B Rist, and J Marius Zöllner. A survey on deep domain adaptation for lidar perception. In *IV*, 2021.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, et al. Attention is all you need. In *NeurIPS*, 2017.
- [Vuong *et al.*, 2019] Quan Vuong, Sharad Vikram, Hao Su, Sicun Gao, and Henrik I Christensen. How to pick the domain randomization parameters for sim-to-real transfer of reinforcement learning policies? *arXiv*, 2019.
- [Walke *et al.*, 2023] Homer Rich Walke, Kevin Black, Tony Z Zhao, et al. Bridgedata v2: A dataset for robot learning at scale. In *CoRL*, 2023.
- [Wang *et al.*, 2022a] Guan Wang, Haoyi Niu, Desheng Zhu, Jianming Hu, Xianyuan Zhan, and Guyue Zhou. A versatile and efficient reinforcement learning approach for autonomous driving. *NeurIPS ML4AD Workshop*, 2022.
- [Wang *et al.*, 2022b] Zihan Wang, Zhangjie Cao, Yilun Hao, and Dorsa Sadigh. Weakly supervised correspondence learning. In *ICRA*, 2022.
- [Wang *et al.*, 2023] Zidan Wang, Takeru Oba, Takuma Yoneda, Rui Shen, Matthew Walter, and Bradley C Stadie. Cold diffusion on the replay buffer: Learning to plan from known good states. In *CoRL*, 2023.
- [Xu *et al.*, 2023a] Kang Xu, Chenjia Bai, Xiaoteng Ma, Dong Wang, Bin Zhao, Zhen Wang, Xuelong Li, and Wei Li. Cross-domain policy adaptation via value-guided data filtering. In *NeurIPS*, 2023.
- [Xu *et al.*, 2023b] Mengda Xu, Zhenjia Xu, Cheng Chi, Manuela Veloso, and Shuran Song. Xskill: Cross embodiment skill discovery. In *CoRL*, 2023.
- [Xue *et al.*, 2023] Zhenghai Xue, Qingpeng Cai, Shuchang Liu, Dong Zheng, Peng Jiang, Kun Gai, and Bo An. State regularized policy optimization on data with dynamics shift. In *NeurIPS*, 2023.
- [Yang *et al.*, 2023] Jonathan Heewon Yang, Dorsa Sadigh, and Chelsea Finn. Polybot: Training one policy across robots while embracing variability. In *CoRL*, 2023.
- [Yu *et al.*, 2017] Wenhao Yu, Jie Tan, C Karen Liu, and Greg Turk. Preparing for the unknown: Learning a universal policy with online system identification. *arXiv*, 2017.
- [Yu *et al.*, 2018] Tianhe Yu, Chelsea Finn, Sudeep Dasari, et al. One-shot imitation from observing humans via domain-adaptive meta-learning. In *RSS*, 2018.
- [Yu *et al.*, 2023] Chen Yu, Weinan Zhang, Hang Lai, Zheng Tian, Laurent Kneip, and Jun Wang. Multi-embodiment legged robot control as a sequence modeling problem. In *ICRA*, 2023.
- [Yue *et al.*, 2019] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *ICCV*, 2019.
- [Zakka *et al.*, 2021] Kevin Zakka, Andy Zeng, Pete Florence, et al. Xirl: Cross-embodiment inverse reinforcement learning. In *CoRL*, 2021.
- [Zhang *et al.*, 2019] Jingwei Zhang, Lei Tai, Peng Yun, et al. Vr-goggles for robots: Real-to-sim domain adaptation for visual control. *RAL*, 2019.
- [Zhang *et al.*, 2020] Qiang Zhang, Tete Xiao, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Learning cross-domain correspondence for control with dynamics cycle-consistency. In *ICLR*, 2020.
- [Zhu *et al.*, 2017] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.
- [Zintgraf *et al.*, 2019] Luisa Zintgraf, Kyriacos Shiarlis, Maximilian Igl, et al. Varibad: A very good method for bayes-adaptive deep rl via meta-learning. In *ICLR*, 2019.