

A survey on imitation learning for contact-rich tasks in robotics

The International Journal of
Robotics Research
2026, Vol. 0(0) 1–36
© The Author(s) 2026



Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/02783649261417694
journals.sagepub.com/home/ijr



Toshiaki Tsuji¹ , Yasuhiro Kato² , Gokhan Solak³ ,
Heng Zhang^{3,4,5} , Tadej Petrič⁶, Francesco Nori⁷ and
Arash Ajoudani³

Abstract

This paper comprehensively surveys research trends in imitation learning (IL) for contact-rich robotic tasks. Contact-rich tasks, which require complex physical interactions with the environment, represent a central challenge in robotics due to their nonlinear dynamics and sensitivity to small positional deviations. The paper examines demonstration collection methodologies, including teaching methods and sensory modalities crucial for capturing subtle interaction dynamics. We then analyze IL approaches, highlighting their applications to contact-rich manipulation. Recent advances in multimodal learning and foundation models have significantly enhanced performance in complex contact tasks across industrial, household, and healthcare domains. Through systematic organization of current research and identification of challenges, this survey provides a foundation for future advancements in contact-rich robotic manipulation.

Keywords

contact rich tasks, imitation learning, machine learning, reinforcement learning, learning from demonstration, impedance control, manipulation

Introduction

Robots are intelligent systems that bring physical effects to real environments. Many basic tasks are contact-rich tasks involving multiple contact states between robots and their environment, and developing these capabilities is one of the core challenges in robotics. Meanwhile, understanding “everyday physics” has long been known to be extremely difficult (Arimoto, 1999). Our daily tasks involve complex interactions of diverse physical phenomena such as friction, elasticity, plastic deformation, and fracture, which exhibit nonlinear and unpredictable behaviors, making the advancement of contact-rich tasks both an old and new research challenge. This is also reflected in the publication trend in our bibliometric analysis of the reviewed references (Figure 1).

Approaches to contact-rich tasks can be broadly categorized into model-based and model-free methods, with model-based approaches being widely studied at the practical level (Xu et al., 2019). However, as research in everyday physics indicates, contact-rich tasks require highly nonlinear models since slight positional deviations can cause significant behavioral changes. As these tasks are sensitive to minor differences in models and parameters, the need for machine learning-based model-free methods increases as task difficulty rises. Nevertheless, contact-rich tasks must be performed without damaging target objects, and data collection in robotic

systems is both challenging and costly, resulting in a limited amount of training data. There are two primary approaches to extracting motion models from demonstration data for skilled tasks: reinforcement learning (RL) and imitation learning (IL) (Argall et al., 2009; Kober and Peters, 2010).

RL has the advantage of autonomously acquiring complex movements involving contact state transitions

¹Graduate School of Science and Engineering, Saitama University, Saitama, Japan

²Center for Education and Research in Information Science and Technology, University of Tokyo, Tokyo, Japan

³Human-Robot Interfaces and Interaction Lab, Istituto Italiano di Tecnologia, Genova, Italy

⁴PhD. Program of National Interest in Robotics and Intelligent Machines (DRIM) and Università di Genova, Genoa, Italy

⁵Edwardson School of Industrial Engineering, Purdue University, West Lafayette, IN, USA

⁶Department of Automatics, Biocybernetics and Robotics, Jožef Stefan Institute, Ljubljana, Slovenia

⁷Robotics lab, Google Deepmind, London, UK

All authors contributed equally to this review paper.

Corresponding author:

Toshiaki Tsuji, Graduate School of Science and Engineering, Saitama University, Saitama 338-8570, Japan.

Email: tsuji@mail.saitama-u.ac.jp




Figure 1. Publication trend. The left part of the years (1971–2009) merged as one bar with an average number. And the other right part (2010–2025) shows the annual number of publications related to IL for contact-rich tasks demonstrating a significant increase in recent years.

through interaction with the environment, and has been extensively studied. Recent trends include combining model-free and model-based RL approaches (Fan et al., 2019; Pong et al., 2018), improving sample efficiency (Shi et al., 2021; Wang et al., 2021), pre-training in simulation followed by fine-tuning on real hardware (Yang et al., 2024), and integration with adaptive impedance control for contact force regulation (Beltran-Hernandez et al., 2020; Martín-Martín et al., 2019; Oikawa et al., 2021). However, RL requires extensive trial-and-error, and learning on physical systems is often limited due to hardware wear and safety concerns. Furthermore, the complexity of contact dynamics widens the simulation-to-reality gap, making transfer more challenging. On the other hand, IL has the advantage of efficiently learning expert skills, including subtle adjustments of contact forces and positions in human dexterous manipulation. Expert skills inherently contain human tacit knowledge and empirical rules, and it is expected that if robots can acquire these in some form, their capabilities can be fundamentally enhanced.

Building upon the potential of IL to leverage human tacit knowledge and empirical rules, recent advances in large language models (LLMs) have begun to extend to other modalities such as images and speech, and in robotics, they are expected to develop into foundational technologies for symbolic representation of action sequences, integration of multimodal knowledge, and eventually technologies that can abstract human tacit knowledge and empirical rules inherent in expert skills. With the proliferation of LLMs, the number of papers on IL in robotics continues to increase. However, there is a lack of systematic organization of recent research trends in this field. Therefore, this paper aims to contribute to the advancement of robotics by conducting and systematically organizing a survey on IL for contact-rich tasks.

Related surveys and our contribution

Several surveys have addressed topics related to learning approaches in contact-rich tasks. Suomalainen et al. (2022) provide a comprehensive survey of robot manipulation in

contact, covering a broad spectrum of approaches including classical control and learning methods. Zhu and Hu (2018) survey learning from demonstration in robotic assembly. More recently, Drolet et al. (2024) compare IL algorithms specifically for bimanual manipulation, while Welte and Rayyes (2025) survey interactive IL for dexterous robotic manipulation. General surveys on IL (Argall et al., 2009; Celemin et al., 2022; Fang et al., 2019; Urain et al., 2024) have also been published.

An extensive review of robot learning for manipulation appears in Kroemer et al. (2021), providing a structured analysis of fundamental challenges, representational choices, and algorithmic frameworks. A systematic survey of robot manipulation examines complex interactions between robots and their environment during physical contact tasks (Suomalainen et al., 2022). A thorough review of robotic assembly strategies encompasses the entire operational procedure from planning to evaluation, emphasizing the importance of integrating multiple technological approaches (Jiang et al., 2022).

Surveys have also been published on IL (Argall et al., 2009; Celemin et al., 2022; Fang et al., 2019; Hua et al., 2021; Urain et al., 2024). A foundational survey on robot learning from demonstration in Argall et al. (2009) established key paradigms and methodologies that influenced the field for years. IL in robotic manipulation was specifically addressed in Fang et al. (2019), synthesizing key approaches and challenges in transferring human skills to robotic systems. The interconnections between RL, IL, and transfer learning were explored in Hua et al. (2021), highlighting their complementary nature. A comprehensive examination of interactive IL in robotics appears in Celemin et al. (2022), emphasizing the crucial role of human–robot interaction in skill acquisition. Most recently, the growing role of deep generative models in robotics was investigated in Urain et al. (2024), particularly focusing on their application in learning from multimodal demonstrations. Since dexterous manipulation requires handling the complexity of contact dynamics that is difficult to formalize mathematically, IL that learns from human empirical knowledge represents a promising approach (An et al., 2025).

While these surveys address related topics—including manipulation, assembly tasks, and IL methodologies, no comprehensive survey exists that specifically investigates the full spectrum of IL approaches for contact-rich tasks across diverse robotic systems and application domains. This survey addresses this gap by providing a systematic organization of IL methods all analyzed from the perspective of contact-rich manipulation. We place particular emphasis on recent developments in multimodal learning, foundation models, and generative approaches that have significantly advanced the field.

Although hardware technologies—including tactile sensors, force/torque sensors, and end-effector mechanisms—play a crucial role in contact-rich

manipulation, comprehensive surveys of these technologies exist elsewhere. Recent reviews cover tactile sensing technologies and their applications in robotic manipulation (Meribout et al., 2024) and design of compliant mechanisms (Samadikhoshkho et al., 2025). This survey focuses specifically on IL methodologies and algorithms, treating sensors and hardware as enabling components rather than primary subjects of investigation.

The closest paper surveyed RL technologies in contact-rich tasks (Elguea-Aguinaco et al., 2023). The significant differences in methodology and effects between RL technology and IL make this contrast valuable for organizing trends in this field.

Our work provides a comprehensive survey that:

- (1) Focuses specifically on IL for contact-rich tasks across diverse application domains (industrial, household, and healthcare), rather than limiting to specific hardware configurations (bimanual, dexterous hands) or task categories (assembly, dexterous manipulation).
- (2) Systematically organizes the full spectrum of IL methodologies including behavior cloning, dynamic movement primitives, generative methods (variational autoencoders, diffusion models), foundation models, inverse reinforcement learning, multimodal IL, and offline RL—all analyzed from the perspective of contact-rich manipulation.
- (3) Emphasizes recent developments (2019–2025) in multimodal learning, foundation models, and generative approaches that have significantly advanced the field, with particular attention to how these methods address contact-rich interaction challenges.
- (4) Provides comprehensive coverage of data modalities and collection methods (*Collecting demonstrations* section) crucial for contact-rich tasks, including tactile sensing, force feedback, and their integration with vision and language.
- (5) Addresses the broader ecosystem including demonstration collection methodologies, synthetic data generation, available datasets, and application cases across multiple domains.

Trends and insights from bibliometric analysis

Our scope and paper structure are supported by the comprehensive bibliometric analysis of the reviewed papers, which show significant trends in both temporal distribution (Figure 1), research topics (Figure 2), and application domains (Figure 3). Notably, there's a clear evolution in methodology from classical control approaches (pre-2019) to deep learning-based methods (2019–2022), and finally to foundation models (2023–2025). This trend reflects the field's progression toward more sophisticated, multimodal approaches. The analysis also highlights a shift in focus




Figure 2. Research topic temporal evolution, illustrating the shift of IL from classical control methods to deep learning and foundation models in contact-rich tasks.

from single-task solutions to more generalizable frameworks, particularly in contact-rich manipulation tasks. The increasing integration of vision, language, and tactile feedback demonstrates the field's movement toward more robust and versatile robotic systems capable of handling complex real-world scenarios. We respond to these trends by covering the generative methods, foundation models and multimodal IL extensively in the *Learning approaches* section.

Survey organization and rationale

This survey follows the natural pipeline of implementing IL for contact-rich robotic tasks (Figure 4), progressing from fundamental understanding through data collection to learning methods and practical applications.

The *Preliminaries* section establishes why contact-rich tasks are fundamentally challenging, examining contact dynamics, insights from human motor control, and key challenges when applying IL. This foundation clarifies why certain data modalities and algorithmic choices become

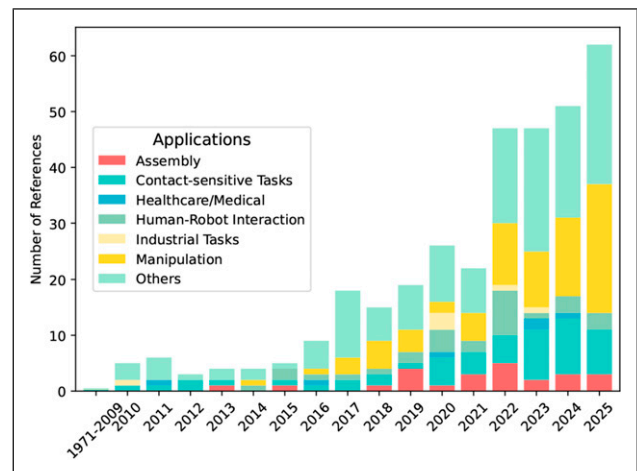


Figure 3. Application distribution over time, including industrial, household/service, and healthcare domains, showing especially a growing trend in contact-rich manipulation in recent years.

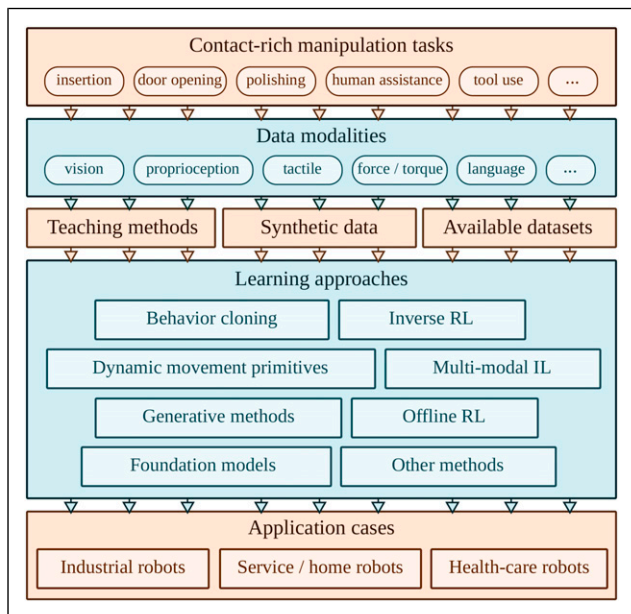


Figure 4. The sections of our survey consist of the contact-rich IL pipeline, as illustrated in five layers from problem definition to end-user application.

critical. The *Collecting demonstrations* section addresses data acquisition—the prerequisite for all IL methods. We cover data modalities (proprioception, force, vision, tactile), teaching methods (kinesthetic, teleoperation, observation), synthetic data generation, and available datasets. Data quality and modality directly constrain which learning methods are applicable. The *Learning approaches* section systematically examines IL paradigms under eight categories: *behavior cloning* (foundation), *dynamic movement primitives* (structured representations), *generative methods* (handling multimodal distributions), *foundation models* (large-scale pre-training), *inverse RL* (reward inference), *multimodal IL* (sensory integration), *offline RL* (learning from datasets), and *other methods* (emerging approaches). The *Application cases* section demonstrates practical implementations across industrial, household, and healthcare robotics, each with unique constraints. Within this section, we also provide guidance on *selection of appropriate algorithms* based on sensor and task requirements. Finally, the *Conclusion* section identifies core technical challenges (hierarchical architectures, multimodal sensing, sim-to-real transfer) and future research directions.

This structure creates a coherent arc from why contact-rich IL is challenging (*Preliminaries* section), to how to collect training data (*Collecting demonstrations* section), what learning methods are available (*Learning approaches* section), where these methods succeed in practice (*Application cases* section), and what remains to be solved (*Conclusion* section). Each section builds essential context for the next while serving as a standalone resource, with extensive cross-referencing strengthening the interconnected narrative. This organization

reflects both logical dependencies and practical workflows, making the survey pedagogically sound and practically useful.

Preliminaries

This section discusses the background of contact-rich tasks in robotics, their relationship with human motor control, and the challenges faced by robots when performing such tasks. First, the *Contact-rich robotics* section defines contact-rich tasks and briefly summarizes their inherent challenges. Next, the *Insights from motor control* section describes how modern adaptive robotic control has been inspired by human adaptive motor control and examines the relationship between the two. Finally, the *Key challenges* section discusses the challenges associated with applying imitation learning (IL) to contact-rich tasks.

Contact-rich robotics/background

In robotics, contact-rich manipulation refers to robotic manipulation tasks that involve continuous and complex interactions between the robot and its environment, often requiring sophisticated control of forces in one to several contact points (Figure 5). This is why direct and indirect force control (e.g., through impedance or admittance control) techniques have been widely exploited to address this problem. Pioneering contributions in this field include the formulation of operational space, hybrid and impedance control (Khatib, 1987; Neville, 1985; Raibert and Craig, 1981), as well as advancements in contact modeling and multi-contact control (Bicchi and Siciliano, 1993; Cutkosky and Kao, 1989; Mason, 1986; Whitney, 1982). More recent studies build upon the foundational theories proposed by these pioneering works to tackle complex contact-rich problems on advanced hardware systems (Khandelwal et al., 2023; Ozdamar et al., 2024).

A major challenge in contact-rich manipulation is the scalability of classical control solutions to varying task conditions and contact dynamics (Suomalainen et al. (2022)). For example, even a seemingly simple assembly task, such as peg-in-hole, involves varying force or impedance requirements depending on factors like the position and orientation of the parts, their material properties, and the clearance between them (Ajoudani et al., 2012). Similarly, in tasks that need continuous contact like wiping a surface or polishing, changes in surface geometry, friction, or compliance call for constant adaptation of the contact-related robot dynamics to maintain stable and effective interaction (Barreiros et al., 2025). Recent advancements in machine learning have accelerated the resolution of this problem, by enabling robots to learn intricate interaction dynamics directly from data (Zhang et al., 2024; Gubbi et al., 2020; Zhang et al., 2025d). This shift from traditional control methods to data-driven learning has yielded significant improvements in robustness and adaptability. Techniques such as RL (Elguea-Aguinaco et al., 2023), IL




Figure 5. Common examples of contact-rich robotic tasks: (a) peg insertion or assembly; (b) human assistance (e.g., dressing); (c) tool-use (e.g., vegetable cutting); (d) wiping, polishing or similar; (e) opening doors or drawers. Image generation: Google Gemini.

(Zhao et al., 2025b), and adaptive model predictive control (Lakshmipathy and Pollard, 2024) are paving the way for robots to achieve human-like dexterity in manipulating contacts. This survey provides an in-depth analysis of these advancements.

Insights from motor control

Humans are adaptively good but not inherently precise at contact-rich tasks (Mathew and Crevecoeur, 2021; Wolpert et al., 2011). They rely on compliance, predictive control, learning, and robust feedback integration to manage contact effectively (Franklin and Wolpert, 2011), but not on precise, optimal control like a model-based system would. This adaptive mastery has long made human motor coordination a powerful source of inspiration for robotics (Kober et al., 2013; Petric et al., 2018). Here, we provide a brief overview of key works focusing on human motor coordination in contact-rich tasks, with the aim of highlighting principles that can inform the design of more robust, adaptive robotic systems.

Motor control refers to the mechanisms by which organisms and robots orchestrate their movements to interact effectively with their environment, particularly in tasks involving significant physical contact. Unlike movements executed freely in space, contact-rich tasks require management of interaction forces, adaptation to varying material properties, and responsiveness to environmental perturbations. For instance, a robot tasked with screwing in a light bulb must precisely regulate the force applied to rotate it without damaging the glass.

Typical robotic controllers rely on predetermined mathematical models, which generally perform well in structured environments but frequently fail in real-world scenarios with uncertainty and variability (Kroemer et al., 2021). In contrast, biological motor control systems manage these complexities naturally. Humans and animals continuously adapt their movements based on sensory feedback, adjusting grip strength or applied force depending on object properties and environmental conditions (Franklin and Wolpert, 2011; Suomalainen et al., 2022). Additionally, biological systems utilize sensorimotor integration, dynamic impedance modulation, predictive control, and adaptive feedback mechanisms, enabling efficient interactions with complex environments (Kober et al., 2013; Petrič et al., 2017).

Roboticians have drawn on these biological principles to develop control strategies aimed at enhancing adaptability in robotic systems. Impedance and admittance control methods, for example, manage interaction forces by dynamically adjusting the robot’s stiffness and damping properties, facilitating more adaptable interactions with the environment (Cui and Trinkle, 2021; Merckaert et al., 2022). Integrating these methods with learning-based approaches, including deep learning and RL, has resulted in hybrid techniques that benefit from both model-based predictions and data-driven adaptability (Aggarwal et al., 2022).

IL methods particularly benefit from insights derived from biological motor control. Human demonstrations naturally encompass subtle adjustments in posture, force modulation, and adaptive responses to dynamic environmental conditions (Fang et al., 2019; Gams et al., 2022). By incorporating principles such as dynamic impedance adaptation and sensorimotor predictive modeling, IL approaches can replicate nuanced aspects of human motor skills effectively (Petric et al., 2018).

However, challenges persist in translating biological motor control insights to robotic systems. Accurately modeling contact dynamics remains difficult due to their inherently nonlinear and sensitive nature to small variations in physical conditions (Jiang et al., 2022). Moreover, robotic sensory systems, particularly tactile and proprioceptive sensors, remain limited compared to their biological counterparts, restricting the precision and adaptability of robotic responses (Jassim et al., 2025). Addressing these challenges requires advancements in sensor technology and computational modeling techniques.

Future research directions will focus on improving sensorimotor integration, refining predictive control models, and enhancing impedance modulation through machine learning. Progress in wearable and tactile sensing technologies and better simulation tools for realistic contact dynamics will further enable the practical implementation of biological motor control principles in robotics (Hua et al., 2021).

Key challenges

In this section we highlight the main challenges which are currently tackled in the context of contact-rich IL.

As already mentioned, one important challenge resides in the *highly nonlinear nature of contact-rich dynamics* which are extremely difficult to capture with classical modeling tools (e.g., the ordinary differential equations generated by Lagrangian mechanics) and often require computationally intractable representations (e.g., the partial differential equations generated by soft-contacts). This challenge often induced researchers to resort to learning-based model-free methods such as RL and IL. Within this context, RL is often applied in simulation and it suffers from the intractable nature of contact models and their limited modeling accuracy (i.e., sim-to-real gap). When applied in real instead, RL poses significant challenges for both the quantity and the nature of the required data: quantity-wise, it exposes hardware to extensive trials-and-errors which in contact-rich tasks induce a significant wear-and-tear; additionally, non-imitation learning requires explorations which by nature have the tendency to go beyond safety limits, which can damage the robot and its surroundings.

On the other hand, IL suffers from the *scarcity of technologies suitable for data-collection of contact-rich tasks*. Despite efforts, tactile sensors remain a technology which is mostly confined to research applications with rare industrial use-cases. Similarly, wearable tactile technologies (e.g., Büscher et al., 2012; Ruppel and Zhang, 2024; Sundaram et al., 2019; Battaglia et al., 2015) seem to be a challenging technology to develop (see Yao et al. (2026) for an in-depth analysis) and their adoption has not yet had the necessary impact either in research or in industrial applications. These technologies become even more necessary since contact-rich tasks are by nature *non-fully-observable* when the primary observation modality is vision. Interaction forces are essential in contact-rich tasks, and yet, they are not directly measurable from images: contact-rich tasks are often susceptible to visual occlusion since the contact area is often obstructed by the body-part in contact with the manipulated object.

Observability is even more hampered when considering the highest form of imitation, often referred to as *third-person* imitation. This is the form of imitation which humans excel at: executing a task or a skill after having seen someone else performing it. As previously pointed out, first-person imitation (i.e., imitation using data of the task executed on the target robot) is in itself challenging due to the scarcity of technologies suitable for data collection of contact-rich tasks. Third-person imitation poses additional challenges. The first challenge consists of bridging the perception gap by translating what the robot sees (third-person human actions) into its own actions (first-person robot movements) and mapping observed goals to its own perspective; this is especially challenging when relying solely on observation data (e.g., images) without access to the interaction forces (e.g., tactile) and applied actions (e.g., applied joint-torques). Another challenge is associated with viewpoint and appearance discrepancy: the significant

differences in viewpoint (third-person human vs first-person robot) and appearance (human arm vs robot arm) make direct image translation or learning difficult. Zare et al. (2024) and Burnwal et al. (2025) explore more in details the difficulties related to imitating without actions (i.e., learning from observations) while Sharma et al. (2019) and Stadie et al. (2017) offer different preservatives on how to approach imitation from different viewpoints but they are limited to relatively simple and non-contact rich application domains.

Finally, in the context of contact-rich imitation the challenge of *data efficiency and generalization* becomes even more challenging due to the richness of involved data and the resulting difficulties in generalizing to new tasks or objects. This is also connected to the difficulties in developing physics-based models that accurately capture and generalize the highly nonlinear nature of contact-rich interactions.

Collecting demonstrations

This section discusses data collection, which is one of the most critical components in data-driven approaches. Robots consist of various sensors and actuators, and different combinations of sensors enable the acquisition of diverse data modalities. The *Data modalities section* describes the data modalities that are particularly useful for contact-rich tasks. The *Teaching methods section* then discusses methods for collecting demonstration data required for imitation learning (IL), that is, techniques for teaching robots. Furthermore, considering the high cost of data collection, which is a major challenge in robot learning, the *Synthetic data generation section* explains approaches for synthetic data generation. Finally, the *Available datasets section* introduces large-scale open datasets as well as datasets that include tactile information, which is especially important for contact-rich tasks.

Data modalities

Diverse data modalities are crucial for capturing the inherent complexity of contact-rich manipulation tasks, as they enable detailed representation of both spatial and dynamic interaction properties. Position data provides essential information for precise spatial alignment and trajectory following, whereas force measurements inform the nuanced adjustments necessary for stable and compliant interactions. Vision-based modalities extend capabilities to tasks involving environmental context and indirect monitoring of contacts, and tactile information offers direct feedback on local interaction dynamics. The integration of these diverse data modalities significantly enhances the robustness and generalization of IL methods in robotics (Ravichandar et al., 2020; Sherwani et al., 2020; Urain et al., 2024).

Positional data from joint angles and end-effector positions are fundamental for accurately replicating demonstrated trajectories. Precise positional information enables robots to achieve desired spatial configurations and smoothly transition between different motion phases. Force data complement positional information by providing necessary details about the magnitude and direction of forces applied during manipulation tasks. This information is critical for tasks requiring delicate adjustments, such as assembly or insertion operations, where appropriate force application can prevent damage to both the manipulated object and the robot itself (Kormushev et al., 2011; Peternel et al., 2015).

Vision-based modalities play a significant role, particularly in IL scenarios where direct interaction feedback may be limited. Visual sensors provide contextual awareness, enabling robots to interpret environmental states, object positions, and movements observed during demonstrations. However, challenges remain, such as visual occlusions and the indirect nature of force inference from visual observations. These limitations can restrict the precision of learned skills, emphasizing the necessity of complementing visual data with other sensory inputs (Dillmann, 2004; Vogt et al., 2017).

Tactile sensing and haptic feedback provide additional insights for interactions between robotic systems and their environments, since it can enable the detection and interpretation of complex contact phenomena, such as friction, slippage, texture differentiation, and subtle deformation, essential for precise force regulation and adaptive manipulation strategies (Edmonds et al., 2017; Higuera et al., 2024; Lambeta et al., 2024). Despite significant advancements in tactile sensor design and integration methodologies, the adoption of these technologies remains primarily limited to research domain. This limitation persists due to unresolved technical challenges, which include sensor durability under repeated mechanical stress, adequate sensitivity to subtle physical interactions, and seamless integration within existing robotic platforms. Nevertheless, current developments in innovative sensor architectures, complemented by data-driven processing approaches, indicate potential enhancements in the robustness, sensitivity, and applicability of tactile and haptic feedback systems, promoting their integration in broader robotic manipulation applications (Ablett et al., 2024; Edmonds et al., 2017; Higuera et al., 2024; Lambeta et al., 2024).

Effective fusion of multiple sensory modalities can further augment robotic capabilities, particularly in complex scenarios characterized by partial observability and dynamic uncertainties. Advanced multimodal integration approaches, including deep learning techniques, probabilistic inference, and filtering methods, facilitate comprehensive state estimation and enhance predictive capabilities. By leveraging complementary positional, force, visual, and tactile information, robots can achieve greater adaptability and

generalization across diverse manipulation tasks. Recent research underscores the benefits of such multimodal integration, demonstrating enhanced performance and robustness in practical robotics applications (Chen et al., 2024; Li and Zou, 2023; Urain et al., 2024).

Emerging data modalities, including electromyography (EMG) signals, soft and flexible sensors, are poised to further revolutionize robotic IL in contact-rich tasks. EMG has already seen applications in robotics, particularly in exoskeleton research (Peternel et al., 2016), providing insights into human muscle activation patterns, allowing robots to mimic not only observable movements but also internal force modulations (Peternel et al., 2014). Soft sensing technologies offer the potential to capture intricate interactions with complex geometries and materials, greatly enhancing sensitivity and versatility. Recent work by Liu et al. (2025b) introduces a novel hand-held device specifically designed to collect robot-free force-based demonstrations, facilitating more accessible data acquisition. Advanced wearable technologies facilitate naturalistic human demonstration capture, promising more intuitive and contextually rich data collection methodologies. The continued evolution of these modalities will likely lead to substantial advancements in the capability, accuracy, and applicability of IL techniques in contact-rich robotic manipulation (Shih et al., 2020; Zhang et al., 2022).

Teaching methods

Robot IL is based on learning human skilled movements and is also called learning from demonstration (LfD) or programming by demonstration (PbD), emphasizing the aspect of learning directly from human demonstrations. The first process is the acquisition of human skilled movements, which shares many technical commonalities with teaching techniques developed in the field of industrial robots.

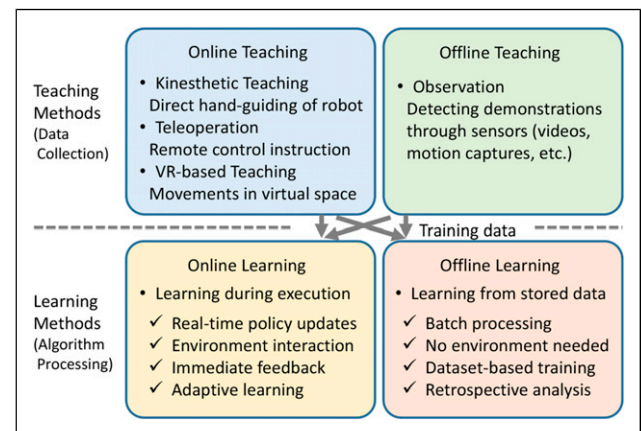


Figure 6. Relationship between robot teaching methods and learning methods.

Figure 6 shows the classification of teaching and learning. For IL, teaching is an essential process for obtaining training data, and it can be divided into online teaching and offline teaching. The difference lies in whether a human provides trajectories to the robot through online operation, or inputs them offline to the computer. The concept of direct teaching is also known (Ravichandar et al., 2020), while it can be interpreted as a teaching method where the robot is directly operated online, so it is classified as online teaching.

Online teaching is mainly classified into three categories: kinesthetic teaching, teleoperation, and VR-based teaching. Kinesthetic teaching involves direct hand-guiding of robots (Zhang et al., 2021a), while teleoperation utilizes remote control instructions. VR-based teaching, which captures movements in virtual space, is also increasingly being introduced. A recent work (Li et al., 2025a) compares the online teaching methods based on downstream learning performances and user satisfaction.

In offline teaching, either designer-calculated trajectories based on equations or trajectory-planning programs are input, or sensors observe human movements and the system creates command trajectories for input to the robot. In IL, only the latter method is used. Such methods, which detect human demonstrations through sensors, are called observation methods.

On the other hand, machine learning is also classified into online learning and offline learning. It is crucial to recognize that online/offline teaching and online/offline learning are distinct concepts that represent different aspects to be evaluated independently in methodological frameworks. These distinctions require careful consideration, as the terminology is frequently conflated in the academic literature.

Traditional industrial robot teaching assumes providing limited motions with high reproducibility in few trials. However, when environmental variations are large or task difficulty is high, it becomes necessary to acquire numerous motions and enhance adaptability through the generalization capabilities of machine learning. IL is effective in such cases. As mentioned in the *Insights from motor control* section, it is often necessary to acquire both position and force from human demonstrations to perform IL for contact-rich tasks. In teleoperation teaching, a leader robot operated by the teacher and a follower robot performing the task are connected, and bilateral control including force feedback is often implemented, allowing the operator's force adjustments to be transmitted to the follower robot (Kormushev et al., 2011). This enables simultaneous teaching of the operator's force and position information. When the follower robot in teleoperation is used directly for motion reproduction, the reproducibility of replay is high because there is no environmental variation when working in the same environment. Delays in teleoperation can degrade task performance, while this problem is avoided by the

operator's ability to compensate for those delays (Sasagawa et al., 2020). On the other hand, in teaching through a robot, the robot's dynamic characteristics can interfere with skilled movements. To reduce operator burden and achieve high skill levels, it is essential to ensure transparency in bilateral control to minimize this interference (Lawrence, 2002).

Meanwhile, observation teaching is also a popular method, in which cameras or motion capture observe human demonstrations, and the system processes these observations (mainly position information) to generate teaching data for the robot (Dillmann, 2004; Vogt et al., 2017). For contact-rich tasks, force sensors need to be embedded in tools (Furuta et al., 2020) or tactile gloves (Edmonds et al., 2017) to acquire force information. Methods that simulate force information during operation through the introduction of virtual reality technology are also useful (Aleotti et al., 2003; Zhang et al., 2018). Additionally, the combination of position and force observations with muscle activity measurement enables learning not only the motion but also stiffness behavior (Peternel et al., 2017). Observation teaching suppresses the degradation of teaching due to robot interference, so the quality of preserved skilled movements is high. However, since dynamics differ between humans and robots, variations in environments involving robots are unavoidable. The robot needs to suppress these variations through some method.

Figure 6 shows that IL can be categorized into the following four categories based on the combinations of online/offline teaching and online/offline learning:

- Interactive imitation
- Demo-augmented reinforcement learning
- Direct imitation
- Observational learning

First, the combination of online teaching and online learning is called interactive imitation. The most prominent example is dataset aggregation (DAGger) (Hoque et al., 2022; Ross et al., 2011), which implements a mechanism where humans correct mistakes in real-time. Second, the combination of offline teaching and online learning is called demo-augmented reinforcement learning. GAIL (Gubbi et al., 2020; Li et al., 2021; Tsurumine et al., 2019; Xiang et al., 2024) performs adversarial matching with expert demonstrations. Deep deterministic policy gradient from demonstrations (DDPGfD) (Vecerik et al., 2017) and demonstrations augmented policy gradient (DAPG) (Rajeswaran et al., 2018) can also be considered as types of IL since RL is initialized with demonstrations. Third, the combination of online teaching and offline learning is called direct imitation. This category includes behavior cloning (BC) and dynamic movement primitive (DMP) methods that use online teaching approaches such as kinesthetic teaching, teleoperation, and VR-based teaching. Finally, the

combination of offline teaching and offline learning is called observational learning. BC from datasets involves learning from large-scale demonstration collections, while vision-based IL realizes learning from visual observations.

Synthetic data generation

In modern robotics, synthetic data generation has become a key enabler for the development of data-intensive learning based systems. It offers diverse and scalable environments that can be crucial for algorithms’ training and validation. When it comes to robotic tasks with physical contacts, the role of synthetic data generation to enable autonomous robotic behaviors become even more significant due to the complexity, variability, and safety-critical nature of contact dynamics, which are difficult to capture, annotate, and scale in real-world data collection. Hence, synthetic data generation offers numerous advantages such as access to contact dynamics like forces and friction properties, fast domain adaptation, and risk mitigation in safety-critical scenarios such as industry and healthcare (James et al., 2020; Liu et al., 2023a; Mees et al., 2022b; Moghani et al., 2025; Yu et al., 2020).

Over the past years, several platforms offering synthetic data generation in contact manipulation have been introduced. Physical simulators such as MuJoCo (Multi-Joint dynamics with Contact) (Todorov et al., 2012), PyBullet (Coumans and Bai, 2016), Isaac Gym (Makoviychuk et al., 2021), and Drake (Tedrake and the Drake Development Team, 2019) are the most commonly used examples in the community, although every year new platforms with even stronger physics simulation capabilities emerge. MuJoCo prioritizes efficiency and smooth contact dynamics, making it a leading choice for model-based control and RL in continuous action spaces. PyBullet, known for its user-friendly interface and extensive community support, excels in accessibility and robotic manipulation, although with less precise contact modeling. Isaac Gym, exploiting NVIDIA’s GPU acceleration, enables high-throughput parallel simulations, ideal for large-scale RL, especially in tasks involving complex contact. Finally, Drake employs hydroelastic contact mechanics, providing a principled and accurate approach, crucial for planning, control, and formal verification in safety-critical or high-fidelity applications. These simulators collectively offer a spectrum of trade-offs between computational speed, physical realism, and implementation complexity for synthetic data generation.

Regardless of the chosen simulation environment, several practical techniques have been introduced for synthetic data generation. Domain randomization (Tobin et al., 2017), for example, is used to vary certain parameters of interest (e.g., object weight, surface texture) to make the learned models more robust to real-world variations. This method

helps algorithms account for such changes and achieve performance in real-world applications that more closely matches their performance in simulation. Another method to bridge the sim-to-real gap is based on physics-based augmentation of the generated data (Yang et al., 2025). It enhances synthetic datasets by incorporating physical models or constraints into the data creation or transformation process.

Despite fundamental progress in these techniques, transferring learned policies or models to real-world robots remains a significant challenge due to the sim-to-real gap, which enlarges as the number and complexity of contact dynamics increase (James et al., 2019; Peng et al., 2018; Tobin et al., 2017). This is why several recent studies have focused on fine-tuning the learned algorithms with real-world data (Chebotar et al., 2019; Finn et al., 2017; Yu et al., 2018). In parallel, more advanced and emerging strategies, such as differentiable simulation (Degraeve et al., 2019; Freeman et al., 2021a), synthetic tactile data generation (Wang et al., 2022a), and learning from hybrid (simulated and real) datasets (Ferguson et al., 2020), are gaining increasing attention for their potential to further reduce the sim-to-real gap in contact-rich manipulation tasks.

Available datasets for contact rich tasks

The field of contact-rich IL heavily relies on diverse and extensive datasets that capture interactions between robots and their environment through various sensory modalities. Several notable datasets have been developed to address the challenges of data scarcity, generalization, and multimodal integration in this domain.

Some datasets are worth mentioning considering the contact-rich nature of the selected tasks, though these datasets do not contain tactile data. BridgeData V2 (Walke et al., 2023) is a large-scale dataset with over 60,000 trajectories collected across 24 environments using a low-cost WidowX 250 robot arm. It integrates RGB images, depth data, and natural language instructions, supporting open-vocabulary task specification for various IL and offline RL methods, with a focus on generalizing skills across different environments and institutions. The DROID (Distributed Robot Interaction Dataset (Khazatsky et al., 2024)) further extends this with an unprecedented scale, featuring 76,000 demonstration trajectories (350 hours of interaction) across 564 scenes, 52 buildings, and 86 tasks. Collected by a distributed network of 50 data collectors across 18 labs worldwide on the Franka Panda robot arm, DROID includes synchronized RGB camera streams, depth information, and language annotations, aiming to enhance policy performance and robustness in “in-the-wild” scenarios. The Open X-Embodiment (OXE) Dataset (O’Neill et al., 2024) is a significant aggregation, combining over 1 million real robot trajectories from 60 existing datasets

across 22 robot embodiments and 21 institutions. This large-scale repository provides diverse robot behaviors, embodiments, and environments in a standardized format, facilitating research into X-embodiment training for generalizable robot policies, including those like RT-X models.

With the increasing availability of reliable tactile sensors, a number of large datasets which contain tactile data are published. The TVL Dataset (Fu et al., 2024a) comprises 44,000 in-the-wild vision-touch pairs, featuring tactile data from DIGIT sensors and visual observations. A significant portion (90%) of its English language labels are pseudo-labeled by GPT-4V, while 10% are human-annotated, aiming to bridge the gap in integrating touch into multimodal generative language models. Similarly, the Touch100k Dataset (Cheng et al., 2025) focuses on GelSight sensors, offering over 100,000 paired touch-language-vision entries with multi-granularity linguistic descriptions. This dataset, curated from existing tactile datasets like TAG and VisGel, utilizes GPT-4V for generating detailed textual descriptions, which are then refined through a multi-step quality enhancement process involving Gemini 2 for consistency assessment. It aims to improve tactile representation learning for tasks such as material property identification and robot grasping. The VisGel Dataset (Li et al., 2019) also provides a large collection of 3 million synchronized visual and tactile images from 12,000 touches on 195 diverse objects, collected using KUKA LBR iiwa robotic arms equipped with GelSight sensors and webcams, to explore cross-modal prediction between vision and touch.

The Sparsh project (Higuera et al., 2024) introduces a curated dataset of approximately 661,000 images from various vision-based tactile sensors (DIGIT, GelSight) for self-supervised learning, alongside the TacBench benchmark. TacBench offers six touch-centric tasks with labeled data for evaluating touch representations, including force estimation, slip detection, pose estimation, grasp stability, and textile recognition, demonstrating the value of pre-trained touch representations for contact-rich manipulation.

Lastly, a dataset was collected for a study on Multimodal and Force-Matched IL with a See-Through Visuotactile Sensor (Ablett et al., 2024). This dataset, created through kinesthetic teaching with a 7-DOF robotic system, includes visuotactile, wrist camera, and relative end-effector pose data, with a focus on improving IL for door-opening tasks through tactile force matching and learned mode switching. These datasets collectively advance the capabilities of robots in handling complex, contact-rich tasks through multimodal sensory integration and scalable learning approaches.

Learning approaches

This section discusses imitation learning (IL) approaches under seven categories that are prominent in the literature, and an additional category for alternative perspectives. Figure 7 presents these categories with a rough comparison of their




Figure 7. Reviewed learning approaches. Comparison of expert data requirements and ranking of generalization and multimodality handling capabilities. Offline and inverse RL approaches have similar generalization capacities. Inverse RL methods require much less expert data; however, they also require environment interactions.

demonstration data requirements and their generalization capabilities. Following sections provide a more refined discussion of each category. We begin with *behavior cloning* (BC), a supervised learning method that demonstrates stable performance in IL. Then, we introduce *dynamic movement primitives* (DMPs), which are commonly used for motion representation in IL. This is followed by the *generative methods*, where robot behaviors are learned by generative models with techniques such as BC. We also discuss *foundation models*, which are trained on large-scale, pre-collected data with a generative model architecture to build general-purpose models. The next section focuses on *inverse RL* (IRL), which estimates a reward function from expert data and learns a policy to maximize the rewards from the estimated function. Next, we discuss *multimodal IL*, which simultaneously processes other modalities, such as haptic feedback, in addition to position trajectories. Then, we introduce *offline RL*, which, as an off-policy method, learns from a pre-collected dataset of expert data in an offline format without interacting with the environment. And the last section introduces the *other methods* related to IL.

Behavior cloning

Behavior cloning (BC) has emerged as a particularly influential paradigm, attracting substantial research interest due to its straightforward implementation and empirical effectiveness. In particular, its ability to efficiently learn from pre-collected data is especially attractive for IL research. BC constitutes a supervised learning paradigm in which an artificial agent is trained to replicate expert behavior using demonstration data (Levine et al., 2016; Ross and Bagnell, 2010). Specifically, this methodology involves training the agent to generate appropriate actions a in response to corresponding state inputs s , learning the mapping between environmental states and expert-demonstrated actions from the data distribution \mathcal{D} . The objective function for BC can be expressed as

$$J(\theta) = \mathbb{E}_{(s,a) \sim \mathcal{D}} [\|\pi_{\theta}(s) - a\|^2] \quad (1)$$

where policy π_{θ} is parameterized by a vector θ . This supervised learning approach is well suited for acquiring

motor skills when expert data are available, such as demonstrations corrected via teleoperation by human operators. For instance, leveraging large-scale datasets of state-action pairs has enabled robots to learn and execute manipulation tasks, such as pick-and-place operations involving objects within kitchen drawers across various domains (Brohan et al., 2023; Zitkovich et al., 2023). These advancements demonstrate the feasibility of learning intricate manipulation sequences that require precise spatial awareness and robust control strategies.

In BC, models capable of processing sequential data have been widely adopted. Representative approaches include recurrent neural network (RNN) (Ito et al., 2022) and long short-term memory networks (LSTM) (Adachi et al., 2018; Funabashi et al., 2020; Kutsuzawa et al., 2018; Rahmatizadeh et al., 2018; Scherzinger et al., 2019). For instance, LSTMs have been successfully applied to tasks involving deformable objects, such as closing the zipper of a bag (Ichiwara et al., 2022). Yang et al. (2016) demonstrated a method for teaching a robot to fold fabric using a time-delay neural network (TDNN) combined with a deep convolutional autoencoder (DCAE). Sequence-to-Sequence (Seq2Seq) models have been utilized for learning contact-intensive manipulation tasks. Kutsuzawa et al. (2018) incorporated a contact dynamics model into a Seq2Seq framework with an embedded LSTM, allowing a robot to scoop and rotate objects using a spatula. Similar models have also been applied to tasks such as toilet cleaning (Yang et al., 2020b), where precise positional correction is required, and door opening (Yang et al., 2023), which demands effective force regulation.

Transformers (Vaswani et al., 2017) have become central to robot IL, as they handle long sequences more effectively than RNNs or LSTMs (Sherstinsky, 2020) through attention mechanisms and positional encodings. The action chunking transformer (ACT) enables robots to learn complex tasks, such as opening a cup lid or putting on shoes, using a low-cost teleoperation system called ALOHA that facilitates high-quality demonstrations. Since IL depends heavily on data quality, advances in hardware platforms such as Mobile ALOHA (Fu et al., 2024b) are expected to further accelerate progress.

Beyond Transformers, other generative models have been applied to BC. The Mamba model (Jia et al., 2024), a state-space model designed for efficient long-sequence processing, enhances generalization with limited data by focusing on salient features. As a motion encoder, it compresses robotic motion sequences while preserving key temporal dynamics for accurate prediction (Tsuji, 2025). Additionally, implicit behavior cloning (Florence et al., 2022), using energy-based models, offers advantages for tasks with discontinuous transitions, such as contact-rich manipulation. Diffusion models have also gained attention, with studies demonstrating their ability to generate force and position trajectories through denoising (Liu et al.,

2025b). Further discussion of research on transformers and diffusion policies is provided in the *Other generative methods* section.

A key limitation of current BC approaches is the lack of adaptability. In contact-rich tasks, when predicted actions fail to achieve desired contact states, systems must autonomously adjust their behavior. Effective adaptation requires feedback mechanisms that inform how to correct actions in response to environmental changes, ensuring robust performance in dynamic settings. In addition, a common issue in IL is compounding error—the accumulation of incorrect actions over time, which can hinder proper task execution. This problem is amplified in techniques like action chunking, where action sequences are generated in blocks, increasing the risk of error propagation.

Several studies have developed the ability for autonomous behavior correction in the context of IL. Ankile et al. (2024b) have incorporated residual RL policy into base policy trained in BC manner to produce chunked actions. This approach aims to have higher frequency closed-loop control in order to correct behavior in the context of assembly task. Another approach is called Corrective Labels for IL (CCIL) (Deshpande et al., 2024) generates some state-action pair to bring the agent back to the expert state to deal with compounding error.

To improve policy learning in BC, some studies incorporate human feedback. Language-conditioned methods guide robots with verbal cues (e.g., “move to the right”) to resume tasks autonomously (Jang et al., 2022; Shi et al., 2024). Additionally, DAGger (Ross et al., 2011) variants like ThriftyDagger (Hoque et al., 2022) and LazyDagger (Hoque et al., 2021) iteratively integrate human corrections during execution, helping reduce distributional shifts and improve generalization without constant supervision. Moreover, DPIIL (Oh and Matsubara, 2024), which gates based on risks assessed considering the speed of expert demonstrations, improves policy performance on tasks with clearance constraints. To address uncertainty arising from multimodality, Diff-Dagger (Lee et al., 2025), which employs Diffusion policy, has also been proposed. Another approach to improving policy robustness is to augment demonstration data. Methods such as disturbances for augmenting robot trajectories (DART) (Laskey et al., 2017) and Bayesian disturbance injection (BDI) (Oh et al., 2021) introduce noise during data collection, requiring the expert to compensate for the disturbances and thereby enriching the demonstration dataset. Currently, refining policies using human feedback has mainly explored positional trajectory adjustments.

Dynamic movement primitives

DMPs are a widely known motion representation method that facilitates generalizability and encourage convergence of the learned trajectories thanks to their dynamical system-based formulation (Schaal et al., 2003). The classical DMP

formulation fits the target trajectory by learning the weights w_i of phase-shifted Gaussian basis functions ϕ_i , known as the *forcing term*

$$f(x) = \frac{\sum \phi_i w_i}{\sum \phi_i} x, \phi_i = e^{-h_i(x-c_i)^2}, \quad (2)$$

where c_i is the center and h_i is the width of each basis function. x is the phase variable that decays with constant rate α_x toward zero in time, described as the *canonical system*

$$\dot{x} = -\alpha_x x. \quad (3)$$

The *transformation system* is formed as a spring-damper system with stiffness β_y and damping α_y that drives the system state y toward the goal state g

$$\ddot{y} = \alpha_y(\beta_y(g - y) - \dot{y}) + f(x). \quad (4)$$

In the original approach, a single DMP was fit to a single one-dimensional trajectory of the position or joint angle. Learning multi-dimensional motions was possible by learning multiple DMPs, each of which describing an independent motion dimension, modulated with the same canonical system. However, this brought the limitation of representing coupled spaces such as quaternions and rotation matrices. Thus, some of the works focused on representing different spaces and manifolds with DMPs. Ude et al. (2014) extended the DMPs to represent quaternions and rotation matrices for non-minimal, singularity-free orientation handling. Abu-Dakka and Kyrki (2020) proposed the geometry-aware DMPs (\mathcal{G} -DMPs) to support symmetric positive definite matrices that enable learning the stiffness, damping and manipulability ellipsoids as part of the motion. They later generalized the \mathcal{G} -DMP formulation to support any Riemannian manifold including the quaternions and multi-dimensional rotation matrices (Abu-Dakka et al., 2024).

Another limitation of the classical DMPs has been capturing the variance in multiple demonstrated examples. Probabilistic movement primitives (ProMPs) (Paraschos et al., 2013) are a widely-adopted DMP-derivative that learns a probabilistic distribution over multiple trajectories. ProMPs can model the variance in the demonstrations and combine multiple learned primitives smoothly. Another DMP derivative framework is kernelized movement primitives (KMPs) (Huang et al., 2019) that replaces the basis functions with a kernel-based non-parametric approach. KMPs have the advantages of efficiently handling high-dimensional data and allowing via-point based motion modulation. Probabilistic DMPs (ProDMP) (Li et al., 2023) unify the DMPs and ProMPs to retain the useful properties of both the dynamical systems and the statistical distributions.

Some of these works (Huang et al., 2019; Li et al., 2023; Paraschos et al., 2013; Ude et al., 2014) are not contact-rich

applications, but they lay the foundation for some of the contact-rich work we mention later. In the following, we limit our discussion primarily to the recent works focusing on contact-rich IL, rather than presenting an exhaustive list of movement primitive (MP) works. We refer the reader to an earlier survey (Saveriano et al., 2023) for more general and historical views on the MP literature. However, we also include other MP works such as those based on Gaussian mixture models (GMMs) (Calinon et al., 2010) instead of the Gaussian basis functions (2). These works do not directly extend DMPs; however, they were historically developed in parallel and aimed to answer similar problems with similar principles (Khansari-Zadeh and Billard, 2011; Ureche et al., 2015).

As discussed earlier, the traditional DMP formulation is limited to a single modality and a single task that is often a position trajectory in joint or task space. However, contact-rich tasks require high-level context awareness that is possible with multimodality. Consequently, the MP-based contact-rich manipulation methods answer this problem by either *parallelization* of the other modalities as separate perception or action modules, or *reformulation* of the MPs.

The modality *parallelization* strategy can be traced back to early MP research (Kober et al., 2015; Nemeč et al., 2013). For example, an early peg-in-a-hole application (Nemeč et al., 2013) recorded a force profile alongside the DMPs. The DMPs were adapted according to an admittance control law to match the force profile. Focusing on more recent works, Cho et al. (2020) train hidden Markov model (HMM) and DMP models in parallel for each motor skill: former to select which MP to apply based on the reaction force/moment signal, latter to encode the position-based motion trajectories. Chang et al. (2022) train separate DMPs for position and force trajectories. Then, the impedance is adapted during execution to balance between position and force tracking. Zhao et al. (2022a) train GMMs for position, velocity and force profiles. Then, they adapt the impedance parameters through online optimization to achieve learned profiles. Escarabajal et al. (2023) accommodate the force profile in parallel using GMMs, while using DMPs for the position trajectory. Yang et al. (2018) use GMMs for learning from multiple demonstrations and train a neural network-based controller to compensate for the dynamic effects.

Through *reformulation*, compliant movement primitives (CMP) (Deniša et al., 2015) add the joint torque modality to model the task-specific dynamics. Petric et al. (2018) improve the CMP framework for safe and autonomous learning of the joint torque profiles. Another reformulation, Bayesian interaction primitives (BIPs) (Campbell et al., 2019) integrate human monitoring modalities to achieve coordination in human-robot interaction (HRI) tasks. Stepputtis et al. (2022) use the BIP policies to modulate the temporal progress of a bimanual multipoint insertion task. They use multimodal sensing (force, proprioception, object

tracking) to identify the task phase, and control two robot arms accordingly. Ugur and Girgin (2020) propose the compliant parametric DMPs that learn haptic feedback trajectories through parametric HMMs and reproduce the desired force profile through a compliance control term. Qian et al. (2025) propose the hierarchical KMPs to generalize a learned motion from known subregions to novel subregions using the correlations between the human and robot positions in object hand-over tasks. Lödige et al. (2025) extend the ProDMP framework to be force aware (FA-ProDMP). FA-ProDMP learns the force-position correlations from multiple demonstrations to solve contact-rich tasks like peg-in-hole.

MP *reformulation* does not only aim to support multi-modality, but also to answer task-specific requirements. Yang et al. (2022) choose rhythmic DMPs to represent robot policies in periodic household tasks such as table wiping, food stirring, and cable wiring. Unlike the classic DMP approach, the robot learns the task through visual keypoints extracted from human video demonstrations. Rhythmic DMPs facilitate reproducing and adjusting periodic actions. Sidiropoulos and Dougeri (2021) propose the reversible DMPs that support backwards reproduction of a learned trajectory, which is a desirable feature to recover from errors in physical interaction and operate in unpredictable environments such as cluttered areas. Escarabajal et al. (2023) use the reversible DMPs to encode the trajectory so that the mechanism can safely retract its action when assisting an injured person. Mesh DMPs (Dalle Vedove et al., 2025) extend the \mathcal{G} -DMPs for reproducing learned motions on complex mesh surfaces to achieve contact-rich tasks like surface polishing.

MPs provide a structural basis for a generalizable motion. The structure confines the parameter search space more than an unstructured multi-layer perceptron, leading to a better sample efficiency. Thus, movement primitives are often used as the policy to be initialized using IL, and further improved using RL. Cho et al. (2020), who train parallel HMM and DMP models, apply RL to further improve both DMP-HMM parameters. In the case that HMMs do not identify a matching skill, a new DMP-HMM pair is learned using RL. Davchev et al. (2022) combine a DMP policy with RL by learning a residual correction policy to account for the contact-rich aspect of physical insertion tasks. They introduce an additive coupling term in the DMP formulation and learn this term as a nonlinear RL-based strategy. They show the advantages of perturbing the DMP policy directly in the task space on both task success and sample efficiency. Zang et al. (2024) combine IL and RL by first learning ProMP policies from demonstrations and then guiding the RL training based on ProMP priors. ProMPs are chosen for their flexibility, smoothness, and generalization properties.

Just like embedding or coupling another model into the DMP framework, DMPs can also be embedded into more expressive models to leverage the benefits of both. Bahl

et al. (2020) embed a DMP-based dynamic system into a neural network to take the advantages of both the generalization capacity of neural networks and the efficiency of dynamical systems. In their architecture, deep neural layers learn the parameters of DMP systems for various raw input from vision or other sensors. The parametrized DMPs then derive the motion trajectories to be executed. Conditional neural movement primitives (CNMPs) (Seker et al., 2019) use conditional neural processes (CNPs) to learn sensorimotor distributions of multiple modalities. The learned CNPs are then conditioned to generate trajectories for new situations.

DMPs have been a central part of robotic learning from demonstration since their introduction. They are actively being extended and used for contact-rich tasks beside more recent and trending methods. Their generalizability, sample-efficiency, and transparency are invaluable for robotic tasks where data collection is costly and risk-sensitive. We identified the common modes in which the DMPs are employed in recent works, such as *reformulation*, *parallelization*, and in combination with RL. That being said, these modes are not mutually exclusive. They have been used together in some of the works we cited above (Cho et al., 2020; Davchev et al., 2022; Escarabajal et al., 2023). We expect DMPs to stay present and get integrated with novel methods like generative models in the future, thanks to their fundamental structure and various proposed extensions.

Generative methods

Variational autoencoder. Within the domain of probabilistic modeling, auto-regressive models constitute a family of architectures that maintain both high expressivity and computational tractability. These models facilitate the decomposition of log likelihood according to the following expression: $\log p(x) = \sum_i \log p_\theta(x_i|x_{<i})$. However, to explicitly learn a compact latent representation of the data, variational autoencoders (VAEs) introduce a parametric inference model $q_\phi(z|x)$ over the latent variables. In lieu of direct log-likelihood optimization, an alternative approach involves the introduction of a parametric inference model $q_\phi(z|x)$ over the latent variables, enabling the optimization of a lower bound on the log-likelihood. The VAE objective function takes the following form:

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x) || p(z)) \leq \log p(x), \quad (5)$$

In this formulation, $\mathbb{E}_{q_\phi(z|x)}$ signifies the expectation computed with respect to the approximate posterior distribution, while D_{KL} corresponds to the Kullback-Leibler divergence measure. Using this formulation, VAE has been used for diverse behavior learning (Wang et al., 2017).

VAE has been employed in IL by training them to reconstruct position or torque commands (Abolghasemi et al.,

2019). VAE is interpreted as a means to obtain compact representations of tactile or visual signals (Royo-Miquel et al., 2023; Yoo et al., 2021). It is used for skill learning via RL (Cong et al., 2022; Van Hoof et al., 2016). Language is also incorporated using the paired variational autoencoder (PVAE) model (Özdemir et al., 2023). Tsuji et al. (2025) employed a VAE-based model to learn contact-rich tasks like wiping, using the same architecture for simulation pre-training and real-world fine-tuning. They also generated contact-maintaining motions via force feedback in latent space, enabling adaptation to surface variations. Liconti et al. (2024) pre-trained a VAE on a large dataset of human hand-object images, which are easier to collect than robot demonstrations. They then trained a task-specific policy using a smaller, task-focused dataset, leveraging the pre-trained decoder to generate actions.

The conditional variational autoencoder (CVAE) is an extension of the VAE that incorporates conditioning variables, allowing for the generation of data that is influenced by specific conditions. In IL, CVAE is utilized to modify the reconstructed behavior based on these conditioning variables, enabling more adaptive and context-aware action generation. Zhang and Demiris (2023) proposed a predictive model which is a CVAE with contrastive optimization, jointly learning visual-tactile representations and latent dynamics of deformable garments. Another approach involves training a CVAE architecture with input torque and a task-specific parameter such as a task ID, allowing the learned model to adapt its behavior according to the given task (Xu et al., 2025). A similar approach has also been applied in combination with movement primitives (Noseworthy et al., 2020), leveraging CVAE’s conditioning capabilities to adjust generated motions dynamically. By incorporating conditioning parameters into the learning process, the generated motion can be modified by altering these parameters, making CVAE particularly effective for scenarios requiring motion correction, such as adjusting actions based on contact states with objects (Ren et al., 2021). Moreover, Mees et al. (2022a) handled the multimodality of free-form imitation data by encoding demonstrations into a latent plan space using a sequence-to-sequence CVAE. Conditioning the policy on these plans allows it to focus entirely on learning uni-modal behaviors.

Encoder-decoder architectures capable of conditioning, such as CVAE, have also been employed in IL using Transformer models and Mamba models. ACT (Zhao et al., 2023a) and Mamba IL (MaIL) (Jia et al., 2024) follow a CVAE-like structure, demonstrating the effectiveness of CVAE-based architectures in IL for contact-rich tasks.

Other generative methods. Recently, in IL, generative methods have been used to model the policy or the environment dynamics by generating synthetic data or actions that mimic expert demonstrations. Beyond VAEs and

foundation models, there are other notable generative approaches in IL.

Diffusion Policy (Chi et al., 2023) is a recent advancement generating robust and multimodal action sequences for robot control. Unlike traditional policy methods, it iteratively denoises action distributions, enabling smoother and more diverse behaviors that can handle complex, real-world tasks. Recent studies have demonstrated its superiority over traditional methods like VAEs and GAIL in contact-rich manipulation tasks. Ankile et al. (2024a) use diffusion policy for automatically expanding dataset size. Prasad et al. (2024) proposed Consistency Policy, which is distilled from a pre-trained diffusion policy while drastically reducing inference time so that can be used in resource-limited cases. For battery disassembly tasks, diffusion policies have been applied to generate robot actions (Kang et al., 2025). In this context, a cross-attention mechanism is employed to process high-dimensional visual inputs and low-dimensional force signals as inputs to the diffusion policy, enabling appropriate action generation. The adaptive compliance policy (ACP) (Hou et al., 2025) controls robot motions using target poses and stiffness matrices generated through the diffusion process. The reactive diffusion policy (RDP) (Xue et al., 2025) introduces a hierarchical structure, where one policy generates latent actions from low-frequency visual observations and another generates actions from high-frequency force signals, thereby enabling responsive force feedback. Wu et al. (2025) combines a diffusion policy with model-based planning for broader generalization.

Action chunking with transformers (ACT) (Zhao et al., 2023a) uses a transformer-based generative model to predict “chunks” of actions (instead of single-step actions) conditioned on past states, enabling long-horizon, temporally consistent policies. It’s particularly effective for robotic manipulation tasks. ACT has been actively explored in IL, particularly for robotic control tasks where long-horizon action consistency and multimodal behavior are critical for contact-rich tasks. Buamanee et al. (2024) propose Bi-ACT model utilizing both positional and force data, enhancing the precision and adaptability of robotic tasks. In transformer-based approaches that leverage both positional and force information, Comp-ACT (Kamijo et al., 2024) has been proposed to predict not only trajectories but also stiffness parameters. To reduce computational load, RoboMT (Rundong et al., 2025) employs a hybrid architecture combining Mamba and Transformers together with adaptive action chunking. Moreover, FACTR (Liu et al., 2025a) integrates curriculum learning into transformer-based policies to effectively process visual and force information.

Wu et al. (2024) introduced GR-1, a GPT-style transformer pre-trained on large-scale video data and fine-tuned for multi-task, language-conditioned visual robot

manipulation. GR-1 takes language instructions, observation images, and robot states as input, and predicts both robot actions and future images.

In addition, [Chen et al. \(2024\)](#) proposed a generative model called Elemental, which is designed to learn from human demonstrations and generate robot actions for contact-rich tasks. Elemental leverages a combination of generative modeling techniques, including diffusion models and variational inference, to capture the underlying structure of the task space. By learning from diverse human demonstrations, Elemental can generate robot actions that closely resemble human-like behavior, enabling effective IL in complex manipulation scenarios. [Zhang et al. \(2025b\)](#) presented a variable impedance controller supervised by a VLM to integrate semantic reasoning from multimodal inputs with low-level compliant control, enabling robots to adapt their impedance parameters in real-time for safe and effective manipulation in contact-rich scenarios. It uses retrieval-augmented generation (RAG) and in-context learning to improve the VLM suggestions with prior experience.

Foundation models

Foundation models have shown remarkable success in generalization and reasoning capabilities in language and vision modalities. Vision and language were followed by audio and navigation modalities ([Firoozi et al., 2023](#); [Kawaharazuka et al., 2024](#)). Apart from the development of transformers ([Vaswani et al., 2017](#)) and other enabling technologies, the availability of internet-scale data was a key requirement for the development of foundation models. Readily available text and image data fueled the generalization capability of the foundation models; however, this is not the case for the robotics data. Unlike the visuolingual data, embodied interaction modalities such as proprioception and force data are not widely publicized. Currently available robotics datasets have sample sizes between 100K and 1 M examples, which are incomparable to those of the LLMs ([Kim et al., 2024b](#)). These limitations are followed by the additional challenges of the robotics field, such as the presence of various robotic platforms, and safe real-time execution requirements of embodied systems ([Firoozi et al., 2023](#)). These challenges are especially important in case of the contact-rich applications as the safety requirements are high and real-time contact-related data is fundamental.

Vision-language models (VLMs) have emerged as powerful tools in robot learning, enabling robots to interpret and act upon human instructions grounded in visual scenes. CLIPORT ([Shridhar et al., 2022](#)) uses two-stream architecture: a semantic stream that encodes RGB input with a frozen CLIP ResNet50 and conditions decoder layers with language features, and a spatial stream that encodes RGB-D input and fuses laterally with the semantic stream. The output is dense pixel-wise features for predicting pick and place actions. Manipulation of open-world objects (MOO)

([Stone et al., 2023](#)) enables robots to follow instructions involving unseen object categories by leveraging a pre-trained VLM to extract object information, which conditions the robot policy alongside the image and language command. RoboFlamingo ([Li et al., 2024b](#)) also uses an existing VLM by incorporating an explicit policy head, and is fine-tuned by IL only on language-conditioned manipulation datasets. [Yan et al. \(2024\)](#) leverage NeRF for 3D pre-training to acquire a unified semantic and geometric representation. By distilling knowledge from pretrained 2D foundation models into 3D space, the method constructs a semantically meaningful 3D representation that incorporates commonsense priors from large-scale datasets, enabling strong generalization to out-of-distribution scenarios. [Duan et al. \(2025\)](#) introduce AHA, an open-source VLM for detecting and explaining robotic manipulation failures using natural language. Trained with FailGen, a scalable framework that generates failure data by perturbing successful simulations, AHA generalizes well to real-world failures across diverse robots and tasks.

In robot action generation using VLMs, a common approach is to use VLMs to extract semantic features from visual and language inputs, which are then fed into a separate policy module for action prediction. Recently, end-to-end vision-language-action (VLA) models that directly map images and language instructions to robot actions have gained increasing attention. Notable examples include RT-1 ([Brohan et al., 2023](#)), which was trained on a dataset collected over 17 months using a fleet of 13 robots, and RT-2 ([Zitkovich et al., 2023](#)), which employs VLMs trained on Internet-scale data. These models have demonstrated impressive generalization capabilities in executing various real-world tasks. [Zhao et al. \(2025a\)](#) proposed VLAS (VLA with speech instruction), which encodes visual and speech inputs, retrieves personalized knowledge via a Voice RAG module, and uses LLaMA to generate action tokens, which are decoded into continuous robot control commands. [Li et al. \(2025b\)](#) proposed hierarchical VLA models that better leverage off-domain data for robotics by first predicting a coarse 2D trajectory with a VLM, which then guides a low-level 3D control policy for precise manipulation. The self-corrected (SC-)VLA framework ([Liu et al., 2024](#)) enhances manipulation robustness by combining a fast action predictor with a slow, reflective system that uses Chain-of-Thought reasoning to correct failures step by step, mimicking human-like reflection. Moreover, diffusion-based models have been introduced to enable VLMs to predict future states and select actions based on this information through a reflective mechanism. ([Feng et al., 2025](#)) Robocat ([Bousmalis et al., 2024](#)) is a generalist transformer agent that natively supports multiple robotic embodiments with different action- and state-spaces. Furthermore, Robocat can self-improve by generating new data by its own model. [Driess et al. \(2023\)](#)

focus on the problem of grounding multimodal prompts for real-world manipulation planning using LLMs. They propose the multimodal sentences that embed images and state embeddings in text prompts for improved embodied intelligence. Kim et al. (2024b) propose an open source VLA outperforming the state-of-the-art in object manipulation tasks with a smaller parameter space. They employ fine-tuning techniques such as low-rank adaptation (Hu et al., 2022) and model quantization (Dettmers et al., 2023) to efficiently adapt the model on an off-the-shelf GPU. Hao et al. (2025) train a policy through tactile information to execute interactive contact and collision tasks based on VLA that outperform the traditional IL methods. Yu et al. (2025) integrate wrist force/torque and tactile cues into the VLA loop for closed-loop contact regulation, improving robustness and success on insertion and drawer-manipulation tasks (Yu et al., 2025). Zhang et al. (2025a) proposed VTLa, a vision-tactile-language-action model that integrates tactile sensing with vision and language inputs to enhance robot manipulation capabilities. By combining these modalities, VTLa improves the robot's ability to understand and interact with its environment, leading to more effective and precise manipulation tasks. Recent studies have introduced VLA models such as π_0 (Black et al., 2025b) and its more hierarchically structured successor, $\pi_{0.5}$ (Black et al., 2025a), which exhibit strong performance on household tasks in chaotic or previously unseen environments. Given the high cost of data collection, recent studies have focused on improving training efficiency. GraspVLA (Deng et al., 2025) generated a billion-frame robotic grasping dataset in simulation to train a VLA model, reducing the impact of the sim-to-real gap through photorealistic rendering and extensive domain randomization. DexVLA (Wen et al., 2025) further proposed an efficient adaptation approach that first acquires general knowledge from cross-embodiment data and then fine-tunes the model for specific embodiments and tasks.

In recent years, the concept of embodied AI/embodied LLM has been discussed for application to robotics. Chen et al. (2025b) investigated embodiment-aware LLM-based multi-agent system (MAS) that operates a heterogeneous multi-robot system composed of drones, legged robots, and wheeled robots with robotic arms in a multi-floor house. When given a household task, the agent needs to understand their respective robots' hardware specifications for task planning and assignment. Zhang et al. (2025e) identified safety vulnerabilities in embodied AI systems that use LLM for physical robots. They introduced Badrobot, an attack method that exploits three key weaknesses: LLM manipulation within robotic systems, misalignment between language outputs and physical actions, and hazardous behaviors from flawed world knowledge. These attacks use voice interactions to make embodied LLMs violate safety and ethical

constraints, highlighting critical security risks in AI-powered physical systems.

Although these works include contact-rich tasks such as object pushing and cloth folding, these are usually solved in a quasi-static way, through position control and pick-place actions. The use of foundation models in dynamic contact-rich tasks where the force needs to be regulated remains a challenge.

Inverse reinforcement learning

Inverse reinforcement learning (IRL) (Ng and Russell, 2000) also known as reward inference (Kroemer et al., 2021), or inverse optimal control (Englert et al., 2017; Englert and Toussaint, 2017; PJ and BDO, 1971). RL has shown a promising solution to obtain an optimal policy for contact-rich tasks in recent years (Elguea-Aguinaco et al., 2023). To learn the optimal policy, a suitable reward function is significant in RL method. However, in many cases, it is engineering-consuming and unfeasible to design a comprehensive reward function, especially in some complex application scenarios, such as high-dimensional physical interactions and multi-objective manipulation tasks. The natural idea is that we can infer the reward function through expert demonstrations. In this section, we introduce the IRL method.

Different from RL learning an optimal policy via a pre-defined reward function by a trial-and-error paradigm (Martin-Martín et al., 2019; Zhang et al., 2024), IRL is an inverse process, which utilizes expert demonstration as the optimal policy to train a reward function that is a crucial part of RL training for generalization. Assuming that the expert demonstration is perfect, that is, the optimal solution under the optimal reward function and the reward function is optimized to make the expert policy obtain a higher reward value and then uses the RL algorithm to obtain the optimal policy under the latest reward function, and iterates this process until convergence.

However, there are challenges in IRL. Many optimal or sub-optimal policies can match the demonstrations and even more reward functions that can explain an optimal policy. Therefore, to optimize the reward functions, additional experience from the trial-and-error process is needed (Nair et al., 2017; Sermanet et al., 2017).

Recent work by Mandi et al. (2022b) introduced a novel multi-task training method that integrates a self-attention model and a temporal contrastive module to improve task disambiguation. Zhang et al. (2021b) learn both variable impedance policy and reward function from expert demonstrations based on IRL framework. Xu et al. (2022) proposed LION net, which only utilizes images as input to learning a task by RL-based control module.

Adversarial IL Adversarial IL (AIL) addresses a key challenge in IRL by resolving the ambiguity inherent in

inferring reward functions from expert demonstrations. Inverse RL typically suffers from the problem that multiple reward functions can explain the same expert behavior. By leveraging an adversarial framework, AIL incorporates a discriminator that distinguishes between expert and agent-generated trajectories. This mechanism forces the learned policy to produce behaviors that closely mimic the expert, effectively providing a robust and stable reward signal. Consequently, AIL reduces the dependency on hand-crafted rewards and enhances the scalability of training in complex, contact-rich tasks. AIL (Ho and Ermon, 2016) is a variant of IRL that leverages adversarial training to learn the reward function. The reward function is implicitly learned through an adversarial optimization process which is formulated as a minimax game, and the discriminator’s output is used to define the reward signal for the policy. AIL has been shown to be effective in learning complex manipulation tasks, such as pick-and-place operations and assembly tasks (Li and Zou, 2023).

Generative adversarial IL. (Ho and Ermon, 2016) is a variant of IRL that leverages adversarial training to learn the reward function. GAIL is based on the generative adversarial networks (GANs) (Goodfellow et al., 2014) framework, where the discriminator is trained to distinguish between expert demonstrations and generated trajectories, while the generator is trained to generate trajectories that are indistinguishable from expert demonstrations. The reward function is then learned by optimizing the discriminator to minimize the classification error. GAIL has been shown to be effective in learning complex manipulation tasks, such as pick-and-place operations and assembly tasks (Li and Zou, 2023).

Recent progress in GAIL for contact-rich tasks (Gubbi et al., 2020; Lee et al., 2022; Li et al., 2021; Xiang et al., 2024) has shown promising results. Specifically, Tsurumine et al. (2019) proposed a GAIL-based approach that incorporates contact information to improve the performance of robotic manipulation tasks. Recent advances in contact-rich manipulation have explored the use of unified frameworks for handling multiple subtasks. For example, as demonstrated in Xiang et al. (2024), expert demonstrations can be leveraged to train policies across different subtasks by sharing a single critic and an identical reward function. This strategy streamlines the learning process by providing consistent value feedback and simplifying reward design, ultimately promoting improved policy convergence and robustness. Li et al. (2021) translate human videos into practical robot demonstrations and train the meta-policy with adaptive loss based on the quality of the translated data. Robot demonstrations are not used but only human videos to train the meta-policy, facilitating data collection. Gubbi et al. (2020) achieve a peg-

in-hole insertion task with a 6 μm peg-hole clearance on the Yaskawa GP8 industrial robot.

Multimodal IL

Context-awareness is a key capability that distinguishes advanced robots acting in unstructured environments from traditional robots acting in controlled environments. It is especially important in contact-rich tasks as their success depends on timely and appropriate response to the highly dynamic contact conditions. Future robots should be able to understand their surroundings profoundly in order to adapt their behavior and answer changing conditions of tasks. Context-awareness is only possible through the inclusion of diverse modalities, since no single modality can cover the diversity of the tasks and conditions in unstructured environments.

A recent survey (Urain et al., 2024) discusses multimodal IL using deep generative models. They classify the recent works w.r.t. the types of generative models. However, the scope of this survey is different than ours as it focuses on multimodal deep generative models and does not focus on the contact-rich tasks. Here, we discuss the multimodal IL methods from the perspective of contact-rich interactions.

As discussed earlier in the *Data modalities section*, proprioceptive data constitutes the basis of robotic manipulation. Thus, it is used as the default input in most methods, with the exceptions of vision-based end-to-end approaches (Levine et al., 2016). Traditionally, the force modality has been a common choice in contact-rich tasks for its natural relevance (Siciliano and Villani, 1999). In the recent works, Stepputtis et al. (2022) use the force modality in addition to the robot and object position to identify the task phase. The phase variable synchronizes the learned behavior for different modalities. Some DMP-based works (Section 4.2) learn the force profile in parallel to the motion primitives either using another DMP (Chang et al., 2022) or a GMM model (Escarabajal et al., 2023). Some of them reformulate the DMP framework to include the force modality (Lödige et al., 2025; Qian et al., 2025). Osa et al. (2018) integrate online trajectory planning with force tracking control in a surgical task (Section 5.3). Liu et al. (2025b) combine force and vision as detailed later in this section. Luo et al. (2021) conduct a large-scale industrial task assessment of fusing force, proprioception, and vision modalities.

The tactile modality can provide more advanced information about the nature of the contact, such as surface friction, shape and curvature, or a spatial array of force readings. George et al. (2025) combine the vision modality with a tactile array with local shape information. They propose visuotactile contrastive pretraining (Rethmeier and Augenstein, 2023) for contact-rich tasks, and show that the multimodal pretraining improves the deployment time performance, even if the tactile encoder is removed after training. Lin et al. (2025) and Huang et al. (2024) separately

propose bimanual teleoperation systems, and study different factors affecting the visuotactile IL performance. [Ablett et al. \(2024\)](#) propose to learn the coupling of tactile and vision modalities through IL for better handling of contact mode switching and avoiding failures due to contact slipping.

Although direct force and tactile feedback are essentially useful for contact-rich tasks, it is still possible to achieve these tasks without them. Impedance control ([Neville, 1985](#)) has traditionally served as the primary approach for achieving indirect force control in robotic systems. [Zhao et al. \(2022a\)](#) use variable impedance control, optimizing the control parameters online to match the learned force profile. [Ugur and Girgin \(2020\)](#) add a compliance term into the DMP formulation for this purpose. [Solak and Jamone \(2019\)](#) use impedance control to apply grasping forces to keep an object in-hand while demonstrating and reproducing learned trajectories. Another way to modulate force is to rely on indirect sensing modalities such as vision; however, those approaches cannot provide precise force control.

The vision modality receives extensive attention in the related work both because of its capacity to capture spatial context and its readily available methods and datasets. Indeed, computer vision has been the forerunner of machine learning research. The language modality has also seen a huge leap lately with the development of the transformer models. Vision and language foundation models with outstanding capacity are getting available to aid robotic research, as discussed in the *Foundation models section*. Consequently, we see the trend also in multimodal contact-rich IL. [Chen et al. \(2024\)](#) combine visual user demonstrations with natural language instructions to learn both the reward and policy functions using a VLM model. [Mees et al. \(2022a\)](#) study the effects of different algorithmic and architectural decisions on IL through vision and language modalities. They evaluate various techniques on the simulation-based visuolingual manipulation benchmark CALVIN ([Mees et al., 2022b](#)). [Xian et al. \(2023\)](#) train two generative models: a transformer-based model for predicting high-level action keypoints, and a diffusion model for generating the trajectory segments between the keypoints. Both models have access to the inputs of vision, language and proprioception. [Shridhar et al. \(2023\)](#) add language processing and 3D voxel modality to the RL Bench environments ([James et al., 2020](#)) to evaluate their transformer-based behavior cloning agent. In the tasks where it is difficult to model the manipulated object, such as in cloth manipulation, the multi-sensory input becomes even more important. [Seita et al. \(2020\)](#) propose an IL system to solve fabric smoothing task through the RGB and depth modalities.

As a special case of vision modality, offline videos have a vast potential due to the ease of collection. The video modality carries valuable info on high-level plans to solve

long horizon tasks; however, it is hard to extract low-level control skills from it. This challenge is more important in the contact-rich tasks as the physical interactions are essential. For this reason, [Wang et al. \(2023\)](#) combine video modality and teleoperation-based proprioception modality: former for learning high-level plans, and latter for learning low-level control. They use videos of people freely playing in an environment to learn latent features about the possible actions, which are then employed to create plans to guide the low-level controller. [Iodice et al. \(2022\)](#) use the impedance control strategy, and estimate the intended arm stiffness of the human directly from video, based on the arm configuration. They then train GMMs to learn the configuration dependent stiffness (CDS) profiles of a sawing task, and reproduce it on a robotic setup.

Multimodality is a key to obtain general solutions for a diverse set of scenarios. We can categorize the works aiming to achieve diversity, or even generality into two groups: unified model and multi-model. The former aims to develop generalist robotic policies ([Bousmalis et al., 2024](#); [Ghosh et al., 2024](#); [O'Neill et al., 2024](#)) that aim to learn a single large model that can be easily fine-tuned to novel tasks, often employing multimodality in both task specification, model input and action-space. On the other hand, the separate model approach avoids the data complexity and data heterogeneity challenges by answering the subproblems through sub-models ([Ichiwara et al., 2023](#); [Wang et al., 2024d](#)). These sub-models can be combined in a hierarchical manner as [Ichiwara et al. \(2023\)](#) proposed. Their method combines modality-specific RNN models using a high-level RNN model. [Wang et al. \(2024d\)](#) handle generalization by sampling the joint probabilities of separate policies as discussed below.

Diffusion models are particularly suitable for multimodal learning due to being optimized at inference time. This aspect enables combining multiple diffusion models when imitating the learned skills. This capacity was first shown on image generation through the composition of separately trained models ([Liu et al., 2022a](#); [Nie et al., 2021](#)). Diffusion policies (DP) ([Chi et al., 2023](#)) formulate diffusion models to learn visuomotor robot policies. DPs are shown to handle the multimodal action distributions gracefully and achieve significant improvement over the state-of-the-art methods. Policy composition (PoCo) method ([Wang et al., 2024d](#)) combines DPs that are learned separately from many different modalities and domains, for different tasks. They achieve this by sampling the product distribution of multiple policies at the inference time. Inference time policy composition has the advantage of learning many small models, instead of a very large model like the generalist approaches. Also, combined DP sampling is more straightforward in comparison to merging RNN policies ([Wang et al., 2024c](#)), which requires combined optimization of the learned model parameters. [Yan et al. \(2024\)](#) make use of diffusion learning for combining vision, language and proprioception modalities; however, they use it for enhancing the

representation learning, rather than as an action policy. Liu et al. (2025b) adds the force modality in a DP. Their model combines the point cloud-based vision and proprioception modalities with force data.

The multimodal approaches to contact-rich IL aim to improve task performance through the inclusion of haptic sensing, extracting knowledge from more available modalities such as vision and language, and developing generalizable solutions to large variety of tasks. The modalities are combined in various ways such as using some modalities to select or modulate the learned skills, hierarchical use of modalities at different levels, learning joint probabilities of the modalities, extracting the task goals or costs, representation learning, pretraining, and more. The method depends on the used modalities and the tasks. The developments in the multimodal learning depend on the collection of different sensory data, and thus, highly benefit from the availability of public multimodal datasets and benchmarks (Please see the *Synthetic data generation and Available datasets* sections). These are also useful to evaluate and compare the diverse set of proposed methods.

Offline reinforcement learning

IL and RL are considered the most promising approaches for robot behavior acquisition. While IL is constrained by the performance of expert demonstrations, RL offers the potential to surpass expert-level skills. This section focuses on RL methods for learning contact-rich tasks. Although current RL algorithms often require large-scale datasets, posing a significant bottleneck, recent advances in robotic dataset collection suggest that both RL and IL will play a central role in achieving robust and scalable robotic behavior. Among RL algorithms, offline RL (Levine et al., 2020) is a prominent approach that seeks to solve tasks using previously collected datasets, without requiring additional interaction with the environment. In contrast, traditional online RL typically relies on extensive trial-and-error and continuous environment interaction to acquire new data. This requirement presents significant challenges in real-world robotic applications, where safety concerns such as hardware damage are critical. Offline RL addresses these issues by learning solely from pre-collected data, thereby avoiding the aforementioned risks. This section focuses on the use of offline RL for acquiring contact-rich manipulation skills.

Offline RL algorithms are often employed for pre-training on diverse, pre-collected datasets, enabling more efficient learning on downstream target tasks. Several approaches have demonstrated the effectiveness of leveraging such offline pre-training. Kumar et al. (2023) proposed pre-training for robots (PTRs), showing that combining diverse offline datasets with a small amount of target task data can significantly improve performance on new tasks. Their results highlight the potential of hybrid approaches that integrate broad prior experience with limited task-specific supervision. In another line of work, Bhateja et al. (2024)

improved learning efficiency by utilizing observation-only datasets that lack action and reward labels such as Ego4D (Grauman et al., 2022) for pre-training. This approach shows that rich visual or state information alone can provide valuable priors for downstream policy learning. These pre-training strategies are closely related to meta-reinforcement learning (Meta-RL), which aims to enable rapid adaptation to new tasks. For instance, Zhao et al. (2022b) proposed an offline Meta-RL framework that leverages demonstration adaptation to quickly adapt to novel tasks, bridging offline pre-training and meta-learning principles.

In addition to pre-training, fine-tuning is key for adapting models to target tasks, and various strategies have been proposed within the offline RL framework. A major challenge is the sim-to-real gap, as simulation data often lacks real-world noise and variability. Zhou et al. (2023) tackled this by using real-world data from safe executions of related tasks, enabling more robust transfer. Zhang et al. (2023) showed that combining offline pre-training via implicit Q-learning (IQL) with online soft actor-critic (SAC) fine-tuning improves performance on dexterous tasks like cherry-picking. However, Rafailov et al. (2023) noted that naïve fine-tuning can cause distribution shifts and instability, and proposed a model-based on-policy method to mitigate this. Similarly, Feng et al. (2023) introduced a fine-tuning approach using online data collection and a constraint balancing return estimates and model uncertainty, aiming for efficient yet safe real-world deployment.

Several studies have advanced learning with offline RL by exploring architectural innovations, data efficiency, and reward design. One notable direction incorporates Q-learning into transformer-based models initially designed for IL. For instance, Q-Transformer (Chebotar et al., 2023) merges transformer structures with temporal-difference learning, enabling sequence modeling of actions in offline settings. Other work has integrated goal conditioning and affordance models to provide structured priors or constraints that guide policy learning in complex environments (Fang et al., 2023). Dong et al. (2024) proposed a two-stage framework that improves policy learning from limited data by decoupling the learning process, enhancing both sample efficiency and generalization. To mitigate distributional shift, policy constraints and conservative learning methods are common, but often rely on complex approximations in continuous action spaces. Luo et al. (2023) addressed this by introducing state-conditioned action quantization, offering a discrete, state-aware approximation. Reward specification remains critical in offline RL. Liu et al. (2023b) proposed using a small amount of expert data to drive learning via intrinsic rewards, reducing or even eliminating the need for explicit extrinsic rewards.

Other methods

In addition to the common approaches discussed above, several other methods have been proposed to tackle unique

challenges in this area, offering alternative perspectives and solutions.

World models offer a promising avenue for IL in contact-rich tasks by enabling robots to learn predictive representations of their environment (Wu et al., 2023). These models, trained on demonstration data, can forecast the outcomes of actions, allowing robots to plan and execute complex manipulations more effectively (Chen et al., 2025a). By capturing the underlying dynamics of contact interactions, world models facilitate the acquisition of robust policies that generalize well to novel situations, addressing key challenges in contact-rich IL such as data scarcity and sim-to-real transfer (Barcellona et al., 2025).

Zero or One-shot IL: such as works (Bonardi et al., 2020; Duan et al., 2017; Lázaro-Gredilla et al., 2019) that enable agents to generalize from minimal or even no task-specific demonstrations, leveraging meta-learning, transfer learning, or compositional reasoning to adapt quickly to unseen tasks. Unlike traditional BC or IRL, which often require extensive demonstrations, zero/one-shot IL focuses on extreme generalization, where policies must infer correct behavior from just one example (one-shot) or none at all (zero-shot) by relying on prior knowledge or task embeddings. This paradigm bridges the gap between IL and few-shot RL, offering a complementary perspective to offline RL and generative methods by prioritizing data efficiency at the expense of upfront training complexity.

IL on Riemannian Manifolds extends traditional methods to non-Euclidean spaces, where actions or states inherently lie on smooth, curved geometries. Work (Zeebstra et al., 2017) leverages geometric priors to ensure stable and physically plausible policy learning. Manifold-aware IL avoids distortions caused by Euclidean approximations, improving performance in contact-rich tasks.

Several studies focus on learning from limited or sub-optimal demonstrations, addressing scenarios where datasets may be sparse or exhibit noise, biases, or non-expert trajectories (Kim et al., 2021). These methods relax the assumption of high-quality, large-scale data required by standard BC or IRL. For small datasets, techniques like data augmentation (Johns, 2021), meta-learning, or hybrid RL/IL frameworks are employed to extract robust policies despite limited supervision, often leveraging coarse-grained abstractions or hierarchical strategies to compensate for missing details. These approaches are critical in real-world settings in which collecting expert-quality data is expensive or impractical, though they may require careful trade-offs between generalization and fidelity to the demonstrations.

Recent advances have explored integrating unsupervised learning with IL to reduce dependency on labeled demonstrations. By leveraging unlabeled data, these methods can learn useful representations and dynamics models that enhance policy learning efficiency and generalization. Techniques like LOTUS (Wan et al., 2024) combine self-supervised representation learning with imitation policies,

enabling agents to extract meaningful patterns from uncurated data before fine-tuning on limited demonstrations. This paradigm offers significant advantages in real-world scenarios where labeled demonstrations are scarce but unlabeled interaction data is abundant. The unsupervised components can learn robust feature spaces and dynamics that make subsequent IL more sample-efficient and adaptable. While sharing some goals with offline RL and generative approaches, unsupervised IL methods uniquely focus on extracting value from completely unlabeled data. This positions them as particularly valuable for scaling IL to complex, open-ended environments.

An emerging direction in IL incorporates explicit optimization objectives into the learning process, blending traditional imitation approaches with mathematical optimization techniques to create more robust and adaptable policies. As demonstrated by Okada et al. (2023), these methods often formulate policy learning as a bilevel optimization problem or integrate constrained optimization layers directly into neural networks, enabling policies to satisfy physical and logical constraints while imitating demonstrations. While sharing some conceptual ground with DMPs and model-based RL, optimization-based IL methods distinguish themselves through their formal mathematical guarantees and explicit constraint satisfaction mechanisms.

Some other methods such as learning sequential structure (Tanwani et al., 2021), intervention learning (Korkmaz and Büyük, 2025), automatic segmentation of the subtasks (Mao et al., 2024; Sugawara et al., 2023), probabilistic activity grammars (Lee et al., 2013), and skill retrieval and adaptation (Guo et al., 2025; Memmel et al., 2025) offer unique approaches to address specific challenges in contact-rich IL, complementing the primary techniques discussed above.

Application cases

This section provides an overview of the applications of IL in contact-rich tasks. Specifically, we present them in the *Industrial robots*, *Service and household robots*, and *Healthcare robots* sections. Figure 8 shows real-robot examples of these application cases.¹ In addition, the last section discusses the *selection of algorithms* according to the available sensors and the tasks to be adapted when applying imitation learning.

Industrial robots

IL has emerged as a powerful framework for enabling industrial robots to acquire complex contact-rich skills efficiently from expert demonstrations. By imitating human expertise, IL reduces the need for extensive manual programming and facilitates rapid adaptation to varying tasks. This approach enhances precision and robustness in operations such as assembly (Scherzinger et al., 2019), insertion




Figure 8. Examples of industrial, household, and healthcare robotics applications: (a) Peg insertion is a common operation in industrial assembly tasks (reproduced from Wang et al. (2022b)*); (b) household robotic companion ADAM adopts dual arms design with a mobile base (adapted from Mora et al. (2024)*); (c) human–robot hand over is a fundamental skill for household and assistive healthcare robots (adapted from Jadeja et al. (2025)*); (d) the robotic setup for learning the surgical bone-grinding task from demonstration (reproduced from Li et al. (2024a)*).

(Wang et al., 2022b), and pick-and-place operations (Li and Zou, 2023), ultimately leading to more flexible and autonomous robotic systems. These tasks, which involve intricate interactions between the robot and its environment, pose significant challenges due to uncertainties in contact dynamics and the need for precise force and motion control. Industrial robots, which traditionally rely on programmed control policies, can greatly benefit from IL techniques to achieve greater adaptability and autonomy (Liu et al., 2022b).

Industrial scenarios are the main application of contact-rich tasks. Such as assembly (peg-in-hole, insertion) (Zhang et al., 2021b), deburring grinding (Onstein et al., 2020), polishing and sanding (Zeng et al., 2023), and deformable object manipulation (Salhotra et al., 2022). Recent advances in IL have enabled industrial robots to learn complex manipulation tasks involving multiple contact transitions. For example, robots have been trained to perform pick-and-place operations with varying object shapes and sizes, demonstrating the ability to adapt to different scenarios (Zhou et al., 2025). IL has also been applied to assembly tasks, where robots learn to assemble components with high precision and efficiency (Wang et al., 2024b). These developments highlight IL’s potential to enhance industrial robots’ capabilities in contact-rich environments (Abu-Dakka and Saveriano, 2020)

Despite advancements in robotic control, several key challenges persist in these applications. These include achieving precise force regulation amid variable contact conditions, integrating real-time reflexive control with sensory feedback (Van Duong, 2025; Zhang et al., 2025c), and ensuring stability while maintaining safety (Hejrati and Mattila, 2023; Zhang et al., 2025d). Addressing these challenges is critical for enhancing the robustness and autonomy of industrial robots in contact-rich tasks.

Service and household robots

Service and household robots have seen limited usage in commercial applications due to the challenges of physically

interacting in unstructured environments with untrained end-users. Most of the existing commercial household robots focus on a specific task such as surface cleaning, lawn-mowing, or window cleaning (Zachiotis et al., 2018), with the exception of Care-O-bot, which is designed to handle various tasks like tool-use and interaction with kitchen appliances using its two arms (Kittmann et al., 2015). Currently commercialized household robots such as the educational, entertainment, social, and toy robots do not require physical interaction (Zachiotis et al., 2018). Thus, we review mainly the non-commercialized research in the following.

One of the main requirements of the service and household robotics is to interact with the end-user. This requirement entails multiple challenges such as the social acceptance, safe human–robot interaction and intuitive interfaces. Both physical and perceived safety hold a central place in contact-based HRI as reviewed by Farajtabar and Charbonneau (2024). Thus, the robots must endow enhanced context-awareness to predict the human intentions and ensure their safety (Li et al., 2022b). Furthermore, providing easy interaction modalities is desirable, even in professional service robots, whose users may have former training (Gonzalez-Aguirre et al., 2021). For this reason, the natural language modality is fundamental for service and household robots. Furthermore, it is desirable for the robot to understand the human physical and emotional states through non-verbal means. Thus, the multimodal methods allowing language-based or image-based goal conditioning are important for the service robots.

The second requirement of physically interacting with unstructured environments is more relevant for our review. Typical contact-rich household activities include maintenance activities such as table wiping (Yang et al., 2022; Zhao et al., 2022a), cloth folding and smoothing (Hoque et al., 2021; Seita et al., 2020; Xiong et al., 2023); interacting with the household objects such as doors (Ablett et al., 2024; Bharadhwaj et al., 2023), drawers (Liu et al., 2024; Yan et al., 2024), kitchen appliances (Mandi et al., 2022a; Wang et al., 2023), and tools (Wang et al., 2024c, 2024d); and kitchen tasks such as fruit or vegetable cutting (Liu et al., 2025b), pouring liquids (Zhang et al., 2021b), and food stirring (Yang et al., 2022). Currently, these tasks are solved using general-purpose mobile robots and manipulators in laboratory settings. Solving these tasks in a long-term cost-effective manner continues to be a key research goal. Due to the abundance of challenging tasks in household environments, it remains as a major motivation for advanced robotics research. Accordingly, existing datasets and benchmarks often include kitchen (Li et al., 2022a; Xiong et al., 2023), tool-use (Wang et al., 2024c), or general household tasks (Fang et al., 2024; James et al., 2020; Yu et al., 2020).

Unlike industrial robots that deal with relatively constrained environments and tasks, the service robots should

support a wider variety of actions. Thus, household robots aiming for generality may benefit from high-DoF dexterous manipulators like dual arms (Huang et al., 2024; Zhao et al., 2023a), multi-fingered hands (Shaw et al., 2024), or both (Lin et al., 2025; Wang et al., 2024a). However, such systems are usually too costly (An et al., 2025) to be made available for personal use. Developing cost-effective dexterous manipulation systems stands as a challenge for contact-rich household applications.

Lastly, health-care robotic applications like monitoring, assisting and rehabilitating robots can be deployed in houses of the people in need, in addition to the clinics (Halicka and Surel, 2022; Yang et al., 2020a). Some of the potential contact-rich tasks in houses are mobility support, rehabilitation exercises, and activities of daily living, such as feeding, dressing and personal hygiene assistance (Yang et al., 2020a). Health-care robotic applications are covered in more detail in the next section.

Health-care robots

One of the areas in the healthcare industry where automation is highly anticipated is surgical procedures. Given that surgery requires advanced skills and significant physical endurance, automation through robotic systems (Osa et al., 2010) is expected to offer substantial benefits. A clinical study has suggested that surgeons can benefit from the assistance of robotic systems in surgical procedures (Nix et al., 2010). Surgical operations involve the manipulation of soft tissues and organs, which necessitates a high level of dexterity. Osa et al. (2018) have developed a technique that mimics not only positional trajectories but also force trajectories based on expert human data. They demonstrated this approach using a dual-arm robotic system for the knot-tightening task. In surgical procedures, both video and kinematic data are recorded for postoperative analysis, resulting in the establishment of a large repository of empirical data. With this in mind, Kim et al. (2024a) employed IL using visual and kinematic data to enable a dual-arm robotic system to robustly perform surgical tasks. They have introduced Surgical Robot Transformer (SRT) which the ACT is integrated into the Da Vinci Research Kit (dVRK) (Cui et al., 2023) and successfully executed various surgical tasks, including tissue manipulation, needle handling, and knot-tying. Moghani et al. (2025) propose a photorealistic surgery simulator to generate the safety-critic and costly data for the evaluation of surgical systems.

Another area in healthcare where automation is highly anticipated is rehabilitation. Patients recovering from conditions such as stroke (Langhorne et al., 2011) or Parkinson's disease (Abbruzzese et al., 2016) require training to restore or maintain their physical functions. In most cases, treatment is administered by therapists, and due to the inherent complexity of working with human patients, automation through robotic systems presents significant challenges.

Robotic-assisted rehabilitation exercises (Batson et al., 2020; Kato et al., 2024; Krebs et al., 1998) involving physical human-robot interaction require careful consideration due to the physical contact between the patient and the robot, making compliant behavior imperative for these tasks. Escarabajal et al. (2023) have developed a method for generating compliant trajectories for passive rehabilitation exercises, taking into account that previous positions along the trajectory are attainable for the patient. Their approach is based on IL, encoding forces using Gaussian mixture regression (GMR), and employing Reversible DMPs. The system enables self-paced rehabilitation exercises through back-and-forth movements along the trajectory in response to the patient's reactions. Lim et al. (2023) investigate the feasibility of using a general-purpose collaborative robot for rehabilitation therapies. IL methods were employed to replicate expert-provided training trajectories that can adapt to the subject's capabilities, facilitating in-home rehabilitation training. Their approach incorporates the concept of HG-Dagger (Kelly et al., 2019), allowing human intervention to ensure that patients do not attempt trajectories that may be difficult to execute. By integrating IL with a system that permits human intervention, this method is expected to be beneficial for in-home rehabilitation.

Selection of appropriate algorithms

The selection of appropriate IL algorithms for contact-rich manipulation fundamentally depends on the interplay between available sensory modalities and task-specific characteristics. This section provides a systematic framework for matching algorithmic approaches to practical deployment scenarios.

Sensor-algorithm compatibility. The availability and quality of sensory feedback profoundly influence the choice of IL architecture. Table 1 summarizes the recommended algorithmic approaches for different sensor configurations. When high-fidelity force and torque sensing is available, algorithms can directly leverage contact mechanics information during policy learning. BC with force features as explicit state variables enables policies to reason about interaction forces, while Dynamic Movement Primitives can be augmented with force coupling terms that modulate trajectory execution based on measured contact forces.

In vision-only scenarios, contact events become implicit in visual observations, necessitating temporal convolutions or recurrent architectures that can infer contact from motion changes. Generative models such as diffusion policies become particularly valuable in handling multimodal action distributions commonly observed in contact-rich tasks. Tactile sensing provides fine-grained contact geometry and slip detection, enabling hybrid approaches that combine global visual context with local tactile feedback. When only proprioceptive information is available, structured representations such as Dynamic Movement Primitives offer reasonable generalization despite limited observability.

Table I. Algorithm selection by sensor modality.

Sensor modality	Recommended approaches	Key advantages	Limitations
Force/Torque	BC with force features, force-coupled DMPs, IRL	Direct contact observation, explicit force reasoning	Requires expensive sensors
Vision (RGB/RGB-D)	Visual BC, state representation learning, diffusion policies	Rich semantic information, scalable	Implicit contact, occlusion issues
Tactile	Conditioned BC, hybrid vision-tactile	Fine-grained contact geometry, slip detection	Limited sensing area, sensor cost
Proprioception only	DMPs, kinesthetic teaching, model-based IL	Simple, reliable	Limited generalization to new objects

Task characteristics and algorithmic requirements. Beyond sensor availability, specific task characteristics impose additional constraints on algorithm selection. High-precision tasks such as peg-in-hole insertion benefit from residual learning architectures that combine nominal geometric controllers with learned corrective policies, providing the structured reasoning that pure IL often lacks. Tasks involving complex contact sequences require hierarchical decomposition, where high-level task segmentation and low-level motor skills are learned separately to reduce combinatorial complexity.

Variable compliance requirements favor impedance learning approaches that directly parameterize stiffness and damping from demonstrations, or energy-based models that implicitly encode compliant behaviors. Multimodal behaviors, where multiple valid strategies exist, require careful algorithmic consideration. Standard BC with mean squared error loss performs poorly in such scenarios, as it averages over modes and produces invalid intermediate actions. Diffusion policies that model the full action distribution or mixture density networks with explicit multimodal representation provide more robust solutions.

Computational and real-time constraints. Real-time constraints impose practical limitations on algorithm complexity. Applications requiring strict real-time performance with control frequencies exceeding 100 Hz favor computationally efficient approaches such as Dynamic Movement Primitives, which offer closed-form solutions with minimal computational overhead. Linear policies and quantized neural networks provide fast inference at the cost of representational capacity. When timing constraints are relaxed, permitting inference times of 50 milliseconds or more, deep networks including transformers and model-predictive control frameworks become viable despite their computational demands.

Conclusion

This paper comprehensively surveys research trends in IL for contact-rich tasks. These tasks require complex physical interactions with the environment, representing a central challenge in robotics due to their nonlinearity and

complexity. Understanding everyday physics has long been recognized as difficult, particularly in contact-rich tasks where slight positional deviations cause significant behavioral changes. Against this background, IL approaches learning from human demonstrations have attracted considerable attention.

The paper analyzes the main IL approaches, including BC, DMP, Generative methods, and IRL. Recent developments highlight how the latest generative models achieve excellent performance even in complex tasks. Each method has distinct strengths, requiring appropriate selection based on task nature and available data. Implementations are progressing in industrial robots (assembly and picking), household robots (everyday manipulation), and medical robots (surgical and physical therapy support).

Core technical challenges

This section identifies fundamental technical challenges commonly encountered in IL applications for contact-rich manipulation. These challenges represent current research gaps that limit the practical deployment and performance of many IL systems. Addressing these challenges is essential for advancing the field and enabling more robust and capable contact-rich manipulation systems.

Design theory for hierarchical architectures. Research on hierarchical architectures, which has been increasing in recent years (Belkhalé et al., 2024; Li et al., 2025b; Xue et al., 2025), is expected to become increasingly important in IL because it enables effective learning and reproduction of both fast reactive control and high-level planning inherent in human demonstrations, significantly improving adaptability and decision-making capabilities especially in contact-rich tasks (Yamashita and Tani, 2008). While various models have been proposed as mentioned above, a unified design principle has not yet been established. The concept of dual-process theory from cognitive science—distinguishing between System 1 (fast, intuitive processing) and System 2 (slow, analytical processing)—is gaining traction in machine learning and represents a promising candidate as a

design principle for hierarchical architectures (Bengio, 2017). As LLMs continue to develop, expectations for System 2 models capable of handling advanced logic are growing, while contact-rich tasks particularly require high-performance System 1 models to process environmental interactions (Tsuji, 2025).

Multimodal sensing. The success of IL requires appropriate selection and integration of data modalities. Multimodal learning that integrates multiple sensory information—not only position data but also force, vision, and tactile information—significantly contributes to improving performance in contact tasks (Ravichandar et al., 2020; Urain et al., 2024). The development of tactile sensing, in particular, enables detection and interpretation of subtle contact phenomena, and the advancement of its hardware and software, along with integration with other sensory information, is key to achieving more delicate manipulation (Ablett et al., 2024; Edmonds et al., 2017; Higuera et al., 2024; Lambeta et al., 2024).

Bridging the gap between simulation and the real world. The transfer from simulation to real machines (sim-to-real) remains an important challenge. Due to the difficulty of modeling contact dynamics, policies learned in simulation often fail to function well in the real world (Peng et al., 2018; Tobin et al., 2017). Various approaches have been proposed to address this problem, including domain randomization (Tobin et al., 2017), physics-based augmentation (Zeng et al., 2020), differentiable simulation (Freeman et al., 2021b), and hybrid dataset utilization combining simulated and real data (Chebotar et al., 2019). While contact-rich RL faces challenges from limited real-world trials due to safety and hardware concerns (Elguea-Aguinaco et al., 2023), IL encounters even stricter constraints as it depends on costly expert demonstrations rather than autonomous exploration. Further development in this area is essential to overcome the challenges of limited demonstration data in contact-rich tasks.

The development of foundation models is likely to influence all three challenges mentioned above, and further research progress is anticipated. To achieve more versatile and adaptive contact-rich tasks, solutions to these technical challenges and the development of integrated approaches will be necessary.

Future directions

While the previous section addresses fundamental challenges commonly encountered in IL systems for contact-rich tasks, this section identifies emerging research areas and specialized topics that will become increasingly important in the coming years. These directions represent opportunities for expanding the scope and applicability of contact-rich IL beyond current mainstream research, opening new avenues for both theoretical advances and practical applications. Each topic deserves dedicated investigation and has the

potential to significantly broaden the impact of IL in robotics.

Safe learning. Safety and robustness are important aspects of the contact-rich manipulation problem. Even though the *safe learning* topic has received increasing attention in the recent years (Brunke et al., 2022; Gu et al., 2024; Zhao et al., 2023b), contact-rich safe learning research is still at an early stage.

Physical human–robot interaction. With the recent advances in learning, control and design methods, the robots become more adequate to interact with humans in less constrained environments. Thus, the importance of *physical HRI* will keep increasing in the future (Farajtabar and Charbonneau, 2024; Li et al., 2022b).

Deformable object and soft robot manipulation. Mainstream robotics research has mostly focused on rigid objects, environments and robots. However, *deformable object manipulation* (Yin et al., 2021; Zhu et al., 2022) stands as a key unsolved problem for real-world robotics applications. Furthermore, *soft robots* carry the inherent advantages of adaptability and compliance (Yasa et al., 2023), which are essential for contact-rich tasks. Notably, the difficulty of modeling deformable objects, environments and robots render IL an invaluable tool for these tasks.

Dexterous manipulation. *Dexterous robot hands* (An et al., 2025; Welte and Rayyes, 2025) bring valued advantages for contact-rich applications, such as flexibility and precision. However, dexterous manipulation requires its own literature surveys as it brings particular challenges such as grasp stability, high-dimensional control and tactile sensing.

ORCID iDs

Toshiaki Tsuji  <https://orcid.org/0000-0002-4532-4514>

Yasuhiro Kato  <https://orcid.org/0000-0001-8343-5314>

Gokhan Solak  <https://orcid.org/0000-0001-6342-1345>

Heng Zhang  <https://orcid.org/0000-0003-4832-9668>

Funding

The authors disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported in part by JSPS KAKENHI (JP25K01191), the Social Collaboration Program (Scalable Robot Learning) between The University of Tokyo and Honda R&D Co., Ltd., the Slovenian Research Agency (ARIS) (N2-0269), and the European Union Horizon Europe project TORNADO (GA 101189557).

Declaration of Conflicting Interests

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

Note

1. The images are used under CC-BY and CC-BY-ND licenses.

References

- Abbruzzese G, Marchese R, Avanzino L, et al. (2016) Rehabilitation for parkinson's disease: current outlook and future challenges. *Parkinsonism & Related Disorders* 22: S60–S64.
- Ablett T, Limoyo O, Sigal A, et al. (2024) Multimodal and force-matched imitation learning with a see-through visuotactile sensor. *IEEE Transactions on Robotics*.
- Abolghasemi P, Mazaheri A, Shah M, et al. (2019) Pay attention!-robustifying a deep visuomotor policy through task-focused visual attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15-20 June 2019, 4254–4262.
- Abu-Dakka FJ and Kyrki V (2020) Geometry-aware dynamic movement primitives. In: 2020 IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 31 May 2020 - 31 August 2020. IEEE, 4421–4426.
- Abu-Dakka FJ and Saveriano M (2020) Variable impedance control and learning—A review. *Frontiers in Robotics and AI* 7: 590681.
- Abu-Dakka FJ, Saveriano M and Kyrki V (2024) A unified formulation of geometry-aware discrete dynamic movement primitives. *Neurocomputing* 598: 128056.
- Adachi T, Fujimoto K, Sakaino S, et al. (2018) Imitation learning for object manipulation based on position/force information using bilateral control. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Madrid, Spain, 01-05 October 2018. IEEE, 3648–3653.
- Aggarwal K, Singh SK, Chopra M, et al. (2022) *Deep Learning in Robotics for Strengthening Industry 4.0.: Opportunities, Challenges and Future Directions*. Springer International Publishing, 1–19. DOI: [10.1007/978-3-030-96737-6_1](https://doi.org/10.1007/978-3-030-96737-6_1).
- Ajouadani A, Tsagarakis N and Bicchi A (2012) Tele-impedance: teleoperation with impedance regulation using a body-machine interface. *The International Journal of Robotics Research* 31(13): 1642–1656.
- Aleotti J, Caselli S and Reggiani M (2003) Toward programming of assembly tasks by demonstration in virtual environments. In: The 12th IEEE International Workshop on Robot and Human Interactive Communication. *Proceedings. ROMAN 2003*. pp. 309–314. DOI: [10.1109/ROMAN.2003.1251863](https://doi.org/10.1109/ROMAN.2003.1251863).
- An S, Meng Z, Tang C, et al. (2025) Dexterous manipulation through imitation learning: a survey. *arXiv preprint arXiv:2504.03515*.
- Ankile L, Simeonov A, Shenfeld I, et al. (2024a) Juicer: Data-efficient imitation learning for robotic assembly. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, United Arab Emirates, 14-18 October 2024. IEEE, 5096–5103.
- Ankile LL, Simeonov A, Shenfeld I, et al. (2024b) From imitation to refinement-residual rl for precise visual assembly. In: CoRL 2024 Workshop on Mastering Robot Manipulation in a World of Abundant Data, Munich, Germany, November 6 - 9, 2024.
- Argall BD, Chernova S, Veloso M, et al. (2009) A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57(5): 469–483.
- Arimoto S (1999) Robotics research toward explication of everyday physics. *The International Journal of Robotics Research* 18(11): 1056–1063.
- Bahl S, Mukadam M, Gupta A, et al. (2020) Neural dynamic policies for end-to-end sensorimotor learning. *Advances in Neural Information Processing Systems* 33: 5058–5069.
- Barcellona L, Zadaianchuk A, Allegro D, et al. (2025) Dream to manipulate: compositional world models empowering robot imitation learning with imagination. In: The Thirteenth International Conference on Learning Representations.
- Barreiros JA, Önl AÖ, Zhang M, et al. (2025) Learning contact-rich whole-body manipulation with example-guided reinforcement learning. *Science Robotics* 10(105): eads6790.
- Batson JP, Kato Y, Shuster K, et al. (2020) Haptic coupling in dyads improves motor learning in a simple force field. In: 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), Montreal, QC, Canada, 20-24 July 2020. IEEE, 4795–4798.
- Battaglia E, Bianchi M, Altobelli A, et al. (2015) Thimblesense: a fingertip-wearable tactile sensor for grasp analysis. *IEEE transactions on haptics* 9(1): 121–133.
- Belkhale S, Ding T, Xiao T, et al. (2024) Rt-h: action hierarchies using language. *arXiv preprint arXiv:2403.01823*.
- Beltran-Hernandez CC, Petit D, Ramirez-Alpizar IG, et al. (2020) Variable compliance control for robotic peg-in-hole assembly: a deep-reinforcement-learning approach. *Applied Sciences* 10(19): 6923.
- Bengio Y (2017) The consciousness prior. *arXiv preprint arXiv:1709.08568*.
- Bharadhwaj H, Gupta A, Tulsiani S, et al. (2023) Zero-shot robot manipulation from passive human videos. *arXiv preprint arXiv:2302.02011*.
- Bhateja C, Guo D, Ghosh D, et al. (2024) Robotic offline rl from internet videos via value-function learning. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), 13-17 May 2024, 16977–16984.
- Bicchi A and Siciliano B (1993) Closure properties of robotic manipulation. *The International Journal of Robotics Research* 12(2): 122–137.
- Black K, Brown N, Darpinian J, et al. (2025a) $\pi_{0.5}$: a vision-language-action model with open-world generalization. In: 9th Annual Conference on Robot Learning, Seoul, Republic of Korea.
- Black K, Brown N, Driess D, et al. (2025b) π_0 : a vision-language-action flow model for general robot control. In: Proceedings of Robotics: Science and Systems, Los Angeles, CA, USA. DOI: [10.15607/RSS.2025.XXI.010](https://doi.org/10.15607/RSS.2025.XXI.010).
- Bonardi A, James S and Davison AJ (2020) Learning one-shot imitation from humans without humans. *IEEE Robotics and Automation Letters* 5(2): 3533–3539.
- Bousmalis K, Vezzani G, Rao D, et al. (2024) Robocat: a self-improving generalist agent for robotic manipulation. *Transactions on Machine Learning Research*.

- Brohan A, Brown N, Carbajal J, et al. (2023) RT-1: robotics transformer for real-world control at scale. In: Proceedings of Robotics: Science and Systems. Daegu, Republic of Korea, Daegu, Republic of Korea, Jul 10 – Jul 14, 2023. DOI: [10.15607/RSS.2023.XIX.025](https://doi.org/10.15607/RSS.2023.XIX.025).
- Brunke L, Greeff M, Hall AW, et al. (2022) Safe learning in robotics: from learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics, and Autonomous Systems* 5(1): 411–444.
- Buamane T, Kobayashi M, Uranishi Y, et al. (2024) Bi-act: bilateral control-based imitation learning via action chunking with transformer. In: 2024 IEEE International Conference on Advanced Intelligent Mechatronics (AIM), Boston, MA, USA, 15-19 July 2024, 410–415. DOI: [10.1109/AIM55361.2024.10637173](https://doi.org/10.1109/AIM55361.2024.10637173).
- Burnwal R, Mehta H, Bhatt NP, et al. (2025) Learning from observation: a survey of recent advances. *arXiv preprint arXiv: 2509.19379*.
- Büscher G, Koiva R, Schürmann C, et al. (2012) Tactile dataglove with fabric-based sensors. In: 2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012), Osaka, Japan, 29 November 2012 - 01 December 2012. IEEE, 204–209.
- Calinon S, D’halluin F, Sauter EL, et al. (2010) Learning and reproduction of gestures by imitation. *IEEE Robotics and Automation Magazine* 17(2): 44–54.
- Campbell J, Stepputtis S and Ben Amor H (2019) Probabilistic multimodal modeling for human-robot interaction tasks. *Robotics: Science and Systems XV*.
- Celemin C, Pérez-Dattari R, Chisari E, et al. (2022) Interactive imitation learning in robotics: a survey. *Foundations and Trends® in Robotics* 10(1–2): 1–197.
- Chang C, Haninger K, Shi Y, et al. (2022) Impedance adaptation by reinforcement learning with contact dynamic movement primitives. In: 2022 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), Sapporo, Japan, 11-15 July 2022. IEEE, 1185–1191.
- Chebotar Y, Handa A, Makoviychuk V, et al. (2019) Closing the sim-to-real loop: adapting simulation randomization with real world experience. In: 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20-24 May 2019, 8973–8979.
- Chebotar Y, Vuong Q, Hausman K, et al. (2023) Q-transformer: scalable offline reinforcement learning via autoregressive q-functions. *Proceedings of The 7th Conference on Robot Learning* 229: 3909–3928.
- Chen L, Moorman NM and Gombolay MC (2024) Elemental: interactive learning from demonstrations and vision-language models for reward design in robotics. In: Forty-Second International Conference on Machine Learning.
- Chen H, Xu J, Sheng L, et al. (2025a) Learning coordinated bi-manual manipulation policies using state diffusion and inverse dynamics models. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), 5644–5651. DOI: [10.1109/ICRA55743.2025.11127387](https://doi.org/10.1109/ICRA55743.2025.11127387).
- Chen J, Yu C, Zhou X, et al. (2025b) EMOS: Embodiment-aware heterogeneous multi-robot operating system with LLM agents. In: The Thirteenth International Conference on Learning Representations.
- Cheng N, Xu J, Guan C, et al. (2025) Touch100k: a large-scale touch-language-vision dataset for touch-centric multimodal representation. *Information Fusion* 124: 103305.
- Chi C, Xu Z, Feng S, et al. (2023) Diffusion policy: visuomotor policy learning via action diffusion. *The International Journal of Robotics Research* 2024;44(10-11):1684-1704. DOI: [10.1177/02783649241273668](https://doi.org/10.1177/02783649241273668)
- Cho NJ, Lee SH, Kim JB, et al. (2020) Learning, improving, and generalizing motor skills for the peg-in-hole tasks based on imitation learning and self-learning. *Applied Sciences* 10(8): 2719.
- Cong L, Liang H, Ruppel P, et al. (2022) Reinforcement learning with vision-proprioception model for robot planar pushing. *Frontiers in Neurobotics* 16: 829437.
- Coumans E and Bai Y (2016) Pybullet, a python module for physics simulation for games. *Robotics and Machine Learning*. <https://pybullet.org>
- Cui J and Trinkle J (2021) Toward next-generation learned robot manipulation. *Science Robotics* 6(54): eabd9461.
- Cui Z, Cartucho J, Giannarou S, et al. (2023) Caveats on the first-generation da vinci research kit: latent technical constraints and essential calibrations. *IEEE Robotics and Automation Magazine*.
- Cutkosky MR and Kao I (1989) Computing and controlling compliance of a robotic hand. *IEEE Transactions on Robotics and Automation* 5(2): 151–165.
- Dalle Vedove M, Abu-Dakka FJ, Palopoli L, et al. (2025) Meshdmp: motion planning on discrete manifolds using dynamic movement primitives. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025. IEEE, 895–901.
- Davchev T, Luck KS, Burke M, et al. (2022) Residual learning from demonstration: adapting dmps for contact-rich manipulation. *IEEE Robotics and Automation Letters* 7(2): 4488–4495.
- Degrave J, Hermans M, Dambre J, et al. (2019) A differentiable physics engine for deep learning in robotics. *Frontiers in Neurobotics* 13:6.
- Deng S, Yan M, Wei S, et al. (2025) GraspVLA: a grasping foundation model pre-trained on billion-scale synthetic action data. In: 9th Annual Conference on Robot Learning, Seoul, Republic of Korea.
- Deniša M, Gams A, Ude A, et al. (2015) Learning compliant movement primitives through demonstration and statistical generalization. *IEEE/ASME Transactions on Mechatronics* 21(5): 2581–2594.
- Deshpande A, Ke L, Pfeifer Q, et al. (2024) Data efficient behavior cloning for fine manipulation via continuity-based corrective labels. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, United Arab Emirates, 14-18 October 2024, 8531–8538.
- Dettmers T, Pagnoni A, Holtzman A, et al. (2023) Qlora: efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems* 36: 10088–10115.

- Dillmann R (2004) Teaching and learning of robot tasks via observation of human performance. *Robotics and Autonomous Systems* 47(2–3): 109–116.
- Dong Q, Kaneko T and Sugiyama M (2024) An offline learning of behavior correction policy for vision-based robotic manipulation. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 13-17 May 2024, 5448–5454.
- Driess D, Xia F, Sajjadi MS, et al. (2023) Palm-e: an embodied multimodal language model. In: International Conference on Machine Learning, Honolulu, HI, USA. PMLR, 8469–8488.
- Drolet M, Stepputtis S, Kailas S, et al. (2024) A comparison of imitation learning algorithms for bimanual manipulation. *IEEE Robotics and Automation Letters*.
- Duan Y, Andrychowicz M, Stadie B, et al. (2017) One-shot imitation learning. *Advances in Neural Information Processing Systems*. Long Beach, CA, USA: Curran Associates, Inc, 30, 1087–1098.
- Duan J, Pumacay W, Kumar N, et al. (2025) AHA: a vision-language-model for detecting and reasoning over failures in robotic manipulation. In: The Thirteenth International Conference on Learning Representations, Singapore.
- Edmonds M, Gao F, Xie X, et al. (2017) Feeling the force: integrating force and pose for fluent discovery through imitation learning to open medicine bottles. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, Canada, 24-28 September 2017. IEEE, 3530–3537.
- Elguea-Aguinaco Í, Serrano-Muñoz A, Chrysostomou D, et al. (2023) A review on reinforcement learning for contact-rich robotic manipulation tasks. *Robotics and Computer-Integrated Manufacturing* 81: 102517.
- Englert P and Toussaint M (2017) Learning manipulation skills from a single demonstration. *The International Journal of Robotics Research* 37: 027836491774379.
- Englert P, Vien N and Toussaint M (2017) Inverse kkt: learning cost functions of manipulation tasks from demonstrations. *The International Journal of Robotics Research* 36: 027836491774598.
- Escarabajal RJ, Pulloquina JL, Zamora-Ortiz P, et al. (2023) Imitation learning-based system for the execution of self-paced robotic-assisted passive rehabilitation exercises. *IEEE Robotics and Automation Letters* 8(7): 4283–4290.
- Fan Y, Luo J and Tomizuka M (2019) A learning framework for high precision industrial assembly. In: 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20-24 May 2019. IEEE, 811–817.
- Fang B, Jia S, Guo D, et al. (2019) Survey of imitation learning for robotic manipulation. *International Journal of Intelligent Robotics and Applications* 3: 362–369.
- Fang K, Yin P, Nair A, et al. (2023) Generalization with lossy affordances: leveraging broad offline data for learning visuomotor tasks. *Proceedings of The 6th Conference on Robot Learning* 205: 106–117.
- Fang HS, Fang H, Tang Z, et al. (2024) Rh20t: a comprehensive robotic dataset for learning diverse skills in one-shot. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 13-17 May 2024. IEEE, 653–660.
- Farajtabar M and Charbonneau M (2024) The path towards contact-based physical human–robot interaction. *Robotics and Autonomous Systems* 182: 104829.
- Feng Y, Hansen N, Xiong Z, et al. (2023) Finetuning offline world models in the real world. *Proceedings of The 7th Conference on Robot Learning* 229: 425–445.
- Feng Y, Han J, Yang Z, et al. (2025) Reflective planning: Vision-language models for multi-stage long-horizon robotic manipulation. In: 9th Annual Conference on Robot Learning, Seoul, Republic of Korea.
- Ferguson S, Liu S, Mandikal V, et al. (2020) Leveraging demonstrations for reinforcement learning with hybrid representations. In: Conference on Robot Learning (CoRL), Virtual Conference, November 16 - 18, 2020.
- Finn C, Yu T, Zhang T, et al. (2017) One-shot visual imitation learning via meta-learning. In: Conference on Robot Learning, Mountain View, CA, USA, 357–368.
- Firoozi R, Tucker J, Tian S, et al. (2023) Foundation models in robotics: applications, challenges, and the future. *The International Journal of Robotics Research* 2024;44(5):701-739. DOI:10.1177/02783649241281508
- Florence P, Lynch C, Zeng A, et al. (2022) Implicit behavioral cloning. *Proceedings of the 5th Conference on Robot Learning* 164: 158–168.
- Franklin DW and Wolpert DM (2011) Computational mechanisms of sensorimotor control. *Neuron* 72(3): 425–442.
- Freeman CD, Frey E, Raichuk A, et al. (2021a) Brax—a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281*.
- Freeman CD, Frey E, Raichuk A, et al. (2021b) Brax—a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281*.
- Fu L, Datta G, Huang H, et al. (2024a) A touch, vision, and language dataset for multimodal alignment. In: Forty-First International Conference on Machine Learning, Vienna, Austria.
- Fu Z, Zhao TZ and Finn C (2024b) Mobile aloha: learning bimanual mobile manipulation with low-cost whole-body teleoperation. In: Conference on Robot Learning, Munich, Germany.
- Funabashi S, Ogasa S, Isobe T, et al. (2020) Variable in-hand manipulations for tactile-driven robot hand via cnn-lstm. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 24 October 2020 - 24 January 2021. IEEE, 9472–9479.
- Furuta D, Kutsuzawa K, Sakaino S, et al. (2020) Motion planning with success judgement model based on learning from demonstration. *IEEE Access* 8: 73142–73150.
- Gams A, Petrič T, Nemeč B, et al. (2022) Manipulation learning on humanoid robots. *Current Robotics Reports* 3(3): 97–109.
- George A, Gano S, Katragadda P, et al. (2025) Vital pretraining: Visuo-tactile pretraining for tactile and non-tactile

- manipulation policies. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 258–264. DOI: [10.1109/ICRA55743.2025.11128336](https://doi.org/10.1109/ICRA55743.2025.11128336).
- Ghosh D, Walke HR, Pertsch K, et al. (2024) Octo: an open-source generalist robot policy. In: Proceedings of Robotics: Science and Systems, Delft, Netherlands.
- Gonzalez-Aguirre JA, Osorio-Oliveros R, Rodríguez-Hernández KL, et al. (2021) Service robots: trends and technology. *Applied Sciences* 11(22): 10702.
- Goodfellow IJ, Pouget-Abadie J, Mirza M, et al. (2014) Generative adversarial nets. *Advances in Neural Information Processing Systems*. Montreal, Canada: Curran Associates, Inc, 27, 2672–2680.
- Grauman K, Westbury A, Byrne E, et al. (2022) Ego4d: around the world in 3,000 hours of egocentric video. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 18973–18990.
- Gu S, Yang L, Du Y, et al. (2024) A review of safe reinforcement learning: methods, theories and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Gubbi S, Kolathaya S and Amrutur B (2020) Imitation learning for high precision peg-in-hole tasks. In: 2020 6th International Conference on Control, Automation and Robotics (ICCAR), Singapore, 20-23 April 2020. IEEE, 368–372.
- Guo Y, Tang B, Akinola I, et al. (2025) SRSA: skill retrieval and adaptation for robotic assembly tasks. In: The Thirteenth International Conference on Learning Representations, Singapore.
- Halicka K and Surel D (2022) Smart living technologies in the context of improving the quality of life for older people: the case of the humanoid rudy robot. *Human Technology* 18(2): 191–208.
- Hao P, Zhang C, Li D, et al. (2025) Tla: Tactile-language-action model for contact-rich manipulation. *arXiv preprint arXiv: 2503.08548*.
- Hejrati M and Mattila J (2023) Nonlinear subsystem-based adaptive impedance control of physical human-robot-environment interaction in contact-rich tasks. *IEEE Robotics and Automation Letters* 8(10): 6083–6090.
- Higuera C, Sharma A, Bodduluri CK, et al. (2024) Sparsh: Self-supervised touch representations for vision-based tactile sensing. In: 8th Annual Conference on Robot Learning, Munich, Germany.
- Ho J and Ermon S (2016) Generative adversarial imitation learning. In: Advances in Neural Information Processing Systems, Volume 29, Barcelona Spain.
- Hoque R, Balakrishna A, Putterman C, et al. (2021) Lazydagger: reducing context switching in interactive imitation learning. In: 2021 IEEE 17th International Conference on Automation Science and Engineering (CASE), Lyon, France, 502–509.
- Hoque R, Balakrishna A, Novoseller E, et al. (2022) Thriftydagger: Budget-aware novelty and risk gating for interactive imitation learning. *Proceedings of the 5th Conference on Robot Learning* 164: 598–608.
- Hou Y, Liu Z, Chi C, et al. (2025) Adaptive compliance policy: learning approximate compliance for diffusion guided control. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 4829–4836. DOI: [10.1109/ICRA55743.2025.11128452](https://doi.org/10.1109/ICRA55743.2025.11128452).
- Hu EJ, Shen Y, Wallis P, et al. (2022) LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations.
- Hua J, Zeng L, Li G, et al. (2021) Learning for a robot: deep reinforcement learning, imitation learning, transfer learning. *Sensors* 21(4): 1278.
- Huang Y, Rozo L, Silvério J, et al. (2019) Kernelized movement primitives. *The International Journal of Robotics Research* 38(7): 833–852.
- Huang B, Wang Y, Yang X, et al. (2024) 3d-vitac: learning fine-grained manipulation with visuo-tactile sensing. In: 8th Annual Conference on Robot Learning, Munich, Germany.
- Ichiwara H, Ito H, Yamamoto K, et al. (2022) Contact-rich manipulation of a flexible object based on deep predictive learning using vision and tactility. In: 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23-27 May 2022. IEEE, 5375–5381.
- Ichiwara H, Ito H, Yamamoto K, et al. (2023) Modality attention for prediction-based robot motion generation: improving interpretability and robustness of using multi-modality. *IEEE Robotics and Automation Letters* 8(12): 8271–8278.
- Iodice F, Wu Y, Kim W, et al. (2022) Learning cooperative dynamic manipulation skills from human demonstration videos. *Mechatronics* 85: 102807.
- Ito H, Yamamoto K, Mori H, et al. (2022) Efficient multitask learning with an embodied predictive model for door opening and entry with whole-body control. *Science Robotics* 7(65): eaax8177.
- Jadeja Y, Shafik M, Wood P, et al. (2025) Enhancing healthcare assistance with a self-learning robotics system: a deep imitation learning-based solution. *Electronics* 14(14): 2823.
- James S, Davison AJ and Johns E (2019) Sim-to-real via sim-to-sim: Data-efficient robotic grasping via randomized-to-canonical adaptation networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 15-20 June 2019, 12627–12637.
- James S, Ma Z, Arrojo DR, et al. (2020) Rlbench: the robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters* 5(2): 3019–3026.
- Jang E, Irpan A, Khansari M, et al. (2022) Be-z: Zero-shot task generalization with robotic imitation learning. *Proceedings of the 5th Conference on Robot Learning* 164: 991–1002.
- Jassim HS, Akhter Y, Aalwahaab DZ, et al. (2025) Recent advances in tactile sensing technologies for human-robot interaction: current trends and future perspectives. *Biosensors and Bioelectronics X* 26: 100669.
- Jia X, Wang Q, Donat A, et al. (2024) MaLL: improving imitation learning with selective state space models. In: 8th Annual Conference on Robot Learning, Munich, Germany.

- Jiang Y, Huang Z, Yang B, et al. (2022) A review of robotic assembly strategies for the full operation procedure: planning, execution and evaluation. *Robotics and Computer-Integrated Manufacturing* 78: 102366.
- Johns E (2021) Coarse-to-fine imitation learning: robot manipulation from a single demonstration. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021 - 05 June 2021. IEEE, 4613–4619.
- Kamijo T, Beltran-Hernandez CC and Hamaya M (2024) Learning variable compliance control from a few demonstrations for bimanual robot with haptic feedback teleoperation system. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, United Arab Emirates, 14–18 October 2024, 12663–12670. DOI: [10.1109/IROS58592.2024.10801731](https://doi.org/10.1109/IROS58592.2024.10801731).
- Kang JH, Joshi S, Huang R, et al. (2025) Robotic compliant object prying using diffusion policy guided by vision and force observations. *IEEE Robotics and Automation Letters* 10(6): 5505–5512.
- Kato Y, Tsuji T and Cikajlo I (2024) Feedback type may change the emg pattern and kinematics during robot supported upper limb reaching task. *IEEE Open Journal of Engineering in Medicine and Biology* 5: 173–179.
- Kawaharazuka K, Matsushima T, Gambardella A, et al. (2024) Real-world robot applications of foundation models: a review. *Advanced Robotics* 38(18): 1232–1254.
- Kelly M, Sidrane C, Driggs-Campbell K, et al. (2019) Hg-dagger: interactive imitation learning with human experts. In: 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019, 8077–8083. DOI: [10.1109/ICRA.2019.8793698](https://doi.org/10.1109/ICRA.2019.8793698).
- Khandelwal A, Kant N and Mukherjee R (2023) Nonprehensile manipulation of a stick using impulsive forces. *Nonlinear Dynamics* 111(1): 113–127.
- Khansari-Zadeh SM and Billard A (2011) Learning stable nonlinear dynamical systems with gaussian mixture models. *IEEE Transactions on Robotics* 27(5): 943–957.
- Khatib O (1987) A unified approach for motion and force control of robot manipulators: the operational space formulation. *IEEE Journal of Robotics and Automation* 3(1): 43–53.
- Khazatsky A, Pertsch K, Nair S, et al. (2024) DROID: a large-scale in-the-wild robot manipulation dataset. In: Proceedings of Robotics: Science and Systems. Delft, Netherlands. DOI: [10.15607/RSS.2024.XX.120](https://doi.org/10.15607/RSS.2024.XX.120).
- Kim GH, Seo S, Lee J, et al. (2021) Demodice: offline imitation learning with supplementary imperfect demonstrations. In: International Conference on Learning Representations.
- Kim JW, Zhao TZ, Schmidgall S, et al. (2024a) Surgical robot transformer (srt): imitation learning for surgical tasks. In: Conference on Robot Learning, Munich, Germany.
- Kim MJ, Pertsch K, Karamcheti S, et al. (2024b) OpenVLA: an open-source vision-language-action model. In: 8th Annual Conference on Robot Learning, Munich, Germany.
- Kittmann R, Fröhlich T, Schäfer J, et al. (2015) Let me introduce myself: i am care-o-bot 4, a gentleman robot. In: *Mensch und Computer 2015—proceedings*. De Gruyter Oldenbourg, 223–232.
- Kober J and Peters J (2010) Imitation and reinforcement learning. *IEEE Robotics and Automation Magazine* 17(2): 55–62.
- Kober J, Bagnell JA and Peters J (2013) Reinforcement learning in robotics: a survey. *The International Journal of Robotics Research* 32(11): 1238–1274.
- Kober J, Gienger M and Steil JJ (2015) Learning movement primitives for force interaction tasks. In: 2015 IEEE International Conference on Robotics and Automation (ICRA), Seattle, WA, USA, 26–30 May 2015. IEEE, 3192–3199.
- Korkmaz Y and Bryık E (2025) Mile: Model-based intervention learning. *arXiv preprint arXiv:2502.13519*.
- Kormushev P, Calinon S and Caldwell DG (2011) Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input. *Advanced Robotics* 25(5): 581–603.
- Krebs HI, Hogan N, Aisen ML, et al. (1998) Robot-aided neurorehabilitation. *IEEE Transactions on Rehabilitation Engineering* 6(1): 75–87.
- Kroemer O, Niekum S and Konidaris G (2021) A review of robot learning for manipulation: challenges, representations, and algorithms. *Journal of Machine Learning Research* 22(30): 1–82.
- Kumar A, Singh A, Ebert FD, et al. (2023) Pre-training for robots: offline RL enables learning new tasks in a handful of trials. In: Proceedings of Robotics: Science and Systems, Daegu, Republic of Korea.
- Kutsuzawa K, Sakaino S and Tsuji T (2018) Sequence-to-sequence model for trajectory planning of nonprehensile manipulation including contact model. *IEEE Robotics and Automation Letters* 3(4): 3606–3613.
- Lakshminpathy A and Pollard NS (2024) Contactmpc: towards online adaptive control for contact-rich dexterous manipulation. In: 2nd Workshop on Dexterous Manipulation: Design, Perception and Control (RSS), Delft, Netherlands, Jul 15 – Jul 19, 2024.
- Lambeta M, Wu T, Sengul A, et al. (2024) Digitizing touch with an artificial multimodal fingertip. *arXiv preprint arXiv:2411.02479*.
- Langhorne P, Bernhardt J and Kwakkel G (2011) Stroke rehabilitation. *The Lancet* 377(9778): 1693–1702.
- Laskey M, Lee J, Fox R, et al. (2017) Dart: noise injection for robust imitation learning. In: Levine S, Vanhoucke V and Goldberg K (eds) *Proceedings of the 1st Annual Conference on Robot Learning, Proceedings of Machine Learning Research*. PMLR, Vol. 78, 143–156.
- Lawrence DA (2002) Stability and transparency in bilateral teleoperation. *IEEE Transactions on Robotics and Automation* 9(5): 624–637.
- Lázaro-Gredilla M, Lin D, Guntupalli JS, et al. (2019) Beyond imitation: Zero-shot task transfer on robots by learning concepts as cognitive programs. *Science Robotics* 4(26): eaav3150.

- Lee K, Su Y, Kim TK, et al. (2013) A syntactic approach to robot imitation learning using probabilistic activity grammars. *Robotics and Autonomous Systems* 61(12): 1323–1334.
- Lee Y, Lim JJ, Anandkumar A, et al. (2022) Adversarial skill chaining for long-horizon robot manipulation via terminal state regularization. In: Proceedings of the 5th Conference on Robot Learning, Auckland, New Zealand, 406–416.
- Lee SW, Kang X and Kuo YL (2025) Diff-dagger: uncertainty estimation with diffusion policy for robotic manipulation. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 4845–4852. DOI: [10.1109/ICRA55743.2025.11127730](https://doi.org/10.1109/ICRA55743.2025.11127730).
- Levine S, Finn C, Darrell T, et al. (2016) End-to-end training of deep visuomotor policies. *Journal of Machine Learning Research* 17(39): 1–40.
- Levine S, Kumar A, Tucker G, et al. (2020) Offline reinforcement learning: tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*.
- Li R and Zou Z (2023) Enhancing construction robot learning for collaborative and long-horizon tasks using generative adversarial imitation learning. *Advanced Engineering Informatics* 58: 102140.
- Li Y, Zhu JY, Tedrake R, et al. (2019) Connecting touch and vision via cross-modal prediction. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15-20 June 2019, 10601–10610.
- Li J, Lu T, Cao X, et al. (2021) Meta-imitation learning by watching video demonstrations. In: International Conference on Learning Representations, Singapore.
- Li C, Xia F, Martín-Martín R, et al. (2022a) Iqibson 2.0: Object-centric simulation for robot learning of everyday household tasks. In: Conference on Robot Learning, Auckland, New Zealand. PMLR, 455–465.
- Li Y, Sena A, Wang Z, et al. (2022b) A review on interaction control for contact robots through intent detection. *Progress in Biomedical Engineering* 4(3): 032004.
- Li G, Jin Z, Volpp M, et al. (2023) Prodmp: a unified perspective on dynamic and probabilistic movement primitives. *IEEE Robotics and Automation Letters* 8(4): 2325–2332.
- Li M, Qi X, Han X, et al. (2024a) Transfer force perception skills to robot-assisted laminectomy via imitation learning from human demonstrations. *CAAI Transactions on Intelligence Technology* 9(4): 903–916.
- Li X, Liu M, Zhang H, et al. (2024b) Vision-language foundation models as effective robot imitators. In: International Conference on Learning Representations, Vienna, Austria.
- Li H, Cui Y and Sadigh D (2025a) How to train your robots? The impact of demonstration modality on imitation learning. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 1113–1120. DOI: [10.1109/ICRA55743.2025.11128520](https://doi.org/10.1109/ICRA55743.2025.11128520).
- Li Y, Deng Y, Zhang J, et al. (2025b) HAMSTER: hierarchical action models for open-world robot manipulation. In: The Thirteenth International Conference on Learning Representations, Singapore.
- Liconti D, Toshimitsu Y and Katzschmann R (2024) Leveraging pretrained latent representations for few-shot imitation learning on an anthropomorphic robotic hand. In: 2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids), Nancy, France, 22-24 November 2024, 181–188.
- Lim JH, He K, Yi Z, et al. (2023) Adaptive learning based upper-limb rehabilitation training system with collaborative robot. In: 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), Sydney, Australia, 24-27 July 2023. IEEE, 1–5.
- Lin T, Zhang Y, Li Q, et al. (2025) Learning visuotactile skills with two multifingered hands. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025. IEEE, 5637–5643.
- Liu N, Li S, Du Y, et al. (2022a) Compositional visual generation with composable diffusion models. In: European Conference on Computer Vision, Israel. Springer, 423–439.
- Liu Z, Liu Q, Xu W, et al. (2022b) Robot learning towards smart robotic manufacturing: a review. *Robotics and Computer-Integrated Manufacturing* 77: 102360.
- Liu B, Zhu Y, Gao C, et al. (2023a) Libero: benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems* 36: 44776–44791.
- Liu J, Zu L, He L, et al. (2023b) Clue: calibrated latent guidance for offline reinforcement learning. *Proceedings of The 7th Conference on Robot Learning* 229: 906–927.
- Liu J, Li C, Wang G, et al. (2024) Self-corrected multimodal large language model for end-to-end robot manipulation. *arXiv preprint arXiv:2405.17418*.
- Liu JJ, Li Y, Shaw K, et al. (2025a) FACTR: Force-Attending curriculum training for contact-rich Policy learning. In: Proceedings of Robotics: Science and Systems, Los Angeles, CA, USA. DOI: [10.15607/RSS.2025.XXI.079](https://doi.org/10.15607/RSS.2025.XXI.079).
- Liu W, Wang J, Wang Y, et al. (2025b) Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 1105–1112. DOI: [10.1109/ICRA55743.2025.11128061](https://doi.org/10.1109/ICRA55743.2025.11128061).
- Lödige PW, Li MX and Lioutikov R (2025) Use the force, bot! - force-aware prodmp with event-based replanning. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 16730–16736. DOI: [10.1109/ICRA55743.2025.11128718](https://doi.org/10.1109/ICRA55743.2025.11128718).
- Luo J, Sushkov O, Pevceviciute R, et al. (2021) Robust multimodal policies for industrial assembly via reinforcement learning and demonstrations: a large-scale study. In: Proceedings of Robotics: Science and Systems, Virtual, July 12–16, 2021. Virtual. DOI: [10.15607/RSS.2021.XVII.088](https://doi.org/10.15607/RSS.2021.XVII.088).
- Luo J, Dong P, Wu J, et al. (2023) Action-quantized offline reinforcement learning for robotic skill learning. *Proceedings of The 7th Conference on Robot Learning* 229: 1348–1361.

- Makoviychuk V, Wawrzyniak L, Guo Y, et al. (2021) Isaac gym: high performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*.
- Mandi Z, Bharadhwaj H, Moens V, et al. (2022a) Cacti: a framework for scalable multi-task multi-scene visual imitation learning. In: CoRL 2022 Workshop on Pre-training Robot Learning, Auckland, New Zealand, Dec 14 2022.
- Mandi Z, Liu F, Lee K, et al. (2022b) Towards more generalizable one-shot visual imitation learning. In: 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23-27 May 2022. IEEE, 2434–2444.
- Mao X, Giudici G, Coppola C, et al. (2024) Dexskills: skill segmentation using haptic data for learning autonomous long-horizon robotic manipulation tasks. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Abu Dhabi, United Arab Emirates, 14-18 October 2024. IEEE, 5104–5111.
- Martín-Martín R, Lee MA, Gardner R, et al. (2019) Variable impedance control in end-effector space: an action space for reinforcement learning in contact-rich tasks. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 03-08 November 2019. IEEE, 1010–1017.
- Mason MT (1986) Mechanics and planning of manipulator pushing operations. *The International Journal of Robotics Research* 5(3): 53–71.
- Mathew J and Crevecoeur F (2021) Adaptive feedback control in human reaching adaptation to force fields. *Frontiers in Human Neuroscience* 15: 15–2021.
- Mees O, Hermann L and Burgard W (2022a) What matters in language conditioned robotic imitation learning over unstructured data. *IEEE Robotics and Automation Letters* 7(4): 11205–11212.
- Mees O, Hermann L, Rosete-Beas E, et al. (2022b) Calvin: a benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters* 7(3): 7327–7334.
- Memmel M, Berg J, Chen B, et al. (2025) STRAP: robot sub-trajectory retrieval for augmented policy learning. In: The Thirteenth International Conference on Learning Representations, Singapore.
- Merckaert K, Convens B, ju Wu C, et al. (2022) Real-time motion control of robotic manipulators for safe human–robot coexistence. *Robotics and Computer-Integrated Manufacturing* 73: 102223.
- Meribout M, Takele NA, Derege O, et al. (2024) Tactile sensors: a review. *Measurement* 238: 115332.
- Moghani M, Nelson N, Ghanem M, et al. (2025) Sufia-bc: generating high quality demonstration data for visuomotor policy learning in surgical subtasks. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 4534–4541. DOI: [10.1109/ICRA55743.2025.11127797](https://doi.org/10.1109/ICRA55743.2025.11127797).
- Mora A, Prados A, Mendez A, et al. (2024) Adam: a robotic companion for enhanced quality of life in aging populations. *Frontiers in Neurorobotics* 18: 1337608.
- Nair A, Chen D, Agrawal P, et al. (2017) Combining self-supervised learning and imitation for vision-based rope manipulation. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), Singapore, 29 May 2017 - 03 June 2017. IEEE, 2146–2153.
- Nemec B, Abu-Dakka FJ, Ridge B, et al. (2013) Transfer of assembly operations to new workpiece poses by adaptation to the desired force profile. In: 2013 16th International Conference on Advanced Robotics (ICAR), Montevideo, Uruguay, 25-29 November 2013. IEEE, 1–7.
- Neville H (1985) Impedance control: an approach to manipulation: part i-iii. *Trans. of the ASME Journal of Dynamic System, Measurement, and Control* 107: 1.
- Ng AY and Russell S (2000) Algorithms for inverse reinforcement learning. In: *Icml*, Vol. 1, 2.
- Nie W, Vahdat A and Anandkumar A (2021) Controllable and compositional generation with latent-space energy-based models. *Advances in Neural Information Processing Systems* 34: 13497–13510.
- Nix J, Smith A, Kurpad R, et al. (2010) Prospective randomized controlled trial of robotic versus open radical cystectomy for bladder cancer: perioperative and pathologic results. *European Urology* 57(2): 196–201.
- Noseworthy M, Paul R, Roy S, et al. (2020) Task-conditioned variational autoencoders for learning movement primitives. In: Conference on Robot Learning. PMLR, 933–944.
- Oh H and Matusbara T (2024) Leveraging demonstrator-perceived precision for safe interactive imitation learning of clearance-limited tasks. *IEEE Robotics and Automation Letters* 9(4): 3387–3394.
- Oh H, Sasaki H, Michael B, et al. (2021) Bayesian disturbance injection: robust imitation learning of flexible policies. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021 - 05 June 2021, 8629–8635. DOI: [10.1109/ICRA48506.2021.9561573](https://doi.org/10.1109/ICRA48506.2021.9561573).
- Oikawa M, Kusakabe T, Kutsuzawa K, et al. (2021) Reinforcement learning for robotic assembly using non-diagonal stiffness matrix. *IEEE Robotics and Automation Letters* 6(2): 2737–2744.
- Okada M, Komatsu M, Okumura R, et al. (2023) Learning compliant stiffness by impedance control-aware task segmentation and multi-objective bayesian optimization with priors. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Detroit, MI, USA, 01-05 October 2023. IEEE, 8155–8162.
- Onstein IF, Semeniuta O and Bjerkeng M (2020) Deburring using robot manipulators: a review. In: 2020 3rd International Symposium on Small-scale Intelligent Manufacturing Systems (SIMS), Gjøvik, Norway, 10-12 June 2020. IEEE, 1–7.
- Osa T, Staub C and Knoll A (2010) Framework of automatic robot surgery system using visual servoing. In: 2010 IEEE/RSJ

- International Conference on Intelligent Robots and Systems, Taipei, Taiwan, 18-22 October 2010, 1837–1842.
- Osa T, Sugita N and Mitsuishi M (2018) Online trajectory planning and force control for automation of surgical tasks. *IEEE Transactions on Automation Science and Engineering* 15(2): 675–691.
- Ozdamar I, Sirintuna D, Arbaud R, et al. (2024) Pushing in the dark: a reactive pushing strategy for mobile robots using tactile feedback. *IEEE Robotics and Automation Letters*.
- Özdemir O, Kerzel M, Weber C, et al. (2023) Language-model-based paired variational autoencoders for robotic language learning. *IEEE Transactions on Cognitive and Developmental Systems* 15(4): 1812–1824.
- O’Neill A, Rehman A, Maddukuri A, et al. (2024) Open x-embodiment: robotic learning datasets and rt-x models: open x-embodiment collaboration. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 13-17 May 2024, 6892–6903. DOI: [10.1109/ICRA57147.2024.10611477](https://doi.org/10.1109/ICRA57147.2024.10611477).
- Paraschos A, Daniel C, Peters JR, et al. (2013) Probabilistic movement primitives. *Advances in Neural Information Processing Systems* 26: Lake Tahoe, CA, USA.
- Peng XB, Andrychowicz M, Zaremba W, et al. (2018) Sim-to-real transfer of robotic control with dynamics randomization. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, QLD, Australia, 21-25 May 2018. IEEE, 3803–3810.
- Pernel L, Petrič T, Oztop E, et al. (2014) Teaching robots to cooperate with humans in dynamic manipulation tasks based on multi-modal human-in-the-loop approach. *Autonomous Robots* 36(1-2): 123–136.
- Pernel L, Petrič T and Babič J (2015) Human-in-the-loop approach for teaching robot assembly tasks using impedance control interface. In: 2015 IEEE International Conference on Robotics and Automation (ICRA), Seattle, WA, USA, 26-30 May 2015. IEEE, 1497–1502.
- Pernel L, Noda T, Petrič T, et al. (2016) Adaptive control of exoskeleton robots for periodic assistive behaviours based on emg feedback minimisation. *PLoS One* 11(2): 1–26.
- Pernel L, Rozo L, Caldwell D, et al. (2017) A method for derivation of robot task-frame control authority from repeated sensory observations. *IEEE Robotics and Automation Letters* 2(2): 719–726.
- Petrič T, Simpson CS, Ude A, et al. (2017) Hammering does not fit fitts’ law. *Frontiers in Computational Neuroscience* 11(May): 1–12.
- Petric T, Gams A, Colasanto L, et al. (2018) Accelerated sensorimotor learning of compliant movement primitives. *IEEE Transactions on Robotics* 34(6): 1636–1642.
- Pj M and Bdo A (1971) Nonlinear regulator theory and an inverse optimal control problem. *Paper Series* 6-C3: 462–467.
- Pong V, Gu S, Dalal M, et al. (2018) Temporal difference models: model-free deep rl for model-based control. In: International Conference on Learning Representations, Vancouver, Canada.
- Prasad A, Lin K, Wu J, et al. (2024) Consistency policy: accelerated visuomotor policies via consistency distillation. In: Proceedings of Robotics: Science and Systems, Delft, Netherlands, Jul 15 – Jul 19, 2024. DOI: [10.15607/RSS.2024.XX.071](https://doi.org/10.15607/RSS.2024.XX.071).
- Qian K, Yue Z and Bai J (2025) Hierarchical kernelized movement primitives for learning human-robot collaborative trajectories in referred object handover. *Applied Intelligence* 55(1): 1–15.
- Rafailov R, Hatch KB, Kolev V, et al. (2023) Moto: offline pre-training to online fine-tuning for model-based robot learning. In: Proceedings of the 7th Conference on Robot Learning, Atlanta, GA, USA, 3654–3671.
- Rahmatizadeh R, Abolghasemi P, Bölöni L, et al. (2018) Vision-based multi-task manipulation for inexpensive robots using end-to-end learning from demonstration. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, QLD, Australia, 21-25 May 2018. IEEE, 3758–3765.
- Raibert MH and Craig JJ (1981) Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control* 103(2): 126–133.
- Rajeswaran A, Kumar V, Gupta A, et al. (2018) Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. *Robotics: Science and Systems XIV*.
- 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.
- Ren A, Veer S and Majumdar A (2021) Generalization guarantees for imitation learning. In: Conference on Robot Learning, London, United Kingdom. PMLR, 1426–1442.
- Rethmeier N and Augenstein I (2023) A primer on contrastive pretraining in language processing: methods, lessons learned, and perspectives. *ACM Computing Surveys* 55(10): 1–17.
- Ross S and Bagnell D (2010) Efficient reductions for imitation learning. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings, Sardinia, Italy, 661–668.
- Ross S, Gordon G and Bagnell D (2011) A reduction of imitation learning and structured prediction to no-regret online learning. In: Gordon G, Dunson D and Dudik M (eds) *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, Proceedings of Machine Learning Research*. PMLR, Vol. 15, 627–635.
- Royo-Miquel J, Hamaya M, Beltran-Hernandez CC, et al. (2023) Learning robotic assembly by leveraging physical softness and tactile sensing. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Detroit, MI, USA, 01-05 October 2023, 6469–6476.
- Rundong W, Yanchun C, Qilong Y, et al. (2025) Robomt: Human-like compliance control for assembly via a bilateral robotic teleoperation and hybrid mamba-transformer framework. *IEEE Robotics and Automation Letters* 10(8): 7771–7778.
- Ruppel P and Zhang J (2024) Elastic tactile sensor glove for dexterous teaching by demonstration. *Sensors* 24(6): 1912.

- Salhotra G, Liu ICA, Dominguez-Kuhne M, et al. (2022) Learning deformable object manipulation from expert demonstrations. *IEEE Robotics and Automation Letters* 7(4): 8775–8782.
- Samadikhoshkho Z, Saive E and Lipsett MG (2025) A review of compliant mechanisms for contact robotics applications. *Robotics and Autonomous Systems* 186: 104902.
- Sasagawa A, Fujimoto K, Sakaino S, et al. (2020) Imitation learning based on bilateral control for human–robot cooperation. *IEEE Robotics and Automation Letters* 5(4): 6169–6176.
- Saveriano M, Abu-Dakka FJ, Kramberger A, et al. (2023) Dynamic movement primitives in robotics: a tutorial survey. *The International Journal of Robotics Research* 42(13): 1133–1184.
- Schaal S, Peters J, Nakanishi J, et al. (2003) Control, planning, learning, and imitation with dynamic movement primitives. In: Workshop on Bilateral Paradigms on Humans and Humanoids: IEEE International Conference on Intelligent Robots and Systems (IROS 2003), Las Vegas, NV, USA, 27–31 October 2003, 1–21.
- Scherzinger S, Roennau A and Dillmann R (2019) Contact skill imitation learning for robot-independent assembly programming. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Macau, China, 03–08 November 2019. IEEE, 4309–4316.
- Seita D, Ganapathi A, Hoque R, et al. (2020) Deep imitation learning of sequential fabric smoothing from an algorithmic supervisor. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, NV, USA, 24 October 2020 - 24 January 2021. IEEE, 9651–9658.
- Seker MY, Imre M, Piater JH, et al. (2019) Conditional neural movement primitives. In: *Proceedings of Robotics: Science and Systems (RSS)*, Freiburg im Breisgau, Germany: Robotics: Science and Systems Foundation. DOI: [10.15607/RSS.2019.XV.071](https://doi.org/10.15607/RSS.2019.XV.071)
- Sermanet P, Xu K and Levine S (2017) Unsupervised perceptual rewards for imitation learning. In: *Proceedings of Robotics: Science and Systems*, Cambridge, MA, USA.
- Sharma P, Pathak D and Gupta A (2019) Third-person visual imitation learning via decoupled hierarchical controller. In H. Wallach, H. Larochelle, A. Beygelzimer, et al. (eds.). *Advances in Neural Information Processing Systems (NeurIPS)*. Curran Associates, Inc: Vancouver, Canada. 32. https://proceedings.neurips.cc/paper_files/paper/2019/file/8a146f1a3da4700cbf03cdc55e2daae6-Paper.pdf
- Shaw K, Agarwal A, Bahl S, et al. (2024) Demonstrating learning from humans on open-source dexterous robot hands. In: *Robotics: Science and Systems*, Delft, Netherlands.
- Sherstinsky A (2020) Fundamentals of recurrent neural network (rnn) and long short-term memory (lstm) network. *Physica D: Nonlinear Phenomena* 404: 132306.
- Sherwani F, Asad MM and Ibrahim BSKK (2020) Collaborative robots and industrial revolution 4.0 (ir 4.0). In: 2020 International Conference on Emerging Trends in Smart Technologies (ICETST), Karachi, Pakistan, 26–27 March 2020. IEEE, 1–5.
- Shi Y, Chen Z, Liu H, et al. (2021) Proactive action visual residual reinforcement learning for contact-rich tasks using a torque-controlled robot. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021 - 05 June 2021. IEEE, 765–771.
- Shi LX, Hu Z, Zhao TZ, et al. (2024) Yell at your robot: improving on-the-fly from language corrections. In: *Proceedings of Robotics: Science and Systems*, Delft, Netherlands.
- Shih B, Shah D, Li J, et al. (2020) Electronic skins and machine learning for intelligent soft robots. *Science Robotics* 5(41): eaaz9239.
- Shridhar M, Manuelli L and Fox D (2022) Cliport: what and where pathways for robotic manipulation. In: *Conference on Robot Learning*, Auckland, New Zealand. PMLR, 894–906.
- Shridhar M, Manuelli L and Fox D (2023) Perceiver-actor: a multi-task transformer for robotic manipulation *Conference on Robot Learning*. PMLR, 785–799.
- Siciliano B and Villani L (1999) *Robot Force Control*. Springer Science & Business Media.
- Sidiropoulos A and Dougeri Z (2021) A reversible dynamic movement primitive formulation. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021 - 05 June 2021. IEEE, 3147–3153.
- Solak G and Jamone L (2019) Learning by demonstration and robust control of dexterous in-hand robotic manipulation skills. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (Iros), Macau, China, 03–08 November 2019. IEEE, 8246–8251.
- Stadie BC, Abbeel P and Sutskever I (2017) Third-person imitation learning. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France.
- Stepputtis S, Bandari M, Schaal S, et al. (2022) A system for imitation learning of contact-rich bimanual manipulation policies. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Kyoto, Japan, 23–27 October 2022. IEEE, 11810–11817.
- Stone A, Xiao T, Lu Y, et al. (2023) Open-world object manipulation using pre-trained vision-language models. In: *Conference on Robot Learning*, Atlanta, GA, USA. PMLR, 3397–3417.
- Sugawara K, Sakaino S and Tsuji T (2023) Unsupervised human motion segmentation based on characteristic force signals of contact events. *IEEE Robotics and Automation Letters* 8(10): 6203–6210.
- Sundaram S, Kellnhofer P, Li Y, et al. (2019) Learning the signatures of the human grasp using a scalable tactile glove. *Nature* 569(7758): 698–702.
- Suomalainen M, Karayiannidis Y and Kyrki V (2022) A survey of robot manipulation in contact. *Robotics and Autonomous Systems* 156: 104224.
- Tanwani AK, Yan A, Lee J, et al. (2021) Sequential robot imitation learning from observations. *The International Journal of Robotics Research* 40(10–11): 1306–1325.
- Tedrake R and the Drake Development Team (2019) Drake: a planning, control, and analysis toolbox for nonlinear dynamical systems. <https://drake.mit.edu>
- Tobin J, Fong R, Ray A, et al. (2017) Domain randomization for transferring deep neural networks from simulation to the

- real world. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vancouver, BC, Canada, 24-28 September 2017. IEEE, 23-30.
- Todorov E, Erez T and Tassa Y (2012) Mujoco: a physics engine for model-based control. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura, Portugal. IEEE, 5026-5033.
- Tsujii T (2025) Mamba as a motion encoder for robotic imitation learning. *IEEE Access* 13: 69941-69949. DOI: [10.1109/ACCESS.2025.3561283](https://doi.org/10.1109/ACCESS.2025.3561283)
- Tsujii C, Coronado E, Osorio P, et al. (2025) Adaptive contact-rich manipulation through few-shot imitation learning with force-torque feedback and pre-trained object representations. *IEEE Robotics and Automation Letters* 10(1): 240-247.
- Tsurumine Y, Cui Y, Yamazaki K, et al. (2019) Generative adversarial imitation learning with deep p-network for robotic cloth manipulation. In: 2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids), Toronto, ON, Canada, 15-17 October 2019. IEEE, 274-280.
- Ude A, Nemec B, Petrić T, et al. (2014) Orientation in cartesian space dynamic movement primitives. In: 2014 IEEE International Conference on Robotics and Automation (ICRA), Hong Kong, China, 31 May 2014 - 07 June 2014. IEEE, 2997-3004.
- Ugur E and Girgin H (2020) Compliant parametric dynamic movement primitives. *Robotica* 38(3): 457-474.
- Urain J, Mandlkar A, Du Y, et al. (2024) Deep generative models in robotics: a survey on learning from multimodal demonstrations. *arXiv preprint arXiv:2408.04380*.
- Ureche ALP, Umezawa K, Nakamura Y, et al. (2015) Task parameterization using continuous constraints extracted from human demonstrations. *IEEE Transactions on Robotics* 31(6): 1458-1471.
- Van Duong L (2025) A tactile reflex arc for physical human-robot interaction. *Mechatronics* 107: 103307.
- Van Hoof H, Chen N, Karl M, et al. (2016) Stable reinforcement learning with autoencoders for tactile and visual data. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Daejeon, Korea (South), 09-14 October 2016. IEEE, 3928-3934.
- Vaswani A, Shazeer N, Parmar N, et al. (2017) *Attention is all you need*. NIPS.
- Vecerik M, Hester T, Scholz J, et al. (2017) Leveraging demonstrations for deep reinforcement learning on robotics problems with sparse rewards. *arXiv preprint arXiv:1707.08817*.
- Vogt D, Stepputtis S, Grehl S, et al. (2017) A system for learning continuous human-robot interactions from human-human demonstrations. In: 2017 IEEE International Conference on Robotics and Automation (ICRA), Singapore, 29 May 2017 - 03 June 2017. IEEE, 2882-2889.
- Walke HR, Black K, Zhao TZ, et al. (2023) Bridgedata v2: a dataset for robot learning at scale. *Proceedings of The 7th Conference on Robot Learning* 229: 1723-1736.
- Wan W, Zhu Y, Shah R, et al. (2024) Lotus: continual imitation learning for robot manipulation through unsupervised skill discovery. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan. IEEE, 537-544.
- Wang Z, Merel JS, Reed SE, et al. (2017) Robust imitation of diverse behaviors. *Advances in Neural Information Processing Systems*. Long Beach, CA, USA: Curran Associates, Inc, 30, 5320-5329.
- Wang C, Lin C, Liu B, et al. (2021) Deep reinforcement learning with shaping exploration space for robotic assembly. In: 2021 3rd International Symposium on Robotics & Intelligent Manufacturing Technology (ISRIMT), Changzhou, China, 24-26 September 2021. IEEE, 345-351.
- Wang S, Lambeta M, Chou PW, et al. (2022a) Tacto: a fast, flexible, and open-source simulator for high-resolution vision-based tactile sensors. *IEEE Robotics and Automation Letters* 7(2): 3930-3937.
- Wang Y, Beltran-Hernandez CC, Wan W, et al. (2022b) An adaptive imitation learning framework for robotic complex contact-rich insertion tasks. *Frontiers in Robotics and AI* 8: 777363.
- Wang C, Fan L, Sun J, et al. (2023) Mimicplay: long-horizon imitation learning by watching human play. *Proceedings of The 7th Conference on Robot Learning* 229: 201-221.
- Wang C, Shi H, Wang W, et al. (2024a) Dexcap: scalable and portable mocap data collection system for dexterous manipulation. In: RSS 2024 Workshop: Data Generation for Robotics, Delft, Netherlands, Jul 15 - Jul 19, 2024.
- Wang C, Su C, Sun B, et al. (2024b) Extended residual learning with one-shot imitation learning for robotic assembly in semi-structured environment. *Frontiers in Neurorobotics* 18: 1355170.
- Wang L, Zhang K, Zhou A, et al. (2024c) Robot fleet learning via policy merging. In: Proceedings of the International Conference on Learning Representations (ICLR), Vienna Austria, May 7th - 11th, 2024.
- Wang L, Zhao J, Du Y, et al. (2024d) Poco: policy composition from and for heterogeneous robot learning. In: Robotics: Science and Systems, Delft, Netherlands.
- Welte E and Rayyes R (2025) Interactive imitation learning for dexterous robotic manipulation: challenges and perspectives—a survey. *arXiv preprint arXiv:2506.00098*.
- Wen J, Zhu Y, Li J, et al. (2025) DexVLA: Vision-language model with plug-in diffusion expert for general robot control. In: 9th Annual Conference on Robot Learning, Seoul, Republic of Korea.
- Whitney DE (1982) Quasi-static assembly of compliantly supported rigid parts. *Journal of Dynamic Systems, Measurement, and Control* 104(1): 65-77.
- Wolpert DM, Diedrichsen J and Flanagan JR (2011) Principles of sensorimotor learning. *Nature Reviews Neuroscience* 12(12): 739-751.
- Wu P, Escontrela A, Hafner D, et al. (2023) Daydreamer: world models for physical robot learning. In: Conference on Robot Learning, Atlanta, GA, USA. PMLR, 2226-2240.
- Wu H, Jing Y, Cheang C, et al. (2024) Unleashing large-scale video generative pre-training for visual robot manipulation. In: International Conference on Learning Representations, Vienna, Austria.

- Wu R, Chen H, Zhang M, et al. (2025) Neural dynamics augmented diffusion policy. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), Atlanta, GA, USA, 19-23 May 2025, 13234–13241. DOI: [10.1109/ICRA55743.2025.11128651](https://doi.org/10.1109/ICRA55743.2025.11128651).
- Xian Z, Gkanatsios N, Gervet T, et al. (2023) Chaineddiffuser: unifying trajectory diffusion and keypose prediction for robotic manipulation. In: Conference on Robot Learning, Atlanta, GA, USA. PMLR, 2323–2339.
- Xiang G, Li S, Shuang F, et al. (2024) Sc-airl: Share-critic in adversarial inverse reinforcement learning for long-horizon task. *IEEE Robotics and Automation Letters* 9(4): 3179–3186.
- Xiong H, Fu H, Zhang J, et al. (2023) Robotube: learning household manipulation from human videos with simulated twin environments. In: Conference on Robot Learning, Atlanta, GA, USA. PMLR, 1–10.
- Xu J, Hou Z, Liu Z, et al. (2019) Compare contact model-based control and contact model-free learning: a survey of robotic peg-in-hole assembly strategies. *arXiv preprint arXiv:1904.05240*.
- Xu X, You M, Zhou H, et al. (2022) Robot imitation learning from image-only observation without real-world interaction. *IEEE/ASME Transactions on Mechatronics* 28(3): 1234–1244.
- Xu B, Ud Din M and Hussain I (2025) Conditional variational auto encoder based dynamic motion for multitask imitation learning. *Scientific Reports* 15(1): 9196.
- Xue H, Ren J, Chen W, et al. (2025) Reactive diffusion policy: slow-fast visual-tactile policy learning for contact-rich manipulation. In: Proceedings of Robotics: Science and Systems. Los Angeles, CA, USA. DOI: [10.15607/RSS.2025.XXI.052](https://doi.org/10.15607/RSS.2025.XXI.052).
- Yamashita Y and Tani J (2008) Emergence of functional hierarchy in a multiple timescale neural network model: a humanoid robot experiment. *PLoS Computational Biology* 4(11): e1000220.
- Yan G, Wu YH and Wang X (2024) Dnact: diffusion guided multi-task 3d policy learning. *arXiv preprint arXiv:2403.04115*.
- Yang PC, Sasaki K, Suzuki K, et al. (2016) Repeatable folding task by humanoid robot worker using deep learning. *IEEE Robotics and Automation Letters* 2(2): 397–403.
- Yang C, Chen C, He W, et al. (2018) Robot learning system based on adaptive neural control and dynamic movement primitives. *IEEE Transactions on Neural Networks and Learning Systems* 30(3): 777–787.
- Yang G, Pang Z, Deen MJ, et al. (2020a) Homecare robotic systems for healthcare 4.0: visions and enabling technologies. *IEEE Journal of Biomedical and Health Informatics* 24(9): 2535–2549.
- Yang PC, Koganti N, Garcia Ricardez GA, et al. (2020b) Context dependent trajectory generation using sequence-to-sequence models for robotic toilet cleaning. In: 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), Naples, Italy, 31 August 2020 - 04 September 2020, 932–937. DOI: [10.1109/RO-MAN47096.2020.9223341](https://doi.org/10.1109/RO-MAN47096.2020.9223341).
- Yang J, Zhang J, Settle C, et al. (2022) Learning periodic tasks from human demonstrations. In: 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23-27 May 2022. IEEE, 8658–8665.
- Yang W, Angleraud A, Pieters RS, et al. (2023) Seq2seq imitation learning for tactile feedback-based manipulation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA), London, United Kingdom, 29 May 2023 - 02 June 2023, 5829–5836. DOI: [10.1109/ICRA48891.2023.10161145](https://doi.org/10.1109/ICRA48891.2023.10161145).
- Yang J, Mark MS, Vu B, et al. (2024) Robot fine-tuning made easy: pre-training rewards and policies for autonomous real-world reinforcement learning. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), Yokohama, Japan, 13-17 May 2024. IEEE, 4804–4811.
- Yang L, Suh H, Zhao T, et al. (2025) Physics-driven data generation for contact-rich manipulation via trajectory optimization. *arXiv preprint arXiv:2502.20382*.
- Yao Z, Wu W, Gao F, et al. (2026) Flexible tactile sensing systems: challenges in theoretical research transferring to practical applications. *Nano-Micro Letters* 18(1): 1–69.
- Yasa O, Toshimitsu Y, Michelis MY, et al. (2023) An overview of soft robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 6(1): 1–29.
- Yin H, Varava A and Kragic D (2021) Modeling, learning, perception, and control methods for deformable object manipulation. *Science Robotics* 6(54): eabd8803.
- Yoo Y, Lee CY and Zhang BT (2021) Multimodal anomaly detection based on deep auto-encoder for object slip perception of mobile manipulation robots. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), Xi'an, China, 30 May 2021 - 05 June 2021, 11443–11449.
- Yu T, Finn C, Xie A, et al. (2018) One-shot imitation from observing humans via domain-adaptive meta-learning *Robotics: Science and Systems (RSS)*.
- Yu T, Quillen D, He Z, et al. (2020) Meta-world: a benchmark and evaluation for multi-task and meta reinforcement learning. In: Conference on Robot Learning. PMLR, 1094–1100.
- Yu J, Liu H, Yu Q, et al. (2025) Forcevla: enhancing vla models with a force-aware moe for contact-rich manipulation. *arXiv preprint arXiv:2505.22159*.
- Zachiotis GA, Andrikopoulos G, Gornez R, et al. (2018) A survey on the application trends of home service robotics. In: 2018 IEEE International Conference on Robotics and Biomimetics (ROBIO), Kuala Lumpur, Malaysia, 12-15 December 2018. IEEE, 1999–2006.
- Zang Y, Wang P, Zha F, et al. (2024) Human skill knowledge guided global trajectory policy reinforcement learning method. *Frontiers in Neurobotics* 18: 1368243.
- Zare M, Kebria PM, Khosravi A, et al. (2024) A survey of imitation learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics* 54(12): 7173–7186.
- Zeeustraten MJ, Havoutis I, Silvério J, et al. (2017) An approach for imitation learning on riemannian manifolds. *IEEE Robotics and Automation Letters* 2(3): 1240–1247.
- Zeng A, Song S, Lee J, et al. (2020) Tossingbot: learning to throw arbitrary objects with residual physics. *IEEE Transactions on Robotics* 36(4): 1307–1319.

- Zeng X, Zhu G, Gao Z, et al. (2023) Surface polishing by industrial robots: a review. *The International Journal of Advanced Manufacturing Technology* 125(9): 3981–4012.
- Zhang F and Demiris Y (2023) Visual-tactile learning of garment unfolding for robot-assisted dressing. *IEEE Robotics and Automation Letters* 8(9): 5512–5519.
- Zhang T, McCarthy Z, Jow O, et al. (2018) Deep imitation learning for complex manipulation tasks from virtual reality teleoperation. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, QLD, Australia, 21-25 May 2018. Ieee, 5628–5635.
- Zhang H, Zhong X, Huang Z, et al. (2021a) Sensor-free method with bp network to achieve drag teaching on the 7-dof collaborative robot. In: 2021 China Automation Congress (CAC), Beijing, China, 22-24 October 2021, 7895–7899. DOI: [10.1109/CAC53003.2021.9727613](https://doi.org/10.1109/CAC53003.2021.9727613).
- Zhang X, Sun L, Kuang Z, et al. (2021b) Learning variable impedance control via inverse reinforcement learning for force-related tasks. *IEEE Robotics and Automation Letters* 6(2): 2225–2232.
- Zhang T, Sun H and Zou Y (2022) An electromyography signals-based human-robot collaboration system for human motion intention recognition and realization. *Robotics and Computer-Integrated Manufacturing* 77: 102359.
- Zhang Y, Ke L, Deshpande A, et al. (2023) Cherry-Picking with reinforcement learning. In: Proceedings of Robotics: Science and Systems, Daegu, Republic of Korea.
- Zhang H, Solak G, Lahr GJG, et al. (2024) Srl-vic: a variable stiffness-based safe reinforcement learning for contact-rich robotic tasks. *IEEE Robotics and Automation Letters* 9(6): 5631–5638.
- Zhang C, Hao P, Cao X, et al. (2025a) Vtla: vision-tactile-language-action model with preference learning for insertion manipulation. *arXiv preprint arXiv:2505.09577*.
- Zhang H, Huang WH, Solak G, et al. (2025b) Omnivic: a self-improving variable impedance controller with vision-language in-context learning for safe robotic manipulation. *arXiv preprint arXiv:2510.17150*.
- Zhang H, Solak G and Ajoudani A (2025c) Bresa: Bio-inspired reflexive safe reinforcement learning for contact-rich robotic tasks. *arXiv preprint arXiv:2503.21989*.
- Zhang H, Solak G, Hjorth S, et al. (2025d) Towards passive safe reinforcement learning: a comparative study on contact-rich robotic manipulation. *arXiv preprint arXiv:2503.00287*.
- Zhang H, Zhu C, Wang X, et al. (2025e) Badrobot: jailbreaking embodied LLMs in the physical world. In: The Thirteenth International Conference on Learning Representations, Singapore.
- Zhao J, Giammarino A, Lamon E, et al. (2022a) A hybrid learning and optimization framework to achieve physically interactive tasks with mobile manipulators. *IEEE Robotics and Automation Letters* 7(3): 8036–8043.
- Zhao TZ, Luo J, Sushkov O, et al. (2022b) Offline meta-reinforcement learning for industrial insertion. In: 2022 International Conference on Robotics and Automation (ICRA), Philadelphia, PA, USA, 23-27 May 2022, 6386–6393.
- Zhao TZ, Kumar V, Levine S, et al. (2023a) Learning fine-grained bimanual manipulation with low-cost hardware. In: Proceedings of Robotics: Science and Systems. Daegu, Republic of Korea, Daegu, Republic of Korea, Jul 10 – Jul 14, 2023. DOI: [10.15607/RSS.2023.XIX.016](https://doi.org/10.15607/RSS.2023.XIX.016).
- Zhao W, He T, Chen R, et al. (2023b) State-wise safe reinforcement learning: a survey. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, Macao, China. 6814–6822.
- Zhao W, Ding P, Min Z, et al. (2025a) VLAS: Vision-language-action model with speech instructions for customized robot manipulation *The Thirteenth International Conference on Learning Representations*.
- Zhao Z, Haldar S, Cui J, et al. (2025b) Touch begins where vision ends: generalizable policies for contact-rich manipulation. In: 3rd RSS Workshop on Dexterous Manipulation: Learning and Control with Diverse Data, Sydney, Singapore.
- Zhou G, Ke L, Srinivasa S, et al. (2023) Real world offline reinforcement learning with realistic data source. In: 2023 IEEE International Conference on Robotics and Automation (ICRA), London, United Kingdom, 29 May 2023 - 02 June 2023, 7176–7183.
- Zhou Z, Yang X and Zhang X (2025) Variable impedance control on contact-rich manipulation of a collaborative industrial mobile manipulator: an imitation learning approach. *Robotics and Computer-Integrated Manufacturing* 92: 102896.
- Zhu Z and Hu H (2018) Robot learning from demonstration in robotic assembly: a survey. *Robotics* 7(2): 17.
- Zhu J, Cherubini A, Dune C, et al. (2022) Challenges and outlook in robotic manipulation of deformable objects. *IEEE Robotics and Automation Magazine* 29(3): 67–77.
- Zitkovich B, Yu T, Xu S, et al. (2023) Rt-2: Vision-language-action models transfer web knowledge to robotic control. *Proceedings of The 7th Conference on Robot Learning* 229: 2165–2183.