


Transfer learning in robotics: An upcoming breakthrough? A review of promises and challenges

The International Journal of
Robotics Research
2024, Vol. 0(0) 1–21
© The Author(s) 2024



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649241273565
journals.sagepub.com/home/ijr



Noémie Jaquier^{1*} , Michael C Welle^{2*}, Andrej Gams³, Kunpeng Yao⁴, Bernardo Fichera⁴, Aude Billard⁴, Aleš Ude^{3,5}, Tamim Asfour¹ and Danica Kragic²

Abstract

Transfer learning is a conceptually-enticing paradigm in pursuit of truly intelligent embodied agents. The core concept—reusing prior knowledge to learn in and from novel situations—is successfully leveraged by humans to handle novel situations. In recent years, transfer learning has received renewed interest from the community from different perspectives, including imitation learning, domain adaptation, and transfer of experience from simulation to the real world, among others. In this paper, we unify the concept of transfer learning in robotics and provide the first taxonomy of its kind considering the key concepts of robot, task, and environment. Through a review of the promises and challenges in the field, we identify the need of transferring at different abstraction levels, the need of quantifying the transfer gap and the quality of transfer, as well as the dangers of negative transfer. Via this position paper, we hope to channel the effort of the community towards the most significant roadblocks to realize the full potential of transfer learning in robotics.

Keywords

Transfer learning, imitation learning, domain adaptation, sim-to-real, task transfer, embodiment transfer

Received 29 November 2023; Revised 2 May 2024; Accepted 4 July 2024

1. The rise of transfer learning in robotics

Transferring prior knowledge to novel unknown tasks is one of the abilities that led humans to become the most innovative species on the planet (Reader et al., 2016). In particular, humans' capability to transfer cognitive (Barnett and Ceci, 2002; Perkins and Salomon, 1992) and motor skills (Schmidt and Young, 1987) from one context to another makes the acquisition of new skills and the resolution of problems possible to a large extent. For instance, the difficulty of learning a new language is significantly influenced by factors such as language distance, native language proficiency, and language attitude (Walqui, 2000) as humans can *transfer* their prior experience, for example, grammatical constructions or words, from their native language into the new language. In addition, transfer is a key concept in education as the context of learning, for example, the classroom, significantly differs from the context, for example, the workplace, where the learned concepts should ultimately be applied (Perkins and Salomon, 1992).

To evolve seamlessly in the real world, robots must feature outstanding cognitive abilities allowing them to perceive their environment, act and react to achieve various

goals, and learn continuously from observation and experience, while coping with changes and uncertainty in the world. The transfer learning paradigm for robotics is a promising avenue to avoid learning from scratch by reusing previously-acquired experience in new situations, similar to humans. The core idea of transfer learning in robotics, illustrated in Figure 1, is simple: The experience of a robot

¹Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Karlsruhe, Germany

²KTH Royal Institute of Technology, Stockholm, Sweden

³Jožef Stefan Institute, Ljubljana, Slovenia

⁴Learning Algorithms and Systems Laboratory, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

⁵Faculty of Electrical Engineering, University of Ljubljana, Slovenia

*These Noémie Jaquier and Michael C Welle authors contributed equally (listed in alphabetical order).

Corresponding authors:

Noémie Jaquier, Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Adenauerring 2, Karlsruhe 76131, Germany. Email: noemie.jaquier@kit.edu

Michael C Welle, KTH Royal Institute of Technology, Stallgatan 9A, Stockholm, Upplands Väsby 19432, Sweden. Email: mwelle@kth.se

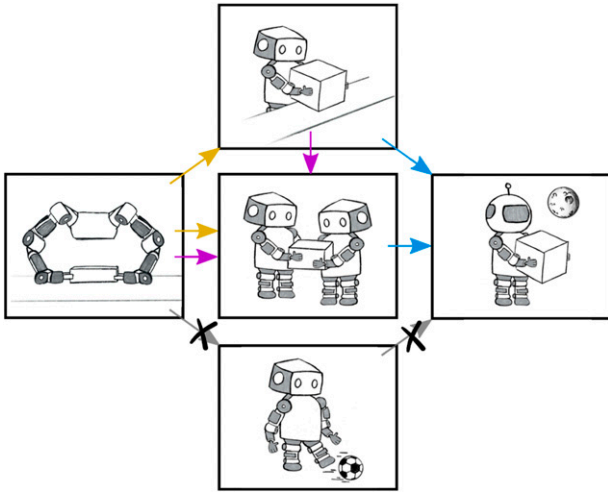


Figure 1. Concept of transfer learning in robotics. The experience of a robot performing a specific task in a specific environment is leveraged to improve the learning of a related task by another robot in a related context. Transfer can occur across embodiments (yellow arrows), across tasks (purple arrows), and/or across environments (blue arrows). It is important to note that successful transfer requires commonalities between the source and target robots, tasks, and environments. For instance, a humanoid robot learning to kick a ball will most likely not benefit from the experience of a dual-arm manipulator systems manipulating a box and vice versa.

performing one task in an environment is leveraged to improve the learning process of a (related) task in a different context, that is, in a different environment or executed by a different robot.

To identify when transfer learning is needed and/or warranted, one need to identify the similarities and differences between the two situations. Following the concepts proposed in imitation learning (Dautenhahn and Nehaniv, 2002), we distinguish between three aspects that can be deemed similar or different: *The tasks, the environments, and the bodies* (a.k.a. the robots in our case).

In the conceptual example of Figure 1, the experience of the two fixed-based manipulators placing a box on a conveyor belt (left) can be transferred to a humanoid robot executing the same task (middle-top). Transfer learning is made possible and easier as the two situations have many aspects in common:

- (1) **Tasks:** Both tasks are similar as both robots are engaged in moving an object with two arms. The notion of similarity with respect to the object may be relative to the robot's own size and payload. For instance, a pair of industrial arms may be able to support objects up to 20 kg whereas a humanoid robot may carry only a quarter of this. Yet, the task and strategy to solve it, namely ensuring coordinated movement of both arms and the trajectory to place an object on a conveyor's belt, remains the same.
- (2) **Environments:** The two environments are similar, as in both cases, the object is to be placed on a conveyor

belt. We can also safely assume that, in both cases, the object is subjected to the same external forces (gravity, friction).

- (3) **Robots:** In this case, the two robots differ. Yet, their differences can be broken into two distinct parts. Both robots are endowed with two arms of similar structure. Solely the second robot is on legs and hence faces the additional challenge of controlling its balance when placing the object on the conveyor belt. This aspect may be resolved without actual learning, for instance, through constrained-based optimization, controlling for additional constraints at the center of mass (Bouyarmane et al., 2018). Alternatively, reinforcement learning or adaptive control may be used to fine tune gains (Khadivar et al., 2023).

Hence, in this case, transfer learning is mostly conducted across the different robots, while taking advantage of their similar bimanual structure. In other cases, when the two robots and/or environments differ importantly, experience may be transferred across tasks. For instance, the bimanual manipulation strategy used when placing a box on a conveyor belt (Figure 1, middle-top) may be transferred to a handover task (middle). Such transfer may also be achieved across different tasks and different embodiments: The bimanual strategy of the two fixed-based manipulators (left) may directly be transferred to a different robot executing a related task, for example, to the humanoid robot handing over the box (middle). Experience may also be transferred across environments. For example, the bimanual manipulation strategies used by the humanoid robot to place a box on the conveyor belt or to handover the box (Figure 1, middle-top and middle) may be reused by a humanoid robot manipulating a box in space (right). In this example, the physical rules of the two environments strongly differ due to the influence of gravity on Earth.

Finally, in some cases, transferring knowledge may not be beneficial or may even impede the robot performance. For instance, we consider transferring experience from the fixed-based bimanual setup placing a box in a conveyor belt to a humanoid robot kicking a ball. In this case, employing transfer learning might not be beneficial. Instead, the large differences between the robots, tasks, and environments might even cause transfer learning to impede the methods employed to fulfill the ball-kicking task. It is important to highlight that, even if some particular actions can be transferred, conflicting goals may lead to negative transfer. This clearly showcases the crucial importance of identifying the commonalities and differences between two situations when applying transfer learning in robotics. Moreover, it highlights the importance of transfer learning metrics that measure not only the transfer quality, but also the transfer gap between robots, tasks, and environments.

While the terms falling under the umbrella of transfer learning in robotics are not exactly agreed upon, the robotics community has intensified its effort to transfer various forms of knowledge across different contexts. Figure 2 (top)

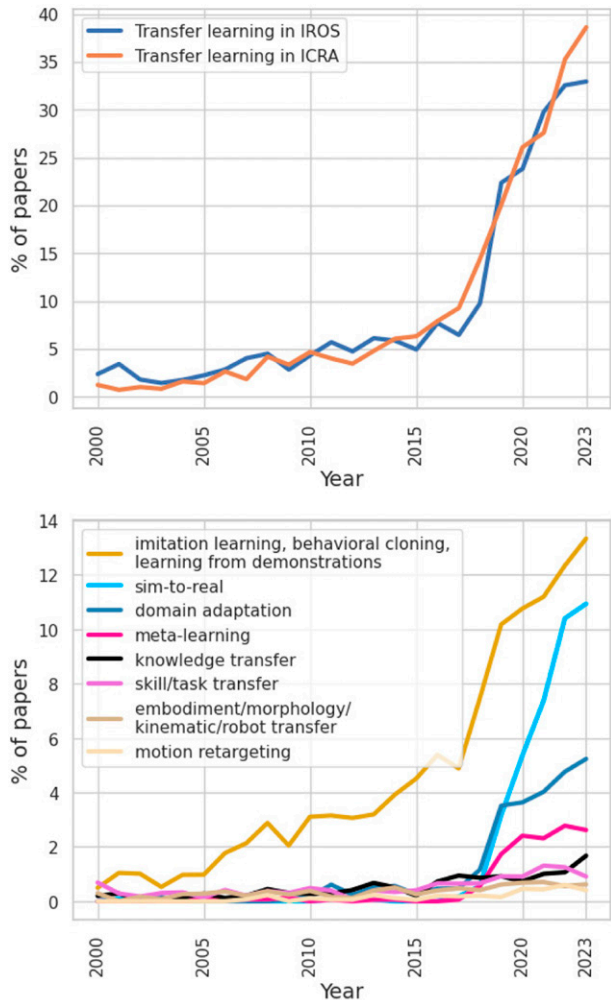


Figure 2. Percentage of papers including words falling under transfer learning in robotics umbrella over the years for the two biggest robotics conferences. The results are based on a systematic search through the content of 44,067 papers, excluding their references. *Top:* Percentage per conference. *Bottom:* Percentage per keywords. Keywords related to transfer across embodiments, tasks, and environments are depicted in variations of yellow, purple, and blue. Knowledge transfer is classified independently and thus is depicted in black.

shows the recent rise in the proportion of published papers at the two biggest robotics conferences—IROS and ICRA—invoking keywords¹ that we consider as falling under the umbrella of transfer learning. This recent interest shows that the community strives for embodied transfer learning, which may be a necessary a priori for truly intelligent systems (Kremelberg, 2019). It is interesting that different keywords related to transfer learning in robotics display different growth rates (see Figure 2, bottom). For instance, terms related to imitation learning, behavioral cloning, and learning from demonstrations display an early rise and a high popularity nowadays, highlighting the effort of the robotics community to tackle transfer between embodiments (human-to-robot or robot-to-robot) from early on (Dautenhahn and Nehaniv, 2002). In contrast, terms related

to transfer across environments (i.e., sim-to-real and domain adaptation) only recently gained interest, while terms related to transfer across tasks remain less documented, suggesting that transfer learning across environments and tasks is still in its infancy.

In this position paper, we contend that transfer learning in robotics has the potential to revolutionize the robot learning paradigm by enabling robots to leverage past experience in novel contexts. However, key challenges remain to be addressed to fulfill this potential. In particular, the fundamental question of *identifying the similarities and differences across tasks, environments, and robots in an automatic manner* remains. This position paper reviews the successes of the field and identifies relevant research questions and promising directions paving the way forward. Starting from the definition of transfer learning in the machine learning field, we propose a unified definition of transfer learning in robotics and subsequently build a novel taxonomy of transfer learning in robotics based on the key concepts of robot, task, and environment (see Section 2). Then, we recount successful applications of transfer learning in robotics and show how they align with the proposed taxonomy (Section 3). In Section 4, we outline challenges, as well as promising research directions to tackle them, including abstraction levels and universal representations for transfer learning in robotics, interpretability, benchmarks and simulations, transfer learning metrics, and dangers of negative transfer. Last but not least, we call for actions to address the most immediate roadblocks in Section 5. Overall, the contributions of this paper are twofold: (1) We provide a unified view of transfer learning in robotics by comprehensively defining the notion of transfer learning in robotics and by introducing the first taxonomy of transfer learning in robotics; (2) We provide a review of promises and challenges in the field of transfer learning in robotics, identifying the most significant roadblocks on the way to unraveling the full potential of transfer learning in robotics.

2. Transfer learning taxonomies: From machine learning to robotics

While the machine learning community has devoted substantial efforts to defining and systematizing transfer learning and categorizing its different instances, transfer learning in robotics is found under various terminologies. This section aims at providing a unified view of transfer learning in robotics. To do so, we take inspiration from machine learning taxonomies and define a taxonomy of transfer learning settings that occur in robotics.

2.1. Taxonomy of transfer learning in machine learning

In this section, we introduce the definitions that are commonly adopted in the transfer learning community, see, for example, Pan and Yang, 2010; Yang et al., 2020;

Zhuang et al., 2020). Transfer learning builds on the two fundamental concepts of domain and task. A *domain* $\mathcal{D} = \{\mathcal{X}, p(X)\}$ consists of a feature space \mathcal{X} and a marginal distribution $p(X)$, with X denoting an instance set of the feature space such that $X = \{\mathbf{x}_i\}_{i=1}^N$, $\mathbf{x}_i \in \mathcal{X}$. Given a specific domain \mathcal{D} , a *task* $\mathcal{T} = \{\mathcal{Y}, f\}$ consists of a label space \mathcal{Y} and a predictive function $f: \mathcal{X} \rightarrow \mathcal{Y}$. The predictive function is used to predict new labels $\mathbf{y} \in \mathcal{Y}$ associated with a new instance $\mathbf{x} \in \mathcal{X}$. It is typically learned from a training dataset $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^M$, with $\mathbf{x}_i \in \mathcal{X}$, $\mathbf{y}_i \in \mathcal{Y}$. While standard machine learning approaches assume that training and test datasets share common domains and tasks, in the case of transfer learning they may instead belong to different spaces, referred to as *source* and *target* spaces. Therefore, it also has the potential to tackle open-set problems (Geng et al., 2021). An example from computer vision is shown in Figure 3.

Transfer learning approaches are commonly categorized into inductive, transductive, and unsupervised transfer learning (Pan and Yang, 2010; Yang et al., 2020; Zhuang et al., 2020). This categorization focuses on the availability of labels independently of the relationships between source and target spaces. Namely, in the inductive setting, labels are available in both source and target spaces, while they are only available in the source space in the transductive setting, and are not available in any space in the unsupervised setting. For a more encompassing categorization that generalizes to robotics, we instead propose to focus on the fundamental concepts of domain and task and on their relationship in the source and target spaces.

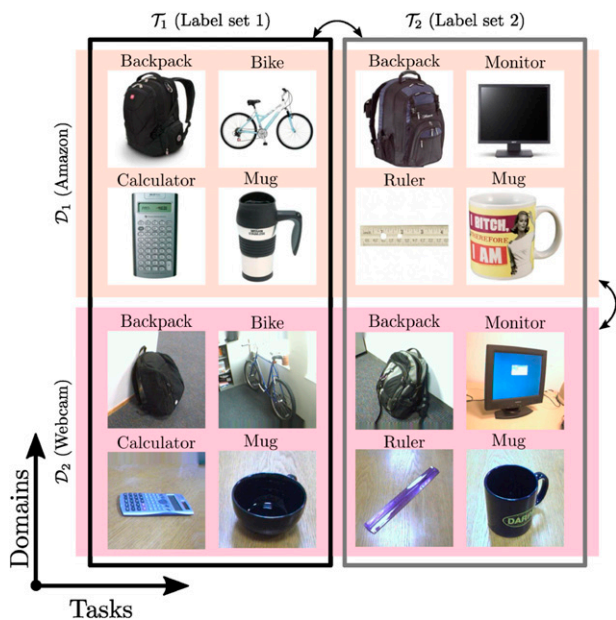


Figure 3. Example of transfer learning on the Office 31 dataset (Saenko et al., 2010). Transfer can occur (1) between the two label sets (tasks \mathcal{T}_1 and \mathcal{T}_2), (2) between the two sources used to obtain the images (domains \mathcal{D}_1 and \mathcal{D}_2), and (3) between both tasks and domains. The different transfer learning instances are illustrated with black arrows.

Definition 2.1. Transfer Learning in Machine Learning. Let $S = \{\mathcal{D}_S, \mathcal{T}_S\}$ a source space and $T = \{\mathcal{D}_T, \mathcal{T}_T\}$ a target space. The objective of transfer learning is to improve the learning of the predictive function f_T over the target domain \mathcal{D}_T by taking advantage of knowledge from the source domain \mathcal{D}_S , and task \mathcal{T}_S , where $\mathcal{D}_S \neq \mathcal{D}_T$ and / or $\mathcal{T}_S \neq \mathcal{T}_T$.

We observe that Definition 2.1 implies the following hierarchical taxonomy of transfer learning settings illustrated in Figure 4. Note that source space may correspond to a union of multiple sub-source tasks and/or domains.

(1) Task transfer learning: $\{\mathcal{D}, \mathcal{T}_S\} \rightarrow \{\mathcal{D}, \mathcal{T}_T\}$. In this setting, the target task differs from the source task, $\mathcal{T}_S \neq \mathcal{T}_T$. This can indicate a difference in the label space, the predictive function or both. Notice that we refer to task transfer learning whenever the source and target domains are identical while the source and target tasks differ. The extent and conditions of the task differences were further categorized according to various transfer learning taxonomies, see, for example, Pan and Yang, 2010; Yang et al., 2020; Zhuang et al., 2020). Task transfer learning approaches include learning strategies involving Gaussian process (GP) prior sharing across different tasks (Bonilla et al., 2007; Lawrence and Platt, 2004). Other strategies focus on sharing the parameters of the model itself rather than the hyperparameters. One important category of algorithms in this domain takes advantage of a modified version of support vector machine (SVM) to transfer knowledge between source and target spaces (Evgeniou and Pontil, 2004; Li et al., 2012). In this modified SVM, the model’s parameters consist of a part shared across the source and target spaces, while the other part is space-specific. The uniqueness of the solution (learning efficiency) and model interpretation make these convex optimization algorithms an interesting solution for robotics. Multilinear relationship networks (Long et al., 2017) leverages labeled data from related source domains by adopting a Bayesian framework for the task-specific portion of the network.

(2) Domain transfer learning: $\{\mathcal{D}_S, \mathcal{T}\} \rightarrow \{\mathcal{D}_T, \mathcal{T}\}$. This form of transfer learning occurs when the source and target tasks are identical, $\mathcal{T}_S = \mathcal{T}_T$, but the source and target domains differ. The condition $\mathcal{D}_S \neq \mathcal{D}_T$ can indicate a difference in the feature space, in the marginal distribution, or in both. It is generally assumed that the domains are related to a certain extent.

The difference in marginal distribution is often tackled by learning a mapping between overlapping instances (also known as “support”) between the source and target domains \mathcal{D}_S and \mathcal{D}_T . Such approaches primarily rely on instance weighting strategies, such as assigning weights to instances or labeled data in \mathcal{D}_S for reuse in \mathcal{D}_S . For instance, kernel

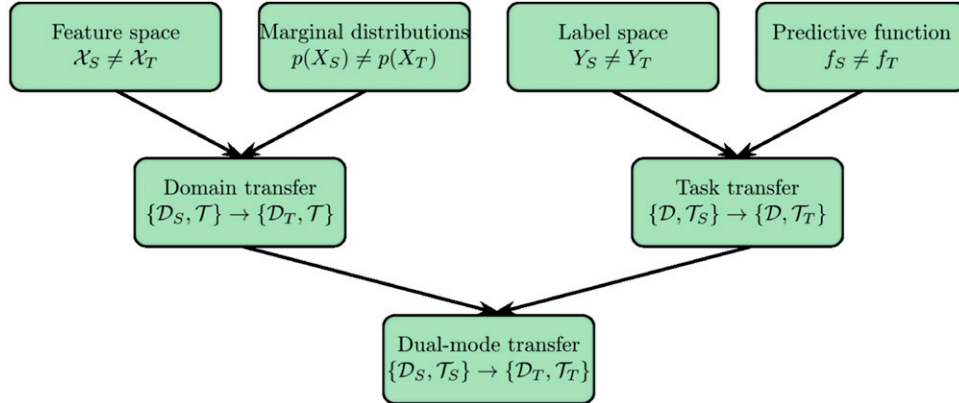


Figure 4. Hierarchical taxonomy of transfer learning in the context of machine learning. Domain transfer occurs when the source and target domain differ and is characterized by a difference in the feature space and/or in the marginal distribution. Task transfer occurs when the source and target task differ and is characterized by a difference in the label space and/or in the predictive function. Dual-mode transfer occurs when both the source and target domain and task differ.

mean matching (Huang et al., 2006) matches source and target domain instance means within a reproducing kernel Hilbert space. In the case of different feature spaces ($\mathcal{X}_S \neq \mathcal{X}_T$), existing approaches aim at reducing domain differences while preserving the properties or structures within the same domain. For instance, structural correspondence learning (Blitzer et al., 2006) utilizes pivot features to establish pseudo tasks connected to the target task and applies multi-task learning techniques to model relationships between pivot features and other features. Spectral feature alignment (Pan et al., 2010) models interdependencies between pivot features and other features using a bipartite graph and identifies novel common features through spectral clustering methods applied to the graph.

- (3) Dual-mode transfer learning: $\{\mathcal{D}_S, \mathcal{T}_S\} \rightarrow \{\mathcal{D}_T, \mathcal{T}_T\}$. In this scenario, both the source and target domains and tasks differ. This is the most challenging setting in transfer learning as every additional difference between the source and target label space, predictive function, feature space, and marginal distribution increases the complexity of the problem. Approaches in the dual-mode are currently mostly related to unsupervised transfer learning scenarios. This remains an underexplored area due to the difficulty of capturing the similarities—or the transferable information (instance, feature, parameter, etc.)—between the source and target spaces.

Figure 3 illustrates the aforementioned transfer learning setting using the transfer dataset Office 31 (Saenko et al., 2010), a classical benchmark for transfer learning. In this case, the domains correspond to the source used to obtain the images (i.e., Amazon and a webcam in Figure 3) and the tasks correspond to the labels of the objects represented in the images.

To extend these concepts to robotics, we must consider the robot as an additional mode. In the next section, we discuss its implications for transfer learning in robotics.

2.2. Taxonomy of transfer learning in robotics

Transfer learning in robotics builds on the three fundamental concepts of robot, environment, and task. A *robot* \mathcal{R} is defined as an embodiment that can act in and thus influence its environment. It encompasses a body with defined morphology, kinematics, dynamics, and sensor modalities. In robotics, the domain \mathcal{D} is generally considered equivalent to the *environment*, which is defined as the virtual or physical world in which the robot lives and interacts. The robot accesses the state of the environment via sensory observations, for example, images, contact forces, auditory and olfactory signals. Informally, the *task* \mathcal{T} refers to what the robot is required to do in the environment. More formally, a task is a discrete or continuous (sub)goal that can be achieved by the robot through (inter)actions within the environment. In general, the goal of the robot is to perform a given task in the environment. The goal of transfer learning in robotics is to leverage prior knowledge from a source space, composed by a robot, a task, and an environment, to improve the performance in a target space, where one or more mode differs from the source space. Formally, we define transfer learning in robotics as an analogy of the machine learning Definition 2.1 as follows.

Definition 2.2. Transfer Learning in Robotics. Let $S = \{\mathcal{R}_S, \mathcal{D}_S, \mathcal{T}_S\}$ a source space and $T = \{\mathcal{R}_T, \mathcal{D}_T, \mathcal{T}_T\}$ a target space. The objective of transfer learning in robotics is to improve the performance of the robot \mathcal{R}_T executing the task \mathcal{T}_T in the environment \mathcal{D}_T by taking advantage of knowledge from the source robot² \mathcal{R}_S , environment \mathcal{D}_S , and task \mathcal{T}_S , where at least one element of the target space T is different from its counterpart in the source space S .

It is important to emphasize that, unlike transfer learning in machine learning which only involves disembodied agents, the agent’s *embodiment*—in other words, the *robot*—is key for transfer learning in robotics. This introduces additional challenges: The presence of a robot not only adds an additional mode to the transfer learning problem and thus to the hierarchical categorization, but also brings numerous robotics-specific issues. Transfer learning methods for robotics must cope with the fact that robots are embodied agents that act and interact in the real world.

Inspired by the hierarchical taxonomy defined for transfer learning in machine learning community, we propose a hierarchical taxonomy for transfer learning in robotics based on the relationship between the source and target robots, environments, and tasks, as shown in Figure 5. Specific illustrative examples of its categories are depicted in Figure 6. Our taxonomy considers the following settings:

- (1) Robot transfer learning: $\{\mathcal{R}_S, \mathcal{D}, \mathcal{T}\} \rightarrow \{\mathcal{R}_T, \mathcal{D}, \mathcal{T}\}$. The goal of this setting is to endow a target robot with the ability to perform a given task known by other source robot(s) in the same environment. Note that the source and target robots may have (very) different morphologies, kinematics, and sensor modalities, leading to different capabilities. For example, Figure 6 illustrates a transfer between the humanoid dual-arm robot ARMAR-III (Asfour et al., 2006) and ALMA (Bellicoso et al., 2019), a four-legged robot equipped with a robotic arm. Moreover, the transfer can happen at different actions levels, for example, at the level of joint or task-space controllers or at the planning level, and at different perceptual levels, for example, across different sensors (Tatiya et al., 2020) or sensory modalities (Lee et al., 2020b). Instances of robot transfer learning are (i) imitation learning (Schaal, 1999), where a teacher human or robot provides demonstrations of a task to a student robot that learns to reproduce the given task in the same environment, (ii) (goal-directed) motion retargeting (Dariush et al., 2008; Yin et al., 2023), whose goal is to learn a mapping between different kinematic structures, and (iii) perceptual

transfer, where the robots are equipped with different sensory modalities (Silva et al., 2020) such as touch, vision, sound, or olfaction.

- (2) Environment transfer learning: $\{\mathcal{R}, \mathcal{D}_S, \mathcal{T}\} \rightarrow \{\mathcal{R}, \mathcal{D}_T, \mathcal{T}\}$. This setting aims at transferring the ability of a robot to perform a given task in a source environment to a different target environment. Its main challenge is to overcome the mismatch between source and target environments in terms of data and environment parameters, for example, underlying dynamics or transition models. For instance, models learned for a specific task performed on earth typically needs to be adapted to perform the same task in an underwater environment or in space. This requires identifying which physical parameters differ between \mathcal{D}_S (the earth) and \mathcal{D}_T (the underwater environment, or space). Typical instances of environment transfer learning are (i) domain adaptation (Bousmalis et al., 2018; Wang and Johnson, 2021), and (ii) sim-to-real transfer (Muratore et al., 2022). The latter is a particular case of the former in which the experience is explicitly transferred from a simulation environment—in which training data are inexpensive and models are fast to train—onto the real world. Sim-to-real transfer is showcased by the first and second rows in Figure 6. The second and third rows indicate transfer between two real-world environments, where the objects composing the physical environment (the cloth or the box) differ.
- (3) Task transfer learning: $\{\mathcal{R}, \mathcal{D}, \mathcal{T}_S\} \rightarrow \{\mathcal{R}, \mathcal{D}, \mathcal{T}_T\}$. This setting aims at leveraging the ability of a robot to perform a given task to learn how to execute a different task in the same environment. The underlying assumption is that the source and target tasks are—to some extent—similar, so that experience can be reused between source and target tasks. For instance, the box tossing and cloth flinging tasks of Figure 6 share similar dynamics characteristics: In both cases, the robot must generate high-velocity dynamic actions to successfully execute the task. Therefore, we may expect that experience about box tossing may be reused by the robot when learning to fling a cloth. Challenges of task transfer learning include inferring

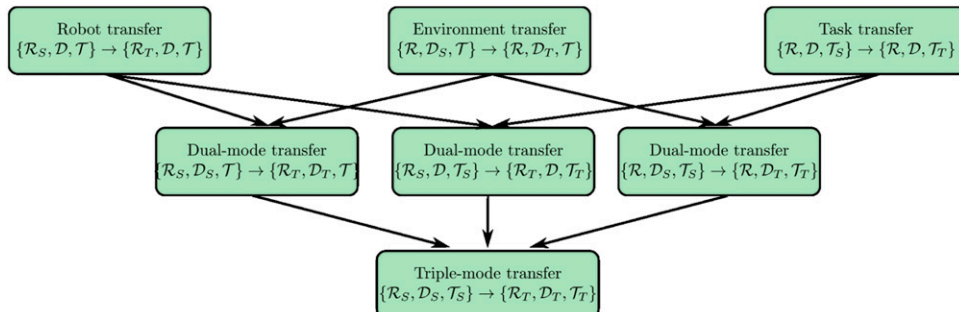


Figure 5. Hierarchical taxonomy of transfer learning in the context of robotics. Robot, environment, and task transfer occur when the source and target robot, environment, and task differ, respectively. Dual-mode and triple-mode transfer occur when two of these modalities and the three of them differ, respectively.

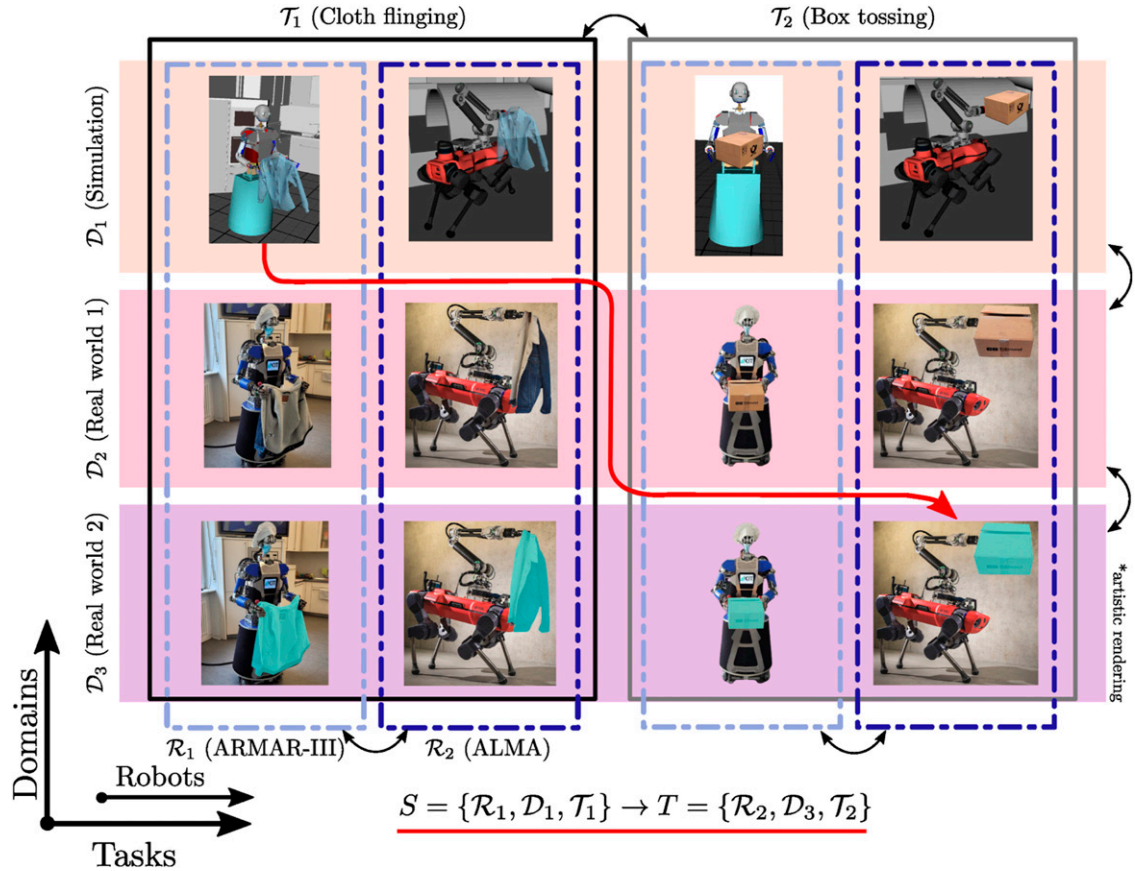


Figure 6. Illustration of the categories of the hierarchical taxonomy for transfer learning in robotics. The humanoid dual-arm robot ARMAR-III (\mathcal{R}_1) and the four-legged robot ALMA equipped with a manipulator (\mathcal{R}_2) execute a cloth flinging task (\mathcal{T}_1) and a box tossing task (\mathcal{T}_2) in three different environments, namely a simulator (\mathcal{D}_1), and in the real world with different object instances (\mathcal{D}_2 and \mathcal{D}_3). Transfer learning can occur (1) between the two robots, (2) between two environments, (3) between the two tasks, and (4-5) between two or three instances thereof. The different transfer learning instances are illustrated with black arrows. Triple-mode transfer learning, which reuses knowledge from a source space to a target space with different robots, environments, and tasks, is depicted by a red arrow.

which part of the source task experience should be transferred and at which level (joint or task-space, planning, etc). Notice that generalizing a given task to an unseen context is a special case of task transfer learning (Li and Figueroa, 2023; Mandlekar et al., 2020). In this case, the model is made compatible with different instances of the same task. Curriculum learning (Narvekar et al., 2020; Shukla et al., 2022) is also a special case of task transfer learning, where a sequence of intermediary tasks of gradually-increasing difficulty is used to learn a complex target task. Finally, task transfer is also particularly considered in the areas of lifelong learning, where task transfer is considered based on a never-ending stream of data, and compositional learning, which focuses on transfer across compositionally-related tasks (Mendez and Eaton, 2021, 2023).

- (4) Dual-mode transfer learning: $\{\mathcal{R}_S, \mathcal{D}_S, \mathcal{T}\} \rightarrow \{\mathcal{R}_T, \mathcal{D}_T, \mathcal{T}\}$, $\{\mathcal{R}_S, \mathcal{D}, \mathcal{T}_S\} \rightarrow \{\mathcal{R}_T, \mathcal{D}, \mathcal{T}_T\}$, or $\{\mathcal{R}, \mathcal{D}_S, \mathcal{T}_S\} \rightarrow \{\mathcal{R}, \mathcal{D}_T, \mathcal{T}_T\}$. This setting is concerned by transferring knowledge between two spaces which differ across two modes. It assumes that the similarities

between source and target spaces can still be leveraged when they share a single common mode. For instance, in Figure 6, it is reasonable to assume that experience acquired in simulation to toss a box may be reused to fling a cloth with the same robot in the real world. Dual-mode transfer learning remains largely unexplored in robotics due to the additional level of complexity compared to the single-mode transfer setting listed above, which has not yet been fully resolved.

- (5) Triple-mode transfer learning: $\{\mathcal{R}_S, \mathcal{D}_S, \mathcal{T}_S\} \rightarrow \{\mathcal{R}_T, \mathcal{D}_T, \mathcal{T}_T\}$. This setting assumes that all three modes of the source and target spaces differ. It is inspired by the human ability to successfully acquire knowledge by observing others executing similar tasks in different environments. For instance, one may observe a chef cooking a pie in a restaurant kitchen and reuse some of her techniques to cook a cake in her own non-professional kitchen. Reusing knowledge from a source space in a target space with different robots, environments, and tasks would endow robots with human-like generalization abilities. This setting is the

most challenging, and requires bridging the gaps between high-level semantic information—indicating the degree of similarity between spaces—and low-level actions. It is the ultimate goal of transfer learning in robotics, as indicated by the red arrow in Figure 6.

Notice that, depending on the relationship between the source and target spaces, our Definition 2.2 intrinsically refers to related fields, some of which received significant attention over the years. In Figure 2 (bottom), we notably observe that imitation learning is the most mentioned transfer learning field followed by sim-to-real and domain adaptation. In this sense, we view transfer learning in robotics as an umbrella term that encompasses “imitation learning,” “learning from demonstrations,” “sim-to-real,” “domain adaption,” “meta-learning,” “knowledge transfer,” “skill transfer,” “motion retargeting,” “embodiment transfer,” “morphology transfer,” and “kinematic transfer,” among others.

3. Successes of transfer learning in robotics

Change of environment or domain as in $\{\mathcal{R}, \mathcal{D}_S, \mathcal{T}\} \rightarrow \{\mathcal{R}, \mathcal{D}_T, \mathcal{T}\}$, change of task as in $\{\mathcal{R}, \mathcal{D}, \mathcal{T}_S\} \rightarrow \{\mathcal{R}, \mathcal{D}, \mathcal{T}_T\}$ and change of the robot as in $\{\mathcal{R}_S, \mathcal{D}, \mathcal{T}\} \rightarrow \{\mathcal{R}_T, \mathcal{D}, \mathcal{T}\}$ have all been addressed with varying success in transfer learning in robotics. The body of literature on the topic is extremely vast, making a comprehensive overview beyond the scope of this paper. On the other hand, research activities that by definition fit into the scope of transfer learning have been addressed before the term took root in robotics. An example of such is imitation learning (Schaal, 1999), where task execution knowledge is transferred from the human to the robot, or generalization, where task execution knowledge is transferred to (at least) a variation of the task (Ude et al., 2010). In the following, we provide examples of transfer learning in robotics, also in the light of such above-mentioned applications.

3.1. Environment transfer

Change of environment conditions (Kramberger et al., 2016) or the complexity of the environment where the task is being executed (Vosylius and Johns, 2023) provide examples of generalization to a declaratively new environment. However, the environment (domain), can be different in other aspects that go beyond just the setting—for example, contact conditions or other physical conditions might not be the same (Muratore et al., 2022). One example of such is transfer from the simulation-to-reality or sim-to-real.

Potentially unjustly, but transfer learning in robotics is often associated exactly with sim-to-real, where typically experience is obtained in one domain—the simulation—and exploited to accelerate learning in the transferred domain—the real world. Several reviews cover sim-to-real transfer learning in robotics, that is, (Muratore et al., 2022; Zhao

et al., 2020), affirming the notion of a huge body of work in this field. The gist of sim-to-real lies in the notion that collecting the data for modern (deep) learning and other AI algorithms in the real world is too expensive in terms of time and resources to scale up (Muratore et al., 2022). Therefore, the data is collected in simulation, despite the difference between the real and simulated domains. This difference, referred to as the “reality gap” (Collins et al., 2019), needs to be overcome for real-world execution, which is done using transfer learning. Since collecting data in the real world is so time-consuming and expensive, researchers might change the domain to a different simulation, ending up with sim-to-sim methodologies. These are applied to demonstrate the behavior of transfer learning methodologies.

Different practices have been proposed for sim-to-real transfer learning, starting with realistic modeling (Muratore et al., 2022). No matter how accurate, modeling will never be fully cover all the aspects of the real world (Muratore et al., 2022), thus other approaches have emerged. Domain randomization, such as randomization of image backgrounds, of physical parameters of objects and robot actions, or of controller parameters (Höfer et al., 2021), is a common approach. By randomizing over, for example, physical parameters, the approach tries to cover the entire spectrum of these parameters in the hope that this includes the parameters that describe the real world. Even so, one-shot transfer learning is seldom successful (Zhao et al., 2020), and additional learning is required, for example using reinforcement learning (Ada et al., 2022), back-propagation (Chen et al., 2018) or both (Lončarević et al., 2022). If there are significantly fewer learning iterations in the target domain, the process is called few-shot transfer learning (Ghadirzadeh et al., 2021). Few-shot transfer learning has notably been applied to sim-to-real transfer in robotics (Bharadhwaj et al., 2019; Shukla et al., 2023). Given that more information can be available in the simulation, the notion of privileged learning was introduced, where the privileged information is used to train a high performance policy, which in turn trains a proprioceptive-only student policy (Lee et al., 2020a). The idea was very successfully demonstrated in quadrupedal locomotion by more than one group (Kumar et al., 2022; Lee et al., 2020a), and is general enough to be applied for very different tasks, such as excavator walking (Egli and Hutter, 2022) and even robotized handling of textiles (Longhini et al., 2022).

3.2. Task transfer

Transferring of robot walking from one domain to the other can be considered more than just domain transfer, as walking itself can be different for different environments. For instance, a pacing gait learned to walk on smooth ground might not be stable enough for walking on mountainous terrains. Thus, walking on mountainous terrains can be seen as a novel task, which may benefit from transferring previously-learned gaits adapted to other terrains. Moreover, walking is not an isolated instance: If the

robot can learn to throw accurately at one target, a modulation of the throwing task to aim at a different target can in fact be considered at the least a different instance of the same task, if not a different task overall.

Such transfers from one (or several) task instances to a new one have been utilized in robotics before, and were often referred to as generalization. In this sense, for example, fast learning from a small set of demonstrations was applied with nonlinear autonomous dynamical system (DS), which have the ability to generalize motions to unseen contexts (Khansari-Zadeh and Billard, 2011). Similarly, a set of dynamical systems in the form of dynamic movement primitives was used to generalize to transfer knowledge from known situations to unknown in positions (Ude et al., 2010) and in torques (Deniša et al., 2016), probabilistic movement primitives (ProMPs) encode complete families of motions (Paraschos et al., 2013), TP-GMMs adapt to changes of predefined local frames (Calinon, 2018), and Mixture Density Networks adapt a learned motion primitive to new targets specified in a different space (Zhou et al., 2020).³ Generalization was even termed inter-task transfer learning (Fernández et al., 2010).

Task transfer has also been tackled via meta-learning. In the meta-learning setting, a model is trained on a variety of tasks so that new tasks are solved by using none (zero-shot) or only a limited amount (few-shot) of additional training data (Finn et al., 2017; Nichol et al., 2018). For instance, MAML (Finn et al., 2017) was shown to generalize to new goal velocities and directions in the half cheetah and ant locomotion tasks of the Gymnasium benchmark (Towers et al., 2023) faster than conventional approaches. Other meta-learning approaches tackle the transfer problem from a different perspective by learning loss functions (Bechtel et al., 2021, 2022). The meta-learned loss functions generalize to different tasks, thus alleviating the need of designing task-specific losses.

Thus, in a broad sense of Definition 2.2, such approaches already propose solutions for $\{\mathcal{R}, \mathcal{D}, \mathcal{T}_S\} \rightarrow \{\mathcal{R}, \mathcal{D}, \mathcal{T}_T\}$, although they were not called transfer learning. Complete skill models were learned from a set of executions also with DNNs (Lončarević et al., 2022). The adaptation of the skill model for a new environment is commonly referred to as transfer learning.

3.3. Robot transfer

Above mentioned approaches use knowledge from several instances of a task. However, learning of even one instance of a task could pose a challenge. Imitation learning, where human skill knowledge was transferred to a robot, has been thoroughly researched as the means for learning of task models and their execution on a robot (Billard et al., 2008; Ravichandar et al., 2020). Imitation learning (IL), also known as programming by demonstration (PbD), is in a strict sense an example where the task and the environment remain the same, but the agent is different, $\{\mathcal{R}_S, \mathcal{D}, \mathcal{T}\} \rightarrow \{\mathcal{R}_T, \mathcal{D}, \mathcal{T}\}$, since one of the agents is in fact

a person. Note that, in some cases, the environment can also change. In PbD one often transfers the demonstrated motion (Ijspeert et al., 2013). However, if only the motion is repeated, the task knowledge might be overlooked and the task correspondence (Heyes, 2001) might not get preserved at all. This may be alleviated by transferring other crucial characteristics of the task, so-called task constraints, such as force patterns (Rozo et al., 2016) and posture-dependent task requirements (Jaquier et al., 2020), or by retargeting the demonstrated motion (Aberman et al., 2020), for example, by leveraging optimization methods (Rakita et al., 2017), learning approaches, or Riemannian geometry (Klein et al., 2022). Task descriptions in the form of reward functions learned from demonstrations are also promising for transferring tasks across different robots. For instance, cross-embodiment inverse reinforcement learning (XIRL) (Zakka et al., 2022) learns a notion of task progress from demonstrations, which is then used as a reward for robots with different embodiments that successfully learn to reproduce the task.

4. Challenges and promising research directions

The aforementioned examples highlight that knowledge can be transferred across several robots, tasks, and environments, thus highlighting the potential of transfer learning for robotics. However, several key questions falling under the areas of identifying the similarities and differences across tasks, environments, and robot to single out what should be transferred and when remain to be answered to realize the full potential of transfer learning in robotics. In this section, we describe the key challenges that currently constitute roadblocks on the way to the future of transfer learning in robotics.

4.1. Abstraction levels in robotics

Humans and some animals, such as great apes, acquire cognitive skills via the concept of social learning (Whiten and Ham, 1992), whose main component is to copy (transfer) behavior from one individual to another. Social learning takes place at different levels depending on the goal and context. In biology, the lowest level of transfer corresponds to *mimicry*, where an individual mimics the actions of another individual superficially, that is, without any underlying understanding of the goal (Genschow et al., 2017). Instead, with a number of methodological differences, *imitation* refers to an individual, that is, the learner, copying the actions of another individual, that is, the teacher, with the aim of achieving the same goal. As opposed to mimicry, imitation implies an explicit understanding of the goal. At the next level, *emulation* refers to the case where the learner aims at achieving the same goal as the teacher without copying their motor actions (Whiten et al., 2004). Combining imitation, emulation, and some other techniques such as object movement reenactment, the

agent ultimately develops an understanding of the world without having to understand the theoretical concept of causality.

These cognitive skill levels can also be roughly identified in robotics, where they intrinsically correspond to different *abstraction levels* (see Figure 7). At the lowest learning level, a robot simply mimics the motion of a teacher without an explicit understanding of the underlying goal. If the teacher and the learner have similar embodiments, the task can simply be abstracted as a joint-level (positions, velocity, or acceleration) trajectory. However, in the case of different embodiments, transferring joint trajectories will result in very different end-effector trajectories. The task’s abstraction level can be increased by specifying, for example, end-effector trajectories and leveraging Cartesian trajectory controllers. To deal with changes in the environment, both the learning and abstraction levels need to be raised. For instance, transferring end-effector trajectories fails if obstacles are present in the environment. In this case, the task needs to be imitated instead of mimicked, that is, the goal must be explicitly identified by the robot. The task can therefore be abstracted using, for example, movement primitives (Ijspeert et al., 2002), thus allowing the specification of the key components of the imitated trajectory, for example, the goal position, while leveraging robot skills such as collision avoidance, localization, and object detection to reproduce the task in different environments. In some cases, the physical capabilities of the teacher and the learner are very different, so the learner cannot achieve a demonstrated task by imitating the teacher. Instead, the learner must infer the goal from the demonstrated task and develop a strategy to achieve the same goal (Schaal, 1999). In other words, the transfer should be conducted at the higher abstraction level corresponding to achieving the goal specifications without imitating the teacher-specific actions. This corresponds to the emulation learning level. Finally, on an even higher abstraction level, the teacher should ideally give only high-level verbal instruction to the robot such as “open the drawer,” or “clean the room.” This requires the robot to have a skill set resembling that of agents with higher cognitive functions. Such capabilities have the potential to facilitate transfer in more challenging settings, such as dual- and triple-mode transfer.

To elaborate on the different levels of “abstraction,” consider the task of transferring a grasp performed by a *source* hand (robot or human) to a *target* robotic hand. This *transfer* can be performed on three levels: (1) The joint angle level (Bouzit, 1996; Kyriakopoulos et al., 1997) involves directly replicating joint angles with minor adjustments if the hands exhibit similar kinematic structures and degrees of freedom (DoFs); (2) The contact level (Maeda et al., 2016; Peer et al., 2008) is applicable when both hands have an equal number of fingers but differ in their kinematic properties (e.g., DoFs, finger lengths). In this scenario, the target hand strives to grasp the object by replicating the contact positions of the source hand; and (3) The outcome level (Mahler et al., 2019) consists of learning new grasps by optimizing the grasp success with different target hands.

It is important to notice that the abstraction level has a direct influence on the capabilities that are required for successful transfer across robots, environments, and tasks. In particular, transfer at a given abstraction level requires abilities ranging from the bottom of the robot capability stack onto the abilities of the current level (see Figure 7-right). For instance, transfer at the level of verbal instructions demands robots not only to have an abstract understanding of the word, but also to be endowed with task and motion planners, a set of robot skills, and low-level controllers to successfully execute the target task on the target robot in the target environment. In this context, recent advances in foundational models are a promising research direction to endow robots with emulated high-level cognitive capabilities (Ahn et al., 2022; Bommasani et al., 2021; Driess et al., 2023). Such foundation models generate semantic plans required to execute a target task based on language and on continuous information collected by the robot (e.g., images, state vectors). Driess et al. (Driess et al., 2023) proposed to address the correspondence problem between tasks at the highest level, that is, from a semantic perspective, by combining a large language model with perceptual inputs in an embodied multimodal model. Transfer between robots, tasks, and environments is then achieved via a large amount of training data and by training the models on several robots, tasks, and environments simultaneously. In other words, the transfer comes—to some




Figure 7. Cognitive skill levels in biology, corresponding abstraction levels in robotics, and associated robotic capability stack. A higher level of abstraction eases the transfer to different agents, environments, and tasks, but requires more and more complex robot’s capabilities. Namely, transfer at a given abstraction level requires the robot to be endowed with abilities ranging from the bottom to the corresponding level of the capability stack.

extent—“for free” thanks to the large scale of foundation models. Importantly, such transfer happens only at the highest level, that is, at the level of semantic planning, while low-level policies and planners are assumed to be given. In other words, transfer is not tackled at lower levels. As a consequence, the difficulty of transfer, as well as the resulting performance, is highly dependent on the capability stack that is made available a priori for each robot. Moreover, training a (still limited) low-level capability stack from scratch, as done, for example, in Ahn et al. (2022), requires months of data collection and is not scalable in the long run.

In this sense, we contend that investigating transfer learning methods across the entire robot capability stack is of utmost importance. In particular, we believe that bridging the gap between high-level semantic task transfer (Driess et al., 2023) and low-level execution of various tasks with different robots in the real world is a crucial challenge for transfer learning in robotics. These require *grounding* the aforementioned transferable high-level representations into the real world via robot sensorimotor experience. Such grounded understanding of the world may enable imitation and emulation learning to be intrinsically linked to the robot’s physical capabilities, thus facilitating the inference of what can be transferred, at which level, and in which situation. Previous works aiming at grounding language in robot sensorimotor behaviors (see, e.g., Cangelosi (2010); Krueger et al. (2011)) may serve as a starting point to tackle this problem. An important challenge is to design grounded representations that allow the expansion of the robot capability stacks at all levels based on similarities between tasks, environments, and robots, thus avoiding cumbersome training of medium- and low-level abilities in novel settings. In addition, designing shared grounded representations as proposed in Krueger et al. (2011); Montesano et al. (2008) is crucial for transfer across different abstraction levels.

4.2. Robotics transformers

As previously mentioned, the use of large pre-trained foundational models (Bommasani et al., 2021) to learn to transfer is enticing. Several large transformer-based models have been adapted for use in robotics, resulting into so-called *Robotic Transformers*. These models take images and natural language instructions as input and aim to output direct robot actions in the form of Cartesian trajectories. Robotics Transformers were popularized by RT-1 (Brohan et al., 2023b), in which both the input sequence of images and the natural language instructions were tokenized, that is, broken down into individual units—words or subwords for language and patches for images—called tokens. RT-1 essentially consists of a combination of existing architectures. Namely, the natural Language instructions are first embedded using the universal sentence encoder (Cer et al., 2018) and passed into a FiLM layer (Perez et al., 2018), which then constitutes the first layer of EfficientNet-B3 (Tan and Le, 2019), thus allowing the fusion of images and

language instructions into tokens. To achieve a closed loop action generation at 3 Hz, the number of tokens is reduced with the TokenLearner (Ryoo et al., 2021). The obtained sequence of tokens, corresponding to the sequence of images, is then finally fed into the transformer architecture (Vaswani et al., 2017), which outputs the action consisting 11 discrete variables of 256 bins (7 variables for the arm and gripper movement, three variables for moving the base, and one variable that switches between controlling the arm, the base, or terminating the episode). The model is trained with a large dataset of approximately 130,000 episodes performing over 700 tasks collected in the real world. Despite the incorporation of semantic reasoning, as well as the considerable amount of training data and model parameters, RT-1 generalization is limited to the combination of seen concepts. Moreover, it is limited to simple robotic tasks, cannot, for example, generate compliant motions or solve complex and dexterous manipulation tasks, and cannot outperform the task demonstrator.

The subsequent RT-2 (Brohan et al., 2023a) is a vision-language-action model based on vision-language models (Chen et al., 2023b; Driess et al., 2023) trained on web-scale data and tuned with robotic actions. The largest RT-2 consists of 55 billions parameters. The increased performance of RT-2 compared to RT-1 and other adjusted baseline models (such as VC-1 (Majumdar et al., 2023), R3M (Nair et al., 2022), MOO (Stone et al., 2023)) is attributed to the vision-language backbone combining co-finetuning the pre-trained model jointly on robotics and web data, so that the model considers more abstract visual concepts as well as robot actions. Interestingly, the largest RT-2 model displays encouraging emergent capabilities, where the model is able to use the high-level concepts acquired from the web-scale data such as relative relations between objects to complete tasks that were not present in that form in the robotic dataset. However, these emerging capabilities only emerged in the largest models, which necessitate a complex cloud infrastructure to be deployed. Therefore, they are currently unsuitable for deployment on robotics platforms and self-sufficient autonomous systems. Moreover, the model is not able to produce motions that are not covered by the large robotics dataset. Furthermore, the size of the model (55 billions parameters) can slow the model inference down to 1 Hz.

Overall, robotics transformers incorporated high-level semantic reasoning capabilities directly into the robotic actions. This is equivalent to fusing the capability stack of Figure 7 into (a) single monolith model. This approach comes at the price of reduced low-level performance and limited capabilities compared to traditional methods that output continuous actions, or direct force control for compliant tasks, among others. In addition, fusing the capability stack also results in big and cumbersome models, which are difficult to deploy. In this sense, decomposing the capability stack may result in scalable models resulting in higher performances in robotics tasks.

Importantly, robotics transformers still offer potential beyond their usage as an end-to-end controller. Such models can potentially be used in a similar fashion as large pre-trained models have been used in machine learning to obtain compact or universal representations. In this context, it is worth highlighting RT-X (Padalkar et al., 2023) the latest iteration of the robotic transformers, which aggregated 60 robotic datasets with 22 different manipulator embodiments and made this data suitable for the robotic transformer architecture. Such open-source tools and datasets are crucial to bootstrap research in transfer learning for robotics.

4.3. Universal representations

Pre-trained representations are widely popular both in machine learning and in robotics (Pari et al., 2022). For instance, ImageNet (Deng et al., 2009) was often used to acquire low-dimensional image representation for picking via suction and parallel gripper (Yen-Chen et al., 2020), contact-rich high-dimensional dexterous manipulation tasks (Shah and Kumar, 2021), and household tasks such as scooping involving tools (Liu et al., 2018). The main advantage of such representations is their flexibility in being leveraged for many different downstream tasks with little adaptation. Although these representations are promising, they still require fine-tuning to be transferred to different settings. In this sense, robotics would benefit from truly universal representations that would be intrinsically transferable between robots, environments, and complex tasks without additional training.

In this context, adapting methods such as *universal domain adaptation* (You et al., 2019) to robotics stands as a particularly promising research direction. Universal domain adaptation removes many assumptions regarding the relationship between source and target label dataset. After extracting features from both domains, the proposed universal adaptation network (UAN) employs (1) an adversarial discriminator to match the source and target feature distributions falling under common labels, (2) a non-adversarial discriminator to obtain the domain similarity, that is, quantify the similarity of an input with the source domain, and (3) a label classifier predicting the probability of the input over to the source classes. Given the domain similarity and the label predicted by the classifier, UAN predicts either a known source label or an unknown class label, thus enabling its use in settings where source and target labels are different. Extending the universal domain adaptation framework beyond classification tasks would be a first promising step towards universal representations for robotics. Such representations may then be directly leveraged for planning and control.

Alternatively, universal representations may be constructed by tasking models with so-called *pretext tasks*, that is, tasks designed solely to acquire representations that are then used in a plethora of downstream tasks. In unsupervised visual representation learning, the pretext task of instance discrimination (Wu et al., 2018) inspired many

representation models based on contrastive learning (Caron et al., 2020; Chen et al., 2020; He et al., 2020). Importantly, the pretext task does not require any labels. In other words, the unsupervised setting removes any assumptions on source and target labels, similarly as in universal domain adaptation. Training models unsupervisedly and jointly on source and target data may be a promising direction to obtain universal representations for transfer learning in robotics.

Moreover, the advent of large language models and visual language models brought a new breed of representation models that have been rapidly applied in all areas of robotics, for example, in planning (Huang et al., 2022; Shah et al., 2023), manipulation (Jiang et al., 2022; Khandelwal et al., 2022; Ren et al., 2023), and navigation (Gadre et al., 2022; Lin et al., 2022; Parisi et al., 2022). Such models are excellent candidates to harvest novel universal representations for robotics.

Some large visual language models jointly account for different modalities by encoding them in a shared latent space. For instance, CLIP (Radford et al., 2021) is pre-trained on a large dataset of images and associated textual descriptions. The model maps both modalities into a shared latent space using a contrastive loss function. The idea of combining representations from different modalities into a shared latent space was also explored in robotics. For instance, Tatiya et al. (2020, 2023) learned a common latent space from haptic feature space of multiple robots. Knowledge from source robots was then transferred through the latent space to facilitate object recognition by a target robot. Lee et al. (2020b) considered specific encoders for RGB, depth, force-torque, and proprioception modalities, which were aggregated into a multimodal representation with a multimodal fusion model. This shared representation was shown to improve the sample efficiency of the manipulation policy for a peg insertion task. The case of missing modalities during inference time was considered in Silva et al. (2020). A perceptual model of the world was trained by assuming that some modalities may not be available at all times. The approach can therefore compensate for missing or corrupted modalities during execution. Such joint representations are crucial to design universal representations for transfer.

To be successfully leveraged in various robotic scenarios, universal representations should be expressive, while remaining simple enough to facilitate downstream applications. This is usually achieved via a dimensionality reduction process by extracting low-dimensional latent representations from data. While this latent space was usually assumed to be Euclidean, that is, flat, recent works have shown the superiority of curved spaces—manifolds like hypersphere, hyperbolic spaces, symmetric spaces, and product of thereof—to learn representations of data exhibiting hierarchical or cyclic structures (Gu et al., 2019; López et al., 2021; Nickel and Kiela, 2017). For instance, the compositionality of visual scenes can be preserved via hyperbolic latent representations, thus improving

downstream performance in point cloud analysis (Montanaro et al., 2022) and unsupervised visual representation learning (Ge et al., 2023). This suggests that re-thinking inductive bias in the form of the geometry of universal representations may also be relevant for robotics applications and for transfer learning in robotics. For example, data associated with robotics taxonomies are better represented in hyperbolic spaces (Jaquier et al., 2024) and manipulation tasks encoded as graphs in the context of visual action planning (Lippi et al., 2023) may benefit from non-Euclidean representations.

4.4. Interpretability

Interpretability and explainability of learning-based approaches are key to safely deploy robots into the real world. In particular, black-box approaches lacking human-level interpretability can severely hinder natural and safe interactions with robots. In this context, transferable universal representations should also be interpretable and explainable. To do so, approaches in the field of visual action planning (Lippi et al., 2023; Wang et al., 2019) proposed to decode the underlying representations into a human-readable format, that is, images. Alternatively, representations can be readily encoded into a human-readable format that is additionally interpretable by many other methods or software architectures. For instance, the universal scene description (USD) (Studio, 2023) was designed to interchange 3D graphics information. This format was recently enhanced by Nvidia to facilitate large, complex digital twins — reflections of the real world that can be coupled to physical robots and synchronized in real time (Nvidia, 2023). USD is made from sets of data structures and APIs, which are then used to represent and modify virtual environments on supported frameworks such as Omniverse (Mittal et al., 2023), Maya (Autodesk, INC, 2019), and Houdini (SideFX, 2022). Such a framework has significant potential to be used for robotics transferability. For instance, it could be leveraged to build joint representations of the world shared across multiple robots, to share knowledge, and even to infer digital twins from sensory readings.

4.5. Benchmarking and simulation

Benchmarks and relevant metrics are key to evaluate and compare methods, thus having the potential to boost the development of innovative novel approaches. For instance, the rapid improvement of deep-learning models benefited from easily-accessible benchmarks that are widely accepted by the community (Cordts et al., 2016; Deng et al., 2009; Krizhevsky, 2009; Lin et al., 2014). The robotics community also benefited from impressive strides towards unified benchmarks with efforts such as the YCB-(Calli et al., 2015) and KIT-(Kasper et al., 2012) object datasets, and with regularly-organized benchmark competitions such as RoboCup (Kitano et al., 1997), ANA Avatar Xprice⁴, and

DARPA challenges⁵. However, they all face robotics unique challenges. First, as robots are real systems evolving in the real world, the deployment of any method can be highly time-consuming. Second, as previously mentioned, transferring methods to robots with different embodiments is non-trivial, which intrinsically hinders benchmarking across different research groups. Last but not least, handcrafted, highly tuned solutions usually outperform more general methods to solve any specific or standardized task as defined in classical benchmarks. This is especially notable for robotic manipulation where accepted benchmarks remain scarce.

The Robothon 2023 task board challenge (So et al., 2022) is an example of recent robotics manipulation benchmark. This board is an assembly of various relevant robotics tasks—including inserting a key into a keyhole and turning it, plugging/unplugging an ethernet connector, and pushing switches, among others—allowing the evaluation of different approaches. As required for a benchmark, the task board is standardized and its specifications are given. However, in such settings, handcrafted, or even pre-recorded, motions can lead to surprisingly high scores. Randomly orienting the board before every trial was later included to discourage such solutions. Within machine learning benchmarks, handcrafted solutions are prevented by dividing the available data into training and test sets, which consists of different samples drawn from a single distribution. Analogously, the Robothon 2023 challenge would require a large number of task boards consisting of the same high-level tasks but differing in their geometric-specific realization. A promising avenue to overcome the impracticability of producing numerous physical task boards would be to leverage simulators.

Modern robotic simulators, such as Mujoco (Todorov et al., 2012), Bullet (Coumans and Bai, 2016–2021), and PhysX⁶ have shown impressive improvements in various areas, including in robotics assembly (Narang et al., 2022). Such simulators have the potential to generate various parametrizations of simulated boards, and thus to create training and test sets similar to machine learning benchmarks. In particular, such sets would be of high relevance for transferability in robotics, as they have the potential to evaluate transferability across (1) robots, (2) environments, that is, different parametrizations, and (3) tasks performed on the same board. The ultimate challenge is to overcome the sim-to-real gap when deploying the developed methods on a real, previously unknown task board using a new robot during live competition. Such benchmarks would boost research in transferability in robotics, as well as provide valuable information on the difficulty and challenges of each transfer setting. It is worth highlighting that a wide range of works and methods have been developed in the field of machine learning in recent years and subsequently compiled into transfer-learning-libraries (Jiang et al., 2020). Such a consortium of methods offers a huge potential to be used in robotics contexts. Importantly, relevant metrics must

be defined to compare different approaches in different transfer settings.

4.6. Metrics for transfer learning in robotics

Two types of metrics are relevant for transfer learning in robotics, namely (1) metrics measuring the *quality of the transfer*, and (2) metrics measuring the *transfer gap*. Metrics measuring the transfer quality aim at quantifying algorithmic performance and allows the comparison of the performances of different algorithms on the target space after transfer. In particular, they can also determine *when* transfer learning is useful. Metrics measuring the transfer gap measure the discrepancy between the source and target spaces. Essentially, they provide a notion of how different the source and target robots, environments, and tasks are.

Transfer quality metrics are crucial to evaluate and compare the performance of transfer learning algorithms. As such, they are key to the development of novel transfer learning methods. Transfer quality in machine learning is commonly measured by comparing the performance achieved on the target task with or without transfer learning. In this case, the quality metric includes not only the final performance (asymptotic performance), but also the initial benefit of the transfer (jumpstart performance), the time to reach a predefined performance threshold (time to threshold), as well as the sensitivity to different hyperparameter settings (Taylor et al., 2007; Taylor and Stone, 2009). In addition to measuring the transfer quality, further analysis of the transfer process can be conducted by comparing the number of required sub-source tasks, the number of demonstrations (Barreto et al., 2018; Zhu et al., 2020b), or the required quality (e.g., suboptimal, expert, oracle) of the source space (Zhu et al., 2020a). Recently, Chen et al. (2023a) highlighted the importance of robust unsupervised evaluation metrics for domain adaptation. Such metrics should be independent of the training method, consistent across hyperparameters and models, and robust to adversarial attacks. Most of the aforementioned transfer quality metrics can and are in fact already used for transfer learning in robotics (Zhu et al., 2023).

Transfer gap metrics provide a measure of the discrepancy between the source and target spaces. Note that such metrics may also be used to measure performance in certain circumstances. Robotics adds an additional challenge to the problem of defining suitable transfer gap metrics for transfer learning: Indeed, transfer learning in robotics can be seen as a three-part transfer problem consisting of transfer across robots, tasks, and environments. Several metrics have been defined and directly optimized to solve each of these sub-problems. In this context, domain adaptation received considerable attention from the machine learning community in recent years. When the distribution of the source and target domains can be reliably estimated, simple divergences, for example, the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951) the Maximum Mean

Discrepancy (MMD) (Gretton et al., 2012) for labeled data, or the $H\Delta H$ divergence (Ben-David et al., 2010) MDD (Zhang et al., 2019), and SND (Saito et al., 2021) for unlabeled data, provide a quantitative estimate of the domain transfer gap. In robotics, a large body of work focuses on the sim-to-real gap—or in other words, the reality gap—as a specific domain gap G_D . The sim-to-real gap is often measured as the capacity of a *realistic* simulator to emulate the real world. Collins et al. (Collins et al., 2019) quantified the reality gap by comparing simulated robot trajectories, for example, using Pybullet (Coumans and Bai, 2016) or Mujoco (Todorov et al., 2012), with real-world trajectories captured by a motion capturing system. Importantly, the simulators accurately model kinematics, but generally struggle with dynamics of robots interacting with objects. Zhang et al. (Zhang et al., 2020) specifically focused on the sim-to-real gap in robotics and predicted the transfer performance of reinforcement learning policies using a probabilistic dynamics model. Limited attention was devoted to designing transfer gap metrics for transfer learning across tasks or robots. In particular, the existing literature related to skill transfer learning in human-robot cooperation (Liu et al., 2020) does not agree on a specific skill transfer metric to measure the task transfer gap G_T . Although various metrics have also been proposed in the context of motion retargeting, see, for example, Gielniak et al. (2013); Penco et al. (2018), quantifying the quality of retargeted motions, as well as robot transfer gap G_R , generally remain open questions. We contend that a suitable transfer gap metric G for robotics should consider all three settings of transfer. For instance, this metric may be defined as a simple combination of individual metrics for robot, environment, and task transfer, for example,

$$G = \lambda_1 G_R + \lambda_2 G_D + \lambda_3 G_T,$$

where $\lambda_1, \lambda_2, \lambda_3$ are weights adjusting the individual metric influence. Such transfer metric has the potential to bootstrap the development of different transfer learning methods for robotics, as it gives insights about the discrepancy of each mode (i.e., robot, environment, and tasks) from source to target space.

4.7. Negative transfer

Importantly, transfer learning is not necessarily beneficial in all settings. Transfer learning algorithms build on systematic similarities between source and target spaces. However, if non-existing similarities are selected by the algorithm, the transfer can have a negative impact on the performance in the target space (Wang, 2021). This phenomena is denoted *negative transfer* (Rosenstein et al., 2005). Negative transfer has notably been studied within the field of meta-learning (Thrun and Pratt, 1998), in which a rapid adaptation to the novel task is assumed to be key for the success of the corresponding models. Preliminary work (Deleu and Bengio, 2018) showed that adaptation using meta-learning

algorithms, such as model-agnostic meta-learning (MAML) (Finn et al., 2017), can significantly reduce the performance on meta-training tasks.

In robotics, negative transfer may occur at the different levels of the robot capability stack. For instance, at the low control level, transferring an inverse dynamic model learned for a source quadrotor to a target quadrotor with significantly different physical properties has been shown to lead to worse performances than using a baseline controller that disregards the inverse dynamics (Sorocky et al., 2020, 2021). Interestingly, the lower levels of the capability stack may be more susceptible to negative transfer as low-level information, for example, inverse dynamic models, may only be transferred across closely-related source and target spaces. In contrast, experience at the higher levels is more general and may be transferred across a larger range of source and target spaces.

We believe that negative transfer remains an under-investigated direction in robotics. In particular, negative transfer may be particularly harmful for robotics. First, negative transfer may lead to potentially-damaging behaviors of the target robot, while safety is a crucial aspect when deploying robots in the real world. Second, negative transfer may lead to transfer learning requiring longer training time than directly learning the desired behavior in the target space, while low training time is crucial for real robots acting in the real world. Therefore, the effects and causes of negative learning remain to be thoroughly studied, as they may be key to develop successful and reliable transfer learning algorithms tailored to robotics.

5. Conclusion: The future of transfer learning in robotics

The rise of transfer learning implies its potential to enable robots to leverage available knowledge to learn and master novel situations efficiently. In this paper, we aimed at unifying the concept of transfer learning in robotics via a novel taxonomy acting as a bedrock for future developments in the field. Building on the successes of transfer learning in robotics, we outlined relevant challenges that have to be solved to realize its full potential. It is important to highlight that these challenges intrinsically relate to determining what can and cannot be transferred. As illustrated in this paper, transfer learning in machine learning largely relies on identifying whether the distribution of data across two domains remain similar. Identifying similarities and differences across robotic tasks amounts to more than comparing distributions. We have emphasized the need to delineate *similarities across tasks, environments, and robots* in an effort to ease identification of commonalities and differences in each case. Automatically identify similarities across two situations in robotics relies importantly on spelling out how much prior knowledge on the physics of the robot and environment is provided. With advances in the

use of foundational models, the amount of prior knowledge readily available may ease this transfer.

We hope that this position paper paves the way towards successful transfer learning between robots, tasks, and environments, as well as their compositions. Reusing knowledge holds the promise of closing the performance gap between humans and robots in overcoming novel challenges and acquiring new skills and concepts.

Declaration of conflicting interests

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

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the European Union’s Horizon Europe Framework Programme under grant agreement No. 101070596 (euROBIN).

ORCID iD

Noémie Jaquier  <https://orcid.org/0000-0003-3565-9414>

Notes

1. The considered keywords are: domain adaptation, domain-adaptation, meta-learning, meta learning, sim2real, sim to real, sim-to-real, knowledge transfer, transfer of knowledge, skill transfer, skills transfer, transfer of skills, task transfer, tasks transfer, transfer of tasks, motion retargeting, motion-retargeting, embodiment transfer, transfer embodiments, morphology transfer, transfer morphology, kinematic transfer, transfer kinematic, robot transfer, robot to robot, imitation learning, learning by imitation, imitation-learning, learning from demonstrations, learning from demonstration, learning by demonstrations, learning by demonstration, behavioral cloning.
2. Note that the source \mathcal{R}_S may be a human instead of a robot. This typically occurs, for example, in imitation learning.
3. Note that demonstrations and reproductions performed with the same robot avoid the need for robot transfer. This includes demonstrations acquired via kinesthetic teaching or teleoperation.
4. <https://www.xprize.org/prizes/avatar>
5. <https://www.darpa.mil/program/darpa-robotics-challenge>
6. <https://developer.nvidia.com/physx-sdk>

References

- Aberman K, Li PU, Lischinski D, et al. (2020) Skeleton-aware networks for deep motion retargeting. *ACM Transactions on Graphics* 39(4): 1–62. doi: [10.1145/3386569.3392462](https://doi.org/10.1145/3386569.3392462).
- Ada SE, Ugur E and Akin HL (2022) Generalization in transfer learning: robust control of robot locomotion. *Robotica* 40(11): 3811–3836. DOI: [10.1017/S0263574722000625](https://doi.org/10.1017/S0263574722000625).
- Ahn M, Brohan A, Brown N, et al. (2022) Do as I can, not as I say: grounding language in robotic affordances. <https://arxiv.org/abs/2204.01691>
- Asfour T, Regenstein K, Azad P, et al. (2006) ARMAR-III: an integrated humanoid platform for sensory-motor control. In:

- IEEE/RAS Intl. Conf. On Humanoid Robots (Humanoids)*, pp. 169–175.
- Autodesk, INC (2019) Maya.
- Barnett SM and Ceci SJ (2002) When and where do we apply what we learn? A taxonomy for far transfer. *Psychological Bulletin* 128(4): 612–637. DOI: [10.1037/0033-2909.128.4.612](https://doi.org/10.1037/0033-2909.128.4.612).
- Barreto A, Borsa D, Quan J, et al. (2018) Transfer in deep reinforcement learning using successor features and generalised policy improvement. In: *International Conference on Machine Learning*. PMLR, pp. 501–510.
- Bechtle S, Molchanov A, Chebotar Y, et al. (2021) Meta learning via learned loss. In: *Intl Conf. On Pattern Recognition (ICPR)*, pp. 4161–4168. DOI: [10.1109/ICPR48806.2021.9412010](https://doi.org/10.1109/ICPR48806.2021.9412010).
- Bechtle S, Righetti L and Meier F (2022) Meta learning via learned loss. <https://arxiv.org/abs/2204.02210>
- Bellicoso CD, Krämer K, Stäuble M, et al. (2019) Alma - articulated locomotion and manipulation for a torque-controllable robot. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, pp. 8477–8483. DOI: [10.1109/ICRA.2019.8794273](https://doi.org/10.1109/ICRA.2019.8794273).
- Ben-David S, Blitzer J, Crammer K, et al. (2010) A theory of learning from different domains. *Machine Learning* 79: 151–175. DOI: [10.1007/s10994-009-5152-4](https://doi.org/10.1007/s10994-009-5152-4).
- Bharadhwaj H, Wang Z, Bengio Y, et al. (2019) A data-efficient framework for training and sim-to-real transfer of navigation policies. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 782–788.
- Billard A, Calinon S, Dillmann R, et al. (2008) *Robot Programming by Demonstration*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1371–1394. DOI: [10.1007/978-3-540-30301-5_60](https://doi.org/10.1007/978-3-540-30301-5_60).
- Blitzer J, McDonald R and Pereira F (2006) Domain adaptation with structural correspondence learning. In: *Proc. Of the Conference on Empirical Methods in Natural Language Processing*, 120–128. <https://aclanthology.org/W06-1615>
- Bommasani R, Hudson DA, Adeli E, et al. (2021) On the opportunities and risks of foundation models. <https://arxiv.org/abs/2108.07258>.
- Bonilla EV, Chai K and Williams C (2007) Multi-task Gaussian process prediction. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 20. https://papers.nips.cc/paper_files/paper/2007/hash/66368270ffd51418ec58bd793f2d9b1b-Abstract.html
- Bousmalis K, Irpan A, Wohlhart P, et al. (2018) Using simulation and domain adaptation to improve efficiency of deep robotic grasping. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, pp. 4243–4250. DOI: [10.1109/ICRA.2018.8460875](https://doi.org/10.1109/ICRA.2018.8460875).
- Bouyarmane K, Chappellet K, Vaillant J, et al. (2018) Quadratic programming for multirobot and task-space force control. *IEEE Transactions on Robotics* 35(1): 64–77.
- Bouzit M (1996) *Design, implementation and testing of a data glove with force feedback for virtual and real objects tele-manipulation, PhD Thesis*. Laboratoire de Robotique de Paris, University of Pierre Et Marie Curie.
- Brohan A, Brown N, Carbajal J, et al. (2023a) RT-2: vision-language-action models transfer web knowledge to robotic control. <https://arxiv.org/abs/2307.15818>.
- Brohan A, Brown N, Carbajal J, et al. (2023b) RT-1: robotics transformer for real-world control at scale. In: *Robotics: Science and Systems (R:SS)*. <https://www.roboticsproceedings.org/rss19/p025.pdf>
- Calinon S (2018) *Robot Learning with Task-Parameterized Generative Models*. Cham: Springer International Publishing, 111–126. DOI: [10.1007/978-3-319-60916-4_7](https://doi.org/10.1007/978-3-319-60916-4_7).
- Calli B, Singh A, Walsman A, et al. (2015) The ycb object and model set: towards common benchmarks for manipulation research. In: *2015 International Conference on Advanced Robotics (ICAR)*. IEEE, 510–517. DOI: [10.1109/ICAR.2015.7251504](https://doi.org/10.1109/ICAR.2015.7251504).
- Cangelosi A (2010) Grounding language in action and perception: from cognitive agents to humanoid robots. *Physics of Life Reviews* 7(2): 139–151. DOI: [10.1016/j.plrev.2010.02.001](https://doi.org/10.1016/j.plrev.2010.02.001).
- Caron M, Misra I, Mairal J, et al. (2020) Unsupervised learning of visual features by contrasting cluster assignments. *Neural Information Processing Systems (NeurIPS)* 33: 9912–9924. https://proceedings.neurips.cc/paper_files/paper/2020/hash/70feb62b69f16e0238f741fab228fec2-Abstract.html
- Cer D, Yang Y, Kong S, et al. (2018) Universal sentence encoder. <https://arxiv.org/abs/1803.11175>.
- Chen T, Murali A and Gupta A (2018) Hardware conditioned policies for multi-robot transfer learning. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 31. https://proceedings.neurips.cc/paper_files/paper/2018/hash/b8cfb77a3d250a4523ba67a65a7d031-Abstract.html
- Chen T, Kornblith S, Norouzi M, et al. (2020) A simple framework for contrastive learning of visual representations. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 119: 1597–1607. <https://proceedings.mlr.press/v119/chen20j.html>
- Chen M, Gao Z, Zhao S, et al. (2023a) A study of unsupervised evaluation metrics for practical and automatic domain adaptation. <https://arxiv.org/abs/2308.00287>
- Chen X, Djolonga J, Padlewski P, et al. (2023b) PaLI-X: on scaling up a multilingual vision and language model. <https://arxiv.org/abs/2305.18565>.
- Collins J, Howard D and Leitner J (2019) Quantifying the reality gap in robotic manipulation tasks. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, pp. 6706–6712. DOI: [10.1109/ICRA.2019.8793591](https://doi.org/10.1109/ICRA.2019.8793591).
- Cordts M, Omran M, Ramos S, et al. (2016) The cityscapes dataset for semantic urban scene understanding. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 3213–3223. DOI: [10.1109/CVPR.2016.350](https://doi.org/10.1109/CVPR.2016.350).
- Coumans E and Bai Y (2016) Pybullet, a python module for physics simulation for games. <https://pybullet.org/>
- Dariussh B, Gienger M, Arumbakkam A, et al. (2008) Online and markerless motion retargeting with kinematic constraints. In: *IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 191–198. DOI: [10.1109/IROS.2008.4651104](https://doi.org/10.1109/IROS.2008.4651104).

- Dautenhahn K and Nehaniv CL (eds) (2002) *Imitation in Animals and Artifacts*. Cambridge, MA, USA: MIT Press. <https://mitpress.mit.edu/9780262527750/imitation-in-animals-and-artifacts/>
- Deleu T and Bengio Y (2018) The effects of negative adaptation in model-agnostic meta-learning. <https://arxiv.org/abs/1812.02159>.
- Deng J, Dong W, Socher R, et al. (2009) Imagenet: a large-scale hierarchical image database. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 248–255. DOI: [10.1109/CVPR.2009.5206848](https://doi.org/10.1109/CVPR.2009.5206848).
- Deniša M, Gams A, Ude A, et al. (2016) Learning compliant movement primitives through demonstration and statistical generalization. *IEEE/ASME Trans. on Mechatronics* 21(5): 2581–2594. DOI: [10.1109/TMECH.2015.2510165](https://doi.org/10.1109/TMECH.2015.2510165).
- Driess D, Xia F, Sajjadi MSM, et al. (2023) PaLM-E: an embodied multimodal language model. *arXiv preprint 2303.03378* URL <https://arxiv.org/abs/2303.03378>
- Egli P and Hutter M (2022) A general approach for the automation of hydraulic excavator arms using reinforcement learning. *IEEE Robotics and Automation Letters* 7(2): 5679–5686. DOI: [10.1109/LRA.2022.3152865](https://doi.org/10.1109/LRA.2022.3152865).
- Evgeniou T and Pontil M (2004) Regularized multi-task learning. In: *Proc. Of the ACM SIGKDD Intl Conf. on Knowledge Discovery and Data Mining*, 109–117. DOI: [10.1145/1014052.1014067](https://doi.org/10.1145/1014052.1014067).
- Fernández F, García J and Veloso M (2010) Probabilistic policy reuse for inter-task transfer learning. *Robotics and Autonomous Systems* 58(7): 866–871. DOI: [10.1016/j.robot.2010.03.007](https://doi.org/10.1016/j.robot.2010.03.007).
- Finn C, Abbeel P and Levine S (2017) Model-agnostic meta-learning for fast adaptation of deep networks. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 70: 1126–1135. <https://proceedings.mlr.press/v70/finn17a.html>
- Gadre SY, Wortsman M, Ilharco G, et al. (2022) Clip on wheels: zero-shot object navigation as object localization and exploration. <https://arxiv.org/abs/2203.10421>.
- Ge S, Mishra S, Kornblith S, et al. (2023) Hyperbolic contrastive learning for visual representations beyond objects. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 6840–6849. DOI: [10.1109/CVPR52729.2023.00661](https://doi.org/10.1109/CVPR52729.2023.00661).
- Geng C, Huang SJ and Chen S (2021) Recent advances in open set recognition: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43(10): 3614–3631. DOI: [10.1109/TPAMI.2020.2981604](https://doi.org/10.1109/TPAMI.2020.2981604).
- Genschow O, van Den Bossche S, Cracco E, et al. (2017) Mimicry and automatic imitation are not correlated. *PLoS One* 12(9): e0183784.
- Ghadirzadeh A, Chen X, Poklukar P, et al. (2021) Bayesian meta-learning for few-shot policy adaptation across robotic platforms. In: *IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 1274–1280. DOI: [10.1109/IROS51168.2021.9636628](https://doi.org/10.1109/IROS51168.2021.9636628).
- Gielen MJ, Liu CK and Thomaz AL (2013) Generating human-like motion for robots. *Intl. Journal of Robotics Research* 32(11): 1275–1301. DOI: [10.1177/0278364913490533](https://doi.org/10.1177/0278364913490533).
- Gretton A, Borgwardt KM, Rasch MJ, et al. (2012) A kernel two-sample test. *Journal of Machine Learning Research* 13(1): 723–773. <https://jmlr.csail.mit.edu/papers/v13/gretton12a.html>
- Gu A, Sala F, Gunel B, et al. (2019) Learning mixed-curvature representations in products of model spaces. In: *Intl. Conf. On Learning Representations (ICLR)*. <https://openreview.net/pdf?id=HJxeWnCcF7>.
- He K, Fan H, Wu Y, et al. (2020) Momentum contrast for unsupervised visual representation learning. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 9729–9738. DOI: [10.1109/CVPR42600.2020.00975](https://doi.org/10.1109/CVPR42600.2020.00975).
- Heyes C (2001) Causes and consequences of imitation. *Trends in Cognitive Sciences* 5(6): 253–261. DOI: [10.1016/s1364-6613\(00\)01661-2](https://doi.org/10.1016/s1364-6613(00)01661-2).
- Höfer S, Bekris K, Handa A, et al. (2021) Sim2real in robotics and automation: applications and challenges. *IEEE Transactions on Automation Science and Engineering* 18(2): 398–400. DOI: [10.1109/TASE.2021.3064065](https://doi.org/10.1109/TASE.2021.3064065).
- Huang J, Gretton A, Borgwardt K, et al. (2006) Correcting sample selection bias by unlabeled data. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 19. https://papers.nips.cc/paper_files/paper/2006/hash/a2186aa7c086b46ad4e8bf81e2a3a19b-Abstract.html
- Huang W, Abbeel P, Pathak D, et al. (2022) Language models as zero-shot planners: extracting actionable knowledge for embodied agents. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 162: 9118–9147. <https://proceedings.mlr.press/v162/huang22a.html>
- Ijspeert A, Nakanishi J and Schaal S (2002) Movement imitation with nonlinear dynamical systems in humanoid robots. In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, volume 2. pp. 1398–1403. DOI: [10.1109/ROBOT.2002.1014739](https://doi.org/10.1109/ROBOT.2002.1014739).
- Ijspeert AJ, Nakanishi J, Hoffmann H, et al. (2013) Dynamical movement primitives: learning attractor models for motor behaviors. *Neural Computation* 25(2): 328–373. DOI: [10.1162/NECO_a_00393](https://doi.org/10.1162/NECO_a_00393).
- Jaquier N, Rozo L and Calinon S (2020) Analysis and transfer of human movement manipulability in industry-like activities. In: *IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 11131–11138. DOI: [10.1109/IROS45743.2020.9341353](https://doi.org/10.1109/IROS45743.2020.9341353).
- Jaquier N, Rozo L, González-Duque M, et al. (2024) Bringing robotics taxonomies to continuous domains via GPLVM on hyperbolic manifolds. In: *Intl. Conf. On Machine Learning (ICML)*. <https://openreview.net/forum?id=ndVXXmxSC5>
- Jiang J, Chen B, Fu B, et al. (2020) Transfer-learning-library. <https://github.com/thuml/Transfer-Learning-Library>.
- Jiang Y, Gupta A, Zhang Z, et al. (2022) Vima: general robot manipulation with multimodal prompts. <https://arxiv.org/abs/2210.03094>
- Kasper A, Xue Z and Dillmann R (2012) The kit object models database: an object model database for object recognition, localization and manipulation in service robotics. *The International Journal of Robotics Research* 31(8): 927–934. DOI: [10.1177/0278364912445831](https://doi.org/10.1177/0278364912445831).

- Khadivar F, Chatzilygeroudis K and Billard A (2023) Self-correcting quadratic programming-based robot control. In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems*.
- Khandelwal A, Weihs L, Mottaghi R, et al. (2022) Simple but effective: clip embeddings for embodied AI. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 14829–14838. DOI: [10.1109/CVPR52688.2022.01441](https://doi.org/10.1109/CVPR52688.2022.01441).
- Khansari-Zadeh SM and Billard A (2011) Learning stable non-linear dynamical systems with Gaussian mixture models. *IEEE Transactions on Robotics* 27(5): 943–957. DOI: [10.1109/TRO.2011.2159412](https://doi.org/10.1109/TRO.2011.2159412).
- Kitano H, Asada M, Kuniyoshi Y, et al. (1997) Robocup: the robot world cup initiative. In: *Proceedings of the First International Conference on Autonomous Agents*, 340–347. DOI: [10.1145/267658.267738](https://doi.org/10.1145/267658.267738).
- Klein H, Jaquier N, Meixner A, et al. (2022) A riemannian take on human motion analysis and retargeting. In: *2022 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*, 5210–5217. DOI: [10.1109/IROS47612.2022.9982127](https://doi.org/10.1109/IROS47612.2022.9982127).
- Kramberger A, Gams A, Nemeč B, et al. (2016) Transfer of contact skills to new environmental conditions. In: *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, 668–675. DOI: [10.1109/HUMANOIDS.2016.7803346](https://doi.org/10.1109/HUMANOIDS.2016.7803346).
- Kremelberg D (2019) Embodiment as a necessary a priori of general intelligence. In: P Hammer, P Agrawal, B Goertzel, et al. (eds) *Artificial General Intelligence*. Springer International Publishing, 132–136. DOI: [10.1007/978-3-030-27005-6_13](https://doi.org/10.1007/978-3-030-27005-6_13).
- Krizhevsky A (2009) Learning multiple layers of features from tiny images. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- Krueger N, Geib C, Piater J, et al. (2011) Object-action complexes: grounded abstractions of sensory-motor processes. *Robotics and Autonomous Systems* 59(10): 740–757. DOI: [10.1016/j.robot.2011.05.009](https://doi.org/10.1016/j.robot.2011.05.009).
- Kullback S and Leibler RA (1951) On information and sufficiency. *The Annals of Mathematical Statistics* 22(1): 79–86.
- Kumar A, Li Z, Zeng J, et al. (2022) Adapting rapid motor adaptation for bipedal robots. In: *IEEE/RSSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 1161–1168. DOI: [10.1109/IROS47612.2022.9981091](https://doi.org/10.1109/IROS47612.2022.9981091).
- Kyriakopoulos KJ, Van Riper J, Zink A, et al. (1997) Kinematic analysis and position/force control of the anthropot dextrous hand. *IEEE Trans. on Systems, Man, and Cybernetics, Part B (Cybernetics)* 27(1): 95–104. DOI: [10.1109/3477.552188](https://doi.org/10.1109/3477.552188).
- Lawrence ND and Platt JC (2004) Learning to learn with the informative vector machine. In: *Intl. Conf. On Machine Learning (ICML)*. <https://icml.cc/Conferences/2004/proceedings/abstracts/178.htm>
- Lee J, Hwangbo J, Wellhausen L, et al. (2020a) Learning quadrupedal locomotion over challenging terrain. *Science Robotics* 5(47): eabc5986. DOI: [10.1126/scirobotics.abc5986](https://doi.org/10.1126/scirobotics.abc5986).
- Lee MA, Zhu Y, Zachares P, et al. (2020b) Making sense of vision and touch: learning multimodal representations for contact-rich tasks. *IEEE Transactions on Robotics* 36(3): 582–596.
- Li T and Figueroa N (2023) Task generalization with stability guarantees via elastic dynamical system motion policies. In: *Conference on Robot Learning (CoRL)*. PMLR. <https://openreview.net/pdf?id=8scj3YORLq>
- Li H, Shi Y, Liu Y, et al. (2012) Cross-domain video concept detection: a joint discriminative and generative active learning approach. *Expert Systems with Applications* 39(15): 12220–12228. DOI: [10.1016/j.eswa.2012.04.054](https://doi.org/10.1016/j.eswa.2012.04.054).
- Lin TY, Maire M, Belongie S, et al. (2014) Microsoft coco: common objects in context. In: *European Conference on Computer Vision (ECCV)*. Springer International Publishing, 740–755. https://link.springer.com/chapter/10.1007/978-3-319-10602-1_48
- Lin B, Zhu Y, Chen Z, et al. (2022) ADAPT: vision-language navigation with modality-aligned action prompts. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 15396–15406. DOI: [10.1109/CVPR52688.2022.01496](https://doi.org/10.1109/CVPR52688.2022.01496).
- Lippi M, Poklukar P, Welle MC, et al. (2023) Enabling visual action planning for object manipulation through latent space roadmap. *IEEE Transactions on Robotics* 39(1): 57–75. DOI: [10.1109/TRO.2022.3188163](https://doi.org/10.1109/TRO.2022.3188163).
- Liu Y, Gupta A, Abbeel P, et al. (2018) Imitation from observation: learning to imitate behaviors from raw video via context translation. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, 1118–1125. DOI: [10.1109/ICRA.2018.8462901](https://doi.org/10.1109/ICRA.2018.8462901).
- Liu Y, Li Z, Liu H, et al. (2020) Skill transfer learning for autonomous robots and human-robot cooperation: a survey. *Robotics and Autonomous Systems* 128: 103515. DOI: [10.1016/j.robot.2020.103515](https://doi.org/10.1016/j.robot.2020.103515).
- Lončarić Z, Simonič M, Ude A, et al. (2022) Combining reinforcement learning and lazy learning for faster few-shot transfer learning. In: *IEEE/RAS Intl. Conf. On Humanoid Robots (Humanoids)*, 285–290. DOI: [10.1109/Humanoids53995.2022.10000095](https://doi.org/10.1109/Humanoids53995.2022.10000095).
- Long M, Cao Z, Wang J, et al. (2017) Learning multiple tasks with multilinear relationship networks. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 30. https://papers.nips.cc/paper_files/paper/2017/hash/03e0704b5690a2dee1861dc3ad3316c9-Abstract.html
- Longhini A, Moletta M, Reichlin A, et al. (2022) EDO-Net: learning elastic properties of deformable objects from graph dynamics. <https://arxiv.org/abs/2209.08996>.
- López F, Pozzetti B, Trettel S, et al. (2021) Symmetric spaces for graph embeddings: a Finsler-Riemannian approach. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 139: 7090–7101. <https://proceedings.mlr.press/v139/lopez21a.html>
- Maeda G, Ewerton M, Koert D, et al. (2016) Acquiring and generalizing the embodiment mapping from human observations to robot skills. *IEEE Robotics and Automation Letters* 1(2): 784–791. DOI: [10.1109/LRA.2016.2525038](https://doi.org/10.1109/LRA.2016.2525038).
- Mahler J, Matl M, Satish V, et al. (2019) Learning ambidextrous robot grasping policies. *Science Robotics* 4(26): eaau4984. DOI: [10.1126/scirobotics.aau4984](https://doi.org/10.1126/scirobotics.aau4984).

- Majumdar A, Yadav K, Arnaud S, et al. (2023) Where are we in the search for an artificial visual cortex for embodied intelligence? <https://openreview.net/pdf?id=NJtSbIWmt2T>
- Mandlekar A, Xu D, Martín-Martín R, et al. (2020) GTI: learning to generalize across long-horizon tasks from human demonstrations. In: *Robotics: Science and Systems (R:SS)*. <https://www.roboticsproceedings.org/rss16/p061.pdf>.
- Mendez JA and Eaton E (2021) Lifelong learning of compositional structures. In: *Intl. Conf. On Learning Representations (ICLR)*. <https://openreview.net/forum?id=ADWd4TJO13G>
- Mendez JA and Eaton E (2023) How to reuse and compose knowledge for a lifetime of tasks: a survey on continual learning and functional composition. In: *Transactions on Machine Learning Research* URL. <https://openreview.net/forum?id=VynY6Bk03b>
- Mittal M, Yu C, Yu Q, et al. (2023) ORBIT: a unified simulation framework for interactive robot learning environments. <https://arxiv.org/abs/2301.04195>
- Montanaro A, Valsesia D and Magli E (2022) Rethinking the compositionality of point clouds through regularization in the hyperbolic space. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 35. https://papers.nips.cc/paper_files/paper/2022/hash/da8f9fc2b555d122369f36a9684415c1-Abstract-Conference.html
- Montesano L, Lopes M, Bernardino A, et al. (2008) Learning object affordances: from sensory–motor coordination to imitation. *IEEE Transactions on Robotics* 24(1): 15–26. DOI: 10.1109/TRO.2007.914848.
- Muratore F, Ramos F, Turk G, et al. (2022) Robot learning from randomized simulations: a review. *Frontiers in Robotics and AI* 9. DOI: 10.3389/frobt.2022.799893.
- Nair S, Mitchell E, Chen K, et al. (2022) Learning language-conditioned robot behavior from offline data and crowd-sourced annotation. *Conference on Robot Learning (CoRL), Proceedings of Machine Learning Research* 164: 1303–1315. <https://proceedings.mlr.press/v164/nair22a.html>
- Narang Y, Storey K, Akinola I, et al. (2022) Factory: fast contact for robotic assembly. <https://arxiv.org/abs/2205.03532>
- Narvekar S, Peng B, Leonetti M, et al. (2020) Curriculum learning for reinforcement learning domains: a framework and survey. *Journal of Machine Learning Research* 21(181): 1–50. <https://jmlr.org/papers/volume21/20-212/20-212.pdf>
- Nichol A, Achiam J and Schulman J (2018) *On First-Order Meta-Learning Algorithms*.
- Nickel M and Kiela D (2017) Poincaré embeddings for learning hierarchical representations *Neural Information Processing Systems (NeurIPS)*, Vol. volume 30. URL https://papers.nips.cc/paper_files/paper/2017/hash/59dfa2df42d9e3d41f5b02bfc32229dd-Abstract.html.
- Nvidia (2023) Universal scene description. <https://developer.nvidia.com/usd>.
- Padalkar A, Pooley A, Jain A, et al. (2023) Open X-Embodiment: robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>.
- Pan SJ and Yang Q (2010) A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22(10): 1345–1359. DOI: 10.1109/TKDE.2009.191.
- Pan SJ, Ni X, Sun JT, et al. (2010) Cross-domain sentiment classification via spectral feature alignment. In: *Proceedings of the Intl Conf. On Worldwide Web*, 751–760. DOI: 10.1145/1772690.1772767.
- Paraschos A, Daniel C, Peters JR, et al. (2013) Probabilistic movement primitives. *Advances in Neural Information Processing Systems* 26. DOI: 10.5555/2999792.2999904.
- Pari J, Shafiullah NM, Arunachalam SP, et al. (2022) The surprising effectiveness of representation learning for visual imitation. In: *Robotics: Science and Systems (R:SS)* URL. <https://www.roboticsproceedings.org/rss18/p010.pdf>
- Parisi S, Rajeswaran A, Purushwalkam S, et al. (2022) The (un) surprising effectiveness of pre-trained vision models for control. *Intl. Conf. on Machine Learning (ICML)* 162: 17359–17371. <https://proceedings.mlr.press/v162/parisi22a.html>
- Peer A, Eickenel S and Buss M (2008) Multi-fingered telemanipulation-mapping of a human hand to a three finger gripper. In: *IEEE Intl. Symposium on Robot and Human Interactive Communication*. (RO-MAN), 465–470. DOI: 10.1109/ROMAN.2008.4600710.
- Penco L, Clement B, Modugno V, et al. (2018) Robust real-time whole-body motion retargeting from human to humanoid. In: *IEEE/RAS Intl. Conf. On Humanoid Robots (Humanoids)*, 425–432. DOI: 10.1109/HUMANOIDS.2018.8624943.
- Perez E, Strub F, De Vries H, et al. (2018) Film: visual reasoning with a general conditioning layer. *AAAI Conf. on Artificial Intelligence* volume 32: 3942–3951.
- Perkins DN and Salomon G (1992) Transfer of learning. In: T Husén and TN Postlethwaite (eds) *The International Encyclopedia of Education*, 47–79.
- Radford A, Kim JW, Hallacy C, et al. (2021) Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning*. PMLR, 8748–8763.
- Rakita D, Mutlu B and Gleicher M (2017) A motion retargeting method for effective mimicry-based teleoperation of robot arms. In: *ACM/IEEE Intl. Conf. on Human-Robot Interaction (HRI)*, 361–370. <https://ieeexplore.ieee.org/document/8534763>.
- Ravichandar H, Polydoros AS, Chernova S, et al. (2020) Recent advances in robot learning from demonstration. *Annual Review of Control, Robotics, and Autonomous Systems* 3(1): 297–330. DOI: 10.1146/annurev-control-100819-063206.
- Reader SM, Morand-Ferron J and Flynn E (2016) Animal and human innovation: novel problems and novel solutions. *Philosophical Transactions of the Royal Society B: Biological Sciences* 371(1690): 20150182. DOI: 10.1098/rstb.2015.0182.
- Ren AZ, Govil B, Yang TY, et al. (2023) Leveraging language for accelerated learning of tool manipulation. In: *Conference on Robot Learning (CoRL)*. PMLR, 1531–1541. <https://openreview.net/pdf?id=nPw7jaGBrCG>.
- Rosenstein MT, Marx Z, Kaelbling LP, et al. (2005) To transfer or not to transfer. In: *NIPS Workshop on Transfer Learning*, Vol. 898.

- Rozo L, Calinon S, Caldwell DG, et al. (2016) Learning physical collaborative robot behaviors from human demonstrations. *IEEE Transactions on Robotics* 32(3): 513–527. DOI: [10.1109/TRO.2016.2540623](https://doi.org/10.1109/TRO.2016.2540623).
- Ryoo M, Piergiovanni A, Arnab A, et al. (2021) Tokenlearner: adaptive space-time tokenization for videos. *Neural Information Processing Systems (NeurIPS)* 34: 12786–12797, <https://proceedings.neurips.cc/paper/2021/hash/6a30e32e56fce5cf381895dfe6ca7b6f-Abstract.html>
- Saenko K, Kulis B, Fritz M, et al. (2010) Adapting visual category models to new domains. In: European Conference on Computer Vision. ECCV. Springer, 213–226. DOI: [10.1007/978-3-642-15561-1_16](https://doi.org/10.1007/978-3-642-15561-1_16).
- Saito K, Kim D, Teterwak P, et al. (2021) Tune it the right way: unsupervised validation of domain adaptation via soft neighborhood density. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 9184–9193.
- Schaal S (1999) Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences* 3(6): 233–242. DOI: [10.1016/S1364-6613\(99\)01327-3](https://doi.org/10.1016/S1364-6613(99)01327-3).
- Schmidt RA and Young DE (1987) Transfer of movement control in motor skill learning. In: Transfer of Learning: Contemporary Research and Applications, 47–79. DOI: [10.1016/b978-0-12-188950-0.50009-6](https://doi.org/10.1016/b978-0-12-188950-0.50009-6).
- Shah R and Kumar V (2021) RRL: resnet as representation for reinforcement learning. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 139: 9465–9476. <https://proceedings.mlr.press/v139/shah21a.html>
- Shah D, Osinski B, Ichter B, et al. (2023) LM-nav: robotic navigation with large pre-trained models of language, vision, and action. *Conference on Robot Learning (CoRL), Proceedings of Machine Learning Research* 205: 44–54. <https://proceedings.mlr.press/v205/shah23a.html>
- Shukla Y, Thierauf C, Hosseini R, et al. (2022) Acute: automatic curriculum transfer from simple to complex environments. In: *Proc. Of the Intl Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, 1192–1200.
- Shukla Y, Kesari B, Goel S, et al. (2023) A framework for few-shot policy transfer through observation mapping and behavior cloning. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 7104–7110.
- SideFX (2022) Sidefx. <https://www.sidefx.com/>
- Silva R, Vasco M, Melo FS, et al. (2020) Playing games in the dark: an approach for cross-modality transfer in reinforcement learning. In: *Proc. Of the Intl Conf. on Autonomous Agents and MultiAgent Systems (AAMAS)*, 1260–1268.
- So P, Wittmann J, Ruhkamp P, et al. (2022) Towards remote robotic competitions: an internet-connected task board and dashboard. <https://arxiv.org/abs/2201.09565>.
- Sorocky MJ, Zhou S and Schoellig AP (2020) Experience selection using dynamics similarity for efficient multi-source transfer learning between robots. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, 2739–2745. DOI: [10.1109/ICRA40945.2020.9196744](https://doi.org/10.1109/ICRA40945.2020.9196744).
- Sorocky MJ, Zhou S and Schoellig AP (2021) To share or not to share? Performance guarantees and the asymmetric nature of cross-robot experience transfer. *IEEE Control Systems Letters* 5(3): 923–928. DOI: [10.1109/lcsys.2020.3005886](https://doi.org/10.1109/lcsys.2020.3005886).
- Stone A, Xiao T, Lu Y, et al. (2023) Open-world object manipulation using pre-trained vision-language models. In: *Conference on Robot Learning (CoRL)*. <https://openreview.net/pdf?id=9al6taqfTzr>
- Studio PU (2023) Universal scene description. <https://github.com/PixarAnimationStudios/OpenUSD>
- Tan M and Le Q (2019) Efficientnet: rethinking model scaling for convolutional neural networks. *Intl. Conf. on Machine Learning (ICML), Proceedings of Machine Learning Research* 97: 6105–6114. <https://proceedings.mlr.press/v97/tan19a.html>
- Tatiya G, Shukla Y, Edeghware M, et al. (2020) Haptic knowledge transfer between heterogeneous robots using kernel manifold alignment. In: *IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 5358–5363. DOI: [10.1109/IROS45743.2020.9340770](https://doi.org/10.1109/IROS45743.2020.9340770).
- Tatiya G, Francis J and Sinapov J (2023) Transferring implicit knowledge of non-visual object properties across heterogeneous robot morphologies. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, 11315–11321. DOI: [10.1109/ICRA48891.2023.10160811](https://doi.org/10.1109/ICRA48891.2023.10160811).
- Taylor ME and Stone P (2009) Transfer learning for reinforcement learning domains: a survey. *Journal of Machine Learning Research* 10(7).
- Taylor ME, Stone P and Liu Y (2007) Transfer learning via inter-task mappings for temporal difference learning. *Journal of Machine Learning Research* 8(9).
- Thrun S and Pratt L (1998) *Learning to Learn: Introduction and Overview*. Springer, 3–17. DOI: [10.1007/978-1-4615-5529-2_1](https://doi.org/10.1007/978-1-4615-5529-2_1).
- Todorov E, Erez T and Tassa Y (2012) Mujoco: a physics engine for model-based control. In: *IEEE/RSJ Intl. Conf. On Intelligent Robots and Systems (IROS)*, 5026–5033. DOI: [10.1109/IROS.2012.6386109](https://doi.org/10.1109/IROS.2012.6386109).
- Towers M, Terry JK, Kwiatkowski A, et al. (2023) Gymnasium. DOI: [10.5281/zenodo.8127026](https://doi.org/10.5281/zenodo.8127026).
- Ude A, Gams A, Asfour T, et al. (2010) Task-specific generalization of discrete and periodic dynamic movement primitives. *IEEE Transactions on Robotics* 26(5): 800–815. DOI: [10.1109/TRO.2010.2065430](https://doi.org/10.1109/TRO.2010.2065430).
- Vaswani A, Shazeer N, Parmar N, et al. (2017) Attention is all you need. In: *Neural Information Processing Systems (NeurIPS)*, Vol. volume 30. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- Vosylius V and Johns E (2023) Where to start? Transferring simple skills to complex environments. *Conference on Robot Learning (CoRL), Proceedings of Machine Learning Research* 205: 471–481. <https://proceedings.mlr.press/v205/vosylius23a.html>
- Walqui A (2000) *Contextual Factors in Second Language Acquisition*. <https://files.eric.ed.gov/fulltext/ED444381.pdf>
- Wang Z (2021) *Mitigating Negative Transfer for Better Generalization and Efficiency in Transfer Learning*. Carnegie Mellon University.
- Wang SJ and Johnson AM (2021) Domain adaptation using system invariant dynamics models. *Conference on Learning for Dynamics and Control (LADC), Proceedings of Machine Learning Research* 144: 1130–1141. <https://proceedings.mlr.press/v144/wang21c.html>

- Wang A, Kurutach T, Liu K, et al. (2019) Learning robotic manipulation through visual planning and acting. <https://www.roboticsproceedings.org/rss15/p74.pdf>.
- Whiten A and Ham R (1992) On the Nature and Evolution of Imitation in the Animal Kingdom: Reappraisal of a Century of Research. *Advances in The Study of Behavior* 21: 239–283.
- Whiten A, Horner V, Litchfield CA, et al. (2004) How do apes ape? *Animal Learning & Behavior* 32: 36–52. DOI: [10.3758/BF03196005](https://doi.org/10.3758/BF03196005).
- Wu Z, Xiong Y, Yu SX, et al. (2018) Unsupervised feature learning via non-parametric instance discrimination. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 3733–3742. DOI: [10.1109/CVPR.2018.00393](https://doi.org/10.1109/CVPR.2018.00393).
- Yang Q, Zhang Y, Dai W, et al. (2020) *Transfer Learning*. Cambridge: Cambridge University Press.
- Yen-Chen L, Zeng A, Song S, et al. (2020) Learning to see before learning to act: visual pre-training for manipulation. In: *IEEE Intl. Conf. On Robotics and Automation (ICRA)*, 7286–7293. DOI: [10.1109/ICRA40945.2020.9197331](https://doi.org/10.1109/ICRA40945.2020.9197331).
- Yin W, Yin H, Baraka K, et al. (2023) Multimodal dance style transfer. *Machine Vision and Applications* 34(4): 48. DOI: [10.1007/s00138-023-01399-x](https://doi.org/10.1007/s00138-023-01399-x).
- You K, Long M, Cao Z, et al. (2019) Universal domain adaptation. In: *IEEE Conf. On Computer Vision and Pattern Recognition (CVPR)*, 2715–2724. DOI: [10.1109/CVPR.2019.00283](https://doi.org/10.1109/CVPR.2019.00283).
- Zakka K, Zeng A, Florence P, et al. (2022) Xirl: cross-embodiment inverse reinforcement learning. *Conference on Robot Learning (CoRL)*, *Proceedings of Machine Learning Research* 164: 537–546. <https://proceedings.mlr.press/v164/zakka22a.html>
- Zhang Y, Liu T, Long M, et al. (2019) Bridging theory and algorithm for domain adaptation. In: *International Conference on Machine Learning*. PMLR, 7404–7413.
- Zhang LM, Plappert M and Zaremba W (2020) Predicting sim-to-real transfer with probabilistic dynamics models. <https://arxiv.org/abs/2009.12864>
- Zhao W, Queralta JP and Westerlund T (2020) Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In: *IEEE Symposium Series on Computational Intelligence*, 737–744. DOI: [10.1109/SSCI47803.2020.9308468](https://doi.org/10.1109/SSCI47803.2020.9308468).
- Zhou Y, Gao J and Asfour T (2020) Movement primitive learning and generalization: using mixture density networks. *IEEE Robotics and Automation Magazine* 27(2): 22–32. DOI: [10.1109/MRA.2020.2980591](https://doi.org/10.1109/MRA.2020.2980591).
- Zhu Z, Lin K, Dai B, et al. (2020a) Learning sparse rewarded tasks from sub-optimal demonstrations.
- Zhu Z, Lin K, Dai B, et al. (2020b) Off-policy imitation learning from observations. *Advances in Neural Information Processing Systems* 33: 12402–12413.
- Zhu Z, Lin K, Jain AK, et al. (2023) Transfer learning in deep reinforcement learning: a survey. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhuang F, Qi Z, Duan K, et al. (2020) A comprehensive survey on transfer learning. *Proceedings of the IEEE* 109(1): 43–76. DOI: [10.1109/JPROC.2020.3004555](https://doi.org/10.1109/JPROC.2020.3004555).