

IMITATION LEARNING FOR HUMAN-ROBOT TEAMING IN BATTERY DISASSEMBLY: ENABLING TECHNOLOGIES AND A COLLABORATIVE MANIPULATION PIPELINE

Filippo SanFilippo¹ and Cecilia Scoccia²

¹Department of Engineering Sciences, University of Agder, 4879 Grimstad, Norway, e-mail: filippo.sanfilippo@uia.no

²Department of Industrial Engineering and Mathematical Sciences, Marche Polytechnic University, 60121 Ancona, Italy, e-mail: c.scoccia@staff.univpm.it

KEYWORDS

Human-Robot Teaming, Imitation Learning, Robot Learning, AI-Driven Robotics.

ABSTRACT

Robotics has progressed from isolated industrial manipulators to systems capable of smooth interaction, collaboration, and ultimately teaming with humans. This transition, from Human-Robot Interaction (HRI) to Human-Robot Collaboration (HRC) and Human-Robot Teaming (HRT) reflects increasing autonomy, shared decision-making and mutual adaptation. Within this evolution, imitation learning (IL) has become a key enabler for intuitive and efficient cooperation, allowing robots to acquire human strategies directly from demonstrations and operate effectively in dynamic, unstructured environments. This review outlines the technological foundations supporting this shift, including advances in sensing, actuation, and both classical and AI-driven control. By synthesising current methods and emerging trends, the paper provides an overview of the state-of-the-art and highlights IL as a promising pathway toward seamless HRT.

INTRODUCTION

Robotic systems are increasingly deployed in industrial environments where tasks are characterised by variability, uncertainty, and stringent safety requirements. As robots move from isolated, pre-programmed operation toward collaborative roles, the ability to learn from and adapt to human expertise becomes essential. Imitation learning (IL) has emerged as a promising paradigm to support this transition, enabling robots to acquire new skills directly from human demonstrations rather than through explicit programming. By leveraging multimodal perception and data-driven models, IL facilitates more intuitive, flexible, and efficient human-robot collaboration (Zare et al., 2024; Celemin et al., 2022).

Battery disassembly for recycling represents a particularly compelling application domain for IL-enabled collab-




Figure 1: The idea of using imitation learning (IL) for human-robot teaming (HRT) in battery disassembly.

oration. The process involves operations such as opening, loosening, cutting, and extracting components from heterogeneous and often unpredictable battery packs. These tasks are not only technically demanding but also pose safety risks related to thermal instability, residual charge, and mechanical fragility (Zheng et al., 2022). Human operators possess the dexterity, contextual awareness, and decision-making capabilities needed to manage such complexity, while robots offer precision, repeatability, and the ability to handle hazardous conditions (Kareer et al., 2025). IL provides a framework to merge these complementary strengths by transferring human strategies to robotic systems and enabling adaptive, shared execution of disassembly tasks (Wang et al., 2022; Tsuji et al., 2025).

This paper offers a perspective on how IL can support the evolution from Human-Robot Interaction (HRI) to more advanced forms of Human-Robot Collaboration (HRC) and physical Human-Robot Collaboration (pHRC), with a specific focus on battery disassembly (Kim et al., 2025). After outlining the fundamental principles of IL, we review key enabling technologies, including sensing, actuation, and control strategies, and discuss both traditional and AI-based approaches relevant to collaborative manipulation (Liu et al., 2025). We then highlight the opportunities and challenges associated with applying IL in battery recycling scenarios, identifying research directions

that can contribute to safer, more efficient, and scalable human–robot teaming (HRT) (Dai and Zeng, 2020; Jiang et al., 2025; Sanfilippo et al., 2025).

IL encompasses a family of techniques that enable robots to acquire new skills by observing or interacting with human demonstrations. Rather than relying on explicit programming, IL allows robots to infer task-relevant behaviours from examples, making it particularly suitable for environments where actions are difficult to formalise or where variability is high (Sasagawa et al., 2020). Classical IL approaches include behaviour cloning, which directly maps observations to actions, and inverse reinforcement learning, which infers the underlying objectives guiding human behaviour. More recent developments integrate deep learning, multimodal sensing and probabilistic models to improve generalization and robustness in real-world scenarios (Ding et al., 2021).

The evolution of human–robot systems provides the operational context in which IL becomes especially valuable. Early HRI focused on communication and coordination between humans and robots operating in adjacent spaces. HRC extends this paradigm by enabling shared workspaces and cooperative task execution, requiring robots to interpret human intent, adapt to dynamic conditions, and ensure safety (Li et al., 2024). pHRC represents a very integrated form of collaboration, where humans and robots may exchange forces, share manipulation tasks, or jointly handle objects. HRT is the the frontier of humans and robots coexistence, sharing space, forces, and decisions. Across this progression, IL offers a mechanism for transferring human expertise into robotic behaviour models that support fluid, intuitive collaboration.

Advances in sensing, actuation, and control have further expanded the applicability of IL in collaborative settings. Multimodal perception, combining vision, force sensing, wearable devices, and environmental sensors—provides the rich data needed to capture human demonstrations with sufficient fidelity. Complementary developments in compliant actuation and adaptive control strategies enable robots to reproduce learned behaviours safely and effectively in close proximity to humans. Together, these technological foundations position IL as a key enabler for next-generation collaborative systems, particularly in domains where tasks are complex, variable, or safety-critical.

This work offers three main contributions to the study of imitation learning for human–robot teaming in battery disassembly. The idea is shown in Figure 1. First, it provides a compact and structured overview of the enabling technologies that support IL in collaborative settings, highlighting the role of multimodal sensing, compliant actuation, and advanced control strategies in capturing and reproducing human expertise. Second, it clarifies how IL facilitates the progression from HRI to HRC and pHRC, outlining the mechanisms through which learned policies can support shared autonomy, adaptive role allocation, and safe physical interaction. Third, the paper introduces a conceptual pipeline that links demonstration acquisition,

policy learning, and collaborative execution, and grounds these ideas in a use case on collaborative panel lifting. This example illustrates how IL can support shared forces, dynamic authority, and robust coordination in a representative battery disassembly task, providing a foundation for future developments in circular-economy robotics.

ENABLING TECHNOLOGIES FOR IL in HRT

Imitation learning relies on a combination of sensing, actuation, and control technologies that allow robots to perceive human demonstrations, extract task-relevant information, and reproduce behaviours safely and adaptively. These enabling technologies form the backbone of IL-based collaboration, especially in domains where tasks are unstructured, variable, or safety-critical.

Sensing Technologies for IL

- **Vision:** is the primary channel for observing human demonstrations. RGB, depth, and stereo cameras enable the robot to track hand/tool motion, recognise objects, and interpret task context. These sensors support pose estimation, gesture recognition, and scene understanding, making them essential for extracting trajectories and identifying steps in manipulation tasks. Their flexibility and non-invasiveness make them suitable for both structured and semi-structured environments (Scoccia et al., 2024).
- **Force and tactile sensing:** demonstrations involve contact-rich interactions that cannot be fully captured through vision alone. Force/torque sensors, tactile arrays, and visuotactile systems provide information about applied forces, grip adjustments, and subtle variations in contact conditions. These signals help IL models encode how humans regulate pressure, adapt to resistance, or modulate compliance-skills that are crucial for safe manipulation.
- **Wearable sensors:** in scenarios where fine-grained motion must be captured or where occlusions are frequent, wearable sensors offer an additional source of high-quality data. Inertial measurement units (IMUs), electromyography (EMG) sensors (Zafar et al., 2025b), and marker-based motion capture systems provide kinematic and muscular information. These modalities are useful for capturing expert demonstrations of delicate or highly dexterous tasks.
- **Environmental sensing:** external sensors such as overhead RGB-D cameras, LiDAR, radar (Avazov et al., 2025), or acoustic arrays contribute global workspace awareness. They support multi-agent tracking, object localization, and detection of unexpected events. New sensing capabilities for HRT are offered by event-based cameras (Zafar et al., 2025d,a,c, 2026). When combined with onboard sensing, they enhance robustness of IL.

Actuation Technologies for IL

Robots must reproduce learned behaviours with a balance of precision and compliance:

- **Soft and hybrid actuators** introduce intrinsic compliance, reducing risks when handling fragile components (Hua et al., 2026).
- **Variable-stiffness and series-elastic actuators** allow dynamic modulation of interaction forces, mirroring human strategies.
- **Rigid actuators** ensure accuracy for tasks such as unscrewing or cutting.

Control Strategies for IL

Control architectures bridge IL outputs and robot dynamics:

- **Impedance and admittance control** ensure compliant interaction (Hua et al., 2025b).
- **Hybrid position/force control** supports tasks requiring precise alignment and controlled contact (Hua et al., 2025a).
- **Shared control** enables humans to guide or override robot actions, i.e., with haptics (Sanfilippo and Zafar, 2024).
- **AI-enhanced control** integrates IL with predictive models for safer adaptation.

Battery disassembly requires precise perception of fasteners, connectors, adhesive joints, and fragile components, as well as early detection of anomalies such as deformation or resistance changes. Multimodal sensing enables IL systems to capture expert strategies for handling these elements safely, making it a foundational component for deploying IL in recycling workflows.

IL ACROSS HRI, HRC AND pHRC

The progression from Human–Robot Interaction (HRI) to Human–Robot Collaboration (HRC) and ultimately to physical Human–Robot Collaboration (pHRC) reflects an increasing level of autonomy, shared control, and physical coupling between human and robot partners. Imitation learning plays a central role in enabling this transition by providing robots with the ability to acquire human-like behaviours and adapt them to collaborative contexts.

In traditional HRI, robots primarily interpret high-level human commands or gestures, with limited autonomy and no shared manipulation. IL enhances this paradigm by allowing robots to learn task-relevant motion patterns, enabling smoother and more intuitive interactions. In HRC, where humans and robots operate in shared workspaces, IL supports intent recognition, trajectory adaptation, and context-aware behaviour, reducing the cognitive load on

the human operator. In pHRC, where forces are exchanged directly and coordination must be continuous, IL becomes essential for learning compliant behaviours, force modulation strategies, and role-switching dynamics that are difficult to encode manually.

Across these stages, IL provides a unifying mechanism for transferring human expertise into robot policies that support fluid, adaptive, and safe collaboration—capabilities that are particularly relevant for complex industrial tasks such as battery disassembly.

IMITATION LEARNING FOR BATTERY DISASSEMBLY

Battery disassembly presents a unique combination of variability, uncertainty, and safety constraints that makes it an ideal yet challenging domain for imitation learning. Battery packs differ widely in geometry, fastening strategies, adhesive use, and internal layout, requiring flexible manipulation strategies that are difficult to encode through traditional programming. Moreover, operations such as opening, loosening, cutting, prying, and extracting modules involve contact-rich interactions where excessive force or misalignment can lead to mechanical damage or safety hazards.

Human operators excel in these tasks thanks to their dexterity, contextual awareness, and ability to adapt to unexpected resistance or anomalies. IL provides a mechanism to transfer these skills to robotic systems by capturing demonstrations that encode not only motion trajectories but also force regulation, compliance, and decision-making cues. Multimodal sensing—combining vision, force/torque feedback, and wearable data—enables the acquisition of rich demonstrations that reflect expert strategies for handling fragile components and responding to uncertain conditions (Pozzi et al., 2025).

In collaborative settings, IL supports shared manipulation, role switching, and adaptive load distribution between human and robot partners. These capabilities are essential for tasks such as lifting heavy battery panels, where the robot must coordinate forces with the human, yield authority when unexpected events occur, and resume support once the anomaly is resolved. IL thus serves as a foundation for developing safe, efficient, and scalable human–robot teaming strategies tailored to the demands of battery recycling.

A conceptual pipeline for IL enabled battery disassembly

The integration of imitation learning into battery disassembly workflows can be structured into a conceptual pipeline that links demonstration capture, policy learning, and collaborative execution. As shown in 2, this pipeline highlights the interplay between sensing, control, and human–robot teaming:

- **Demonstration acquisition** through multimodal sensing.




Figure 2: Pipeline for IL enabled battery disassembly.

- **Data processing and segmentation** into meaningful action primitives.
- **Policy learning** IL algorithms (e.g., behaviour cloning, inverse reinforcement learning, diffusion-based policies) learn mappings from states to actions, capturing both kinematic and force-based strategies.
- **Safety-aware execution** Learned policies are integrated with compliant control, force limits, and anomaly detection to ensure safe operation near humans and fragile components.
- **Human-robot teaming** The robot collaborates with the human by sharing forces, adapting trajectories, and dynamically adjusting authority based on task context.

This pipeline provides a structured foundation for deploying IL in real disassembly scenarios and directly motivates the panel lifting use case presented in the next section.

IL FOR HRT: PANEL LIFTING AS A USE CASE IN BATTERY DISASSEMBLY

To ground the concepts reviewed in this paper, a concrete mathematical formulation for a representative task in battery disassembly is presented: the collaborative lifting of a heavy battery panel by a human operator and a robotic arm acting as a team. This task encapsulates all three dimensions of Human-Robot Teaming (HRT) simultaneously, i.e., shared physical space, shared forces, and shared decision-making authority. This therefore serves as a unifying example for the control and imitation learning (IL) principles discussed throughout the review.

Task and Interaction Variables

A battery panel of mass m must be lifted from the floor, transported to a working height, held stationary while a

disassembly operation is performed, and then lowered to a resting position. The human and robot each grasp one end of the panel. At every instant the relevant observable state is the panel height y and vertical velocity $\dot{y}$, giving state vector $\mathbf{s} = [y, \dot{y}]^\top$. The panel dynamics along the vertical axis are:

$$m \ddot{y} = f_h + f_r - f_g, \quad (1)$$

where $f_h \geq 0$ and $f_r \geq 0$ are the upward forces of the human and robot respectively, and $f_g = mg$ is the gravitational load. For clarity we consider only vertical translation. Rotational dynamics induced by unequal grasp forces are neglected but can be incorporated through a coupled torque equation.

A defining characteristic of HRT as distinct from classical human-robot collaboration is that roles are not fixed. The same human and robot can alternate between leading and following, between carrying most of the load and sharing it equally, depending on the context, the state of the task, and the capabilities of each agent at that moment. This fluidity of roles is precisely what makes HRT powerful and what makes it difficult to engineer with classical controllers. Three scalar variables capture this dynamic at any instant:

- **Load-sharing ratio** $\alpha \in [0, 1]$: the fraction of the gravitational load f_g carried by the robot. When $\alpha = 1$ the robot fully compensates gravity; when $\alpha = 0$ the human carries the entire weight. In practice α shifts over time as fatigue, payload uncertainty, or task demands change.
- **Authority** $\lambda(t) \in [0, 1]$: the degree to which the robot directs the trajectory at a given instant. $\lambda = 0$ means the robot is purely reactive to the human; $\lambda = 1$ means the robot dictates the motion and the human provides passive guidance. Crucially, λ can transition continuously between these extremes within a single task execution, reflecting the dynamic role exchange that is the hallmark of HRT.
- **Internal force** $w_{\text{int}} = |f_h - f_r|$: the component of force that loads the panel without contributing to its motion. It arises when the two agents apply unequal forces and serves as a direct metric of coordination quality, a well-matched team keeps $w_{\text{int}} \approx 0$ regardless of which agent is currently leading.

These three variables are not independent: high robot authority typically implies high α (the robot both leads and carries), and poor coordination between α and the actual forces applied inflates w_{int} . An IL system that learns from expert demonstrations implicitly learns these relationships and their evolution over time, without requiring them to be explicitly programmed.

Demonstration Scenarios

A single demonstration of a nominal lift is insufficient to capture the full range of teaming behaviours a robot

must exhibit in practice. We define four scenarios, each eliciting a qualitatively different pattern of $(\lambda, \alpha, f_h, f_r)$ over time. Together they form a demonstration corpus that spans the space of collaborative lifting behaviours.

Scenario 1: Co-Lift

In the Co-Lift scenario both agents contribute equally to the trajectory ($\lambda \approx 0.5$) and share the gravitational load symmetrically ($\alpha \approx 0.5$ – 0.7). Neither agent leads; the motion emerges from balanced cooperation. This is the baseline demonstration of a functional human-robot team and produces training data in which $f_h \approx f_r$ throughout the lift, so $w_{\text{int}} \approx 0$. From an IL perspective this scenario teaches the robot the fundamental mapping from panel state to a supportive, matched force response.

Scenario 2: Human-Lead

In the Human-Lead scenario the human directs the pace and destination of the lift ($\lambda \approx 0.25$) while the robot primarily compensates gravity and follows the human’s force cues. Robot authority is low but load support remains high ($\alpha \approx 0.9$), reflecting the practical reality that even when the human leads the motion, fatigue reduction requires the robot to carry most of the weight. This scenario is representative of situations in which the human possesses task-specific knowledge the robot does not, such as the location of a partially detached connector or a non-standard panel orientation, and must retain directional control. The demonstration data show large f_h modulations driving the trajectory, with f_r tracking f_h with a lag that encodes robot compliance. An IL system trained on this scenario learns to be responsive to human intent rather than to impose its own trajectory.

Scenario 3: Robot-Lead

In the Robot-Lead scenario the robot takes initiative ($\lambda \approx 0.75$) and carries almost the entire gravitational load ($\alpha \approx 0.95$), while the human provides light oversight and corrective guidance. This mode is appropriate when the robot has a reliable model of the required motion, for instance, when executing a pre-planned placement into a battery rack at a known position, and the human’s role shifts from driver to supervisor. The demonstration data are characterised by large, consistent f_r driving the trajectory and small f_h providing minor corrections. Training on this scenario teaches the robot to take initiative and sustain a trajectory autonomously while remaining sensitive to human override signals.

Scenario 4: Anomaly

The Anomaly scenario captures the most demanding teaming behaviour: a mid-task fault, such as a stuck fastener, an unexpected panel resistance, or a safety alert, that requires the robot to immediately yield authority to the human. At the moment of fault detection at time t_{fault} , the robot

rapidly drives $\lambda \rightarrow 0$ and reduces f_r to near zero, transferring full physical and directional control to the human:

$$\lambda(t) \rightarrow 0, \quad f_r(t) \rightarrow 0, \quad t \gtrsim t_{\text{fault}}. \quad (2)$$

Once the human resolves the anomaly and signals clearance, authority and load support are gradually restored. This scenario is critical for IL because the yielding behaviour cannot easily be derived from first principles: it requires the robot to recognise a contextual signal (elevated resistance, unusual force patterns) and respond with a qualitative change in behaviour rather than a quantitative force adjustment. Demonstrations that include anomaly events expose the IL system to the full distribution of teaming states, including fault recovery, ensuring the learned policy is robust to conditions that fall outside the nominal lift trajectory.

IL Formulation

The four scenarios above collectively define a demonstration corpus $\mathcal{D}$ from which the robot learns a unified policy. The learning proceeds in three stages.

Demonstration Collection

Expert pairs (human-human or human-robot with a tele-operated robot) perform the lift under each scenario while sensors record the panel state and the forces at both grasp points. The full corpus is:

$$\mathcal{D} = \bigcup_{s \in \mathcal{S}} \{(\mathbf{s}^{(i)}, f_r^{(i)}, \lambda^{(i)})\}_{i=1}^{N_s}, \quad (3)$$

where $\mathcal{S} = \{\text{Co-Lift, Human-Lead, Robot-Lead, Anomaly}\}$ and N_s is the number of time steps in scenario s . No target heights, phase labels, or timing constraints are included: the temporal structure of each lift is entirely encoded in the demonstrated state-force trajectories.

Policy Training via Behaviour Cloning

Following the behaviour cloning paradigm (Pomerleau, 1988), the robot policy $\pi_\theta : (\mathbf{s}, \lambda) \mapsto f_r$ is trained by minimising the mean squared error between demonstrated and predicted forces:

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N \left(\pi_\theta(\mathbf{s}^{(i)}, \lambda^{(i)}) - f_r^{(i)} \right)^2, \quad (4)$$

where $N = \sum_s N_s$ is the total number of samples across all scenarios. Including λ as a policy input is essential: it allows the same network to produce high, sustained forces under Robot-Lead and near-zero forces under Anomaly, conditioned on the current authority state rather than learning a separate policy for each scenario.

Policy Execution

At runtime the robot queries the trained policy at each time step and applies the output force directly:

$$f_r = \pi_\theta(s, \lambda) - K_d \dot{y}, \quad (5)$$

where the velocity damping term $K_d \dot{y}$ provides closed-loop stability without requiring the policy to explicitly model velocity regulation. A safety floor $f_r \geq \alpha_{\min} f_g$ ensures the robot never fully releases the panel regardless of policy output. The trajectory that results is not prescribed: it emerges from the interaction between the learned policy, the live human intent encoded in f_h , and the panel dynamics, exactly as it did during the expert demonstrations.

From Single Demonstrations to Generalisable Policies

A policy trained only on behaviour cloning is known to suffer from *distribution shift*: at test time, small errors cause the robot to enter states not seen during training, leading to compounding failures (Ross et al., 2011). For the collaborative lifting task this is particularly acute because the human’s behaviour at test time will never exactly replicate the demonstration.

Several strategies mitigate this risk in the HRT context. Dataset aggregation methods such as DAgger (Ross et al., 2011) correct the policy online by querying the expert on states visited during robot execution. Residual policies (Silver et al., 2018) add a learned corrective term on top of a stable gravity-compensation base, so the worst-case failure mode is a held panel rather than a dropped one. Adversarial training approaches learn a policy that is robust to the worst-case human behaviour within a defined uncertainty set (Ho and Ermon, 2016). Augmenting the demonstration corpus with perturbation recoveries, brief intentional disturbances during demonstration, followed by expert correction, directly teaches the robot to recover from the distributional errors it will encounter at test time.

Table 1 summarises how IL compares to a classical engineered controller across the dimensions most relevant to battery disassembly. IL transfers the burden of task specification from the engineer to the demonstrator. For battery disassembly, where panels vary in weight, connectors vary in position, and operators vary in style, this adaptability is not a convenience, it is a requirement.

CONCLUSION AND FUTURE DIRECTIONS

The application of imitation learning (IL) to human–robot teaming (HRT) in battery disassembly raises a set of challenges that reflect the complexity of the task and the constraints of real industrial environments. Battery packs vary significantly in geometry, fastening strategies, adhesive use, and internal layout, making generalization a central difficulty for IL-based systems. Demonstrations collected on one battery type may not transfer directly to another, and the robot must be able to adapt to differences in stiffness, resistance, and component fragility. This variability

is compounded by the safety-critical nature of disassembly: excessive forces, misalignment, or delayed reactions to unexpected resistance can damage components or create hazardous conditions. Ensuring that learned policies remain safe under distribution shift, sensor noise, or human variability is therefore a fundamental requirement. Another challenge concerns the quality and quantity of demonstrations. IL methods must therefore become more data-efficient, capable of learning from sparse, imperfect, or partially observed demonstrations. At the same time, human–robot coordination introduces its own complexities. Effective teaming requires the robot to interpret human intent, modulate its contribution, and adjust authority dynamically, especially when unexpected events occur. These behaviours are difficult to encode explicitly and must emerge from a combination of learned policies and robust control strategies. Looking ahead, several research directions can strengthen the integration of IL into battery disassembly workflows. Advances in multimodal learning may enable richer fusion of vision, force, tactile, and wearable data, improving the robot’s ability to interpret demonstrations and respond to anomalies. Simulation and digital twins offer a promising avenue for generating synthetic demonstrations, reducing the burden on human experts and enabling large-scale data augmentation. Hybrid approaches that combine IL with reinforcement learning or model predictive control may improve adaptability and robustness, allowing robots to refine their behaviour online while maintaining safety guarantees. In conclusion, IL provides a powerful framework for enabling adaptive, intuitive, and safe HRT in battery disassembly. By leveraging multimodal sensing, compliant actuation, and advanced control strategies, IL allows robots to acquire human-like manipulation skills and collaborate effectively in tasks that are variable, contact-rich, and safety-critical.

Table 1: Classical Control vs. Imitation Learning for Collaborative Lifting

	Classical control		Imitation learning
Target heights	Pre-specified	by designer	Implicit in demonstration
Task timing	Pre-specified	by designer	Implicit in demonstration
Load sharing α	Fixed choice	design	Learned from expert pairs
Authority λ	Fixed or rule-based		Learned from demonstration
Anomaly response	Requires rules	explicit	Learned if demonstrated
New panel variants	Requires engineering	re-	Re-demonstrate and retrain
Operator variability	Not modelled		Averaged over corpus

The panel lifting use case illustrates how IL can support shared forces, dynamic authority, and robust coordination, offering a concrete example of how learned policies can enhance HRT. As research continues to address challenges related to generalization, safety, and data efficiency, IL is poised to become a key enabler for scalable and intelligent robotic systems in circular economy applications.

ACKNOWLEDGEMENT

This research is supported by the Artificial Intelligence, Biomechanics, and Collaborative Robotics research group, University of Agder (UiA), Norway.

References

- Nurilla Avazov, Rym Hicheri, Filippo Sanfilippo, and Matthias Pätzold. Time: Trajectory-informed movement emulation for arm gesture recognition using mmwave radar and optical sensor fusion. In *2025 IEEE 36th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–6. IEEE, 2025.
- Carlos Celemin, Rodrigo Pérez-Dattari, Eugenio Chisari, Giovanni Franzese, Leandro de Souza Rosa, Ravi Prakash, Zlatan Ajanović, Marta Ferraz, Abhinav Valada, and Jens Kober. Interactive imitation learning in robotics: A survey. *Foundations and Trends® in Robotics*, 10(1-2):1–197, 2022.
- Fei Dai and Rui Zeng. Physical examination of abdomen. *Handbook of Clinical Diagnostics*, pages 205–225, 2020.
- Ze Yang Ding, Junn Yong Loo, Vishnu Monn Baskaran, Surya Girinatha Nurzaman, and Chee Pin Tan. Predictive uncertainty estimation using deep learning for soft robot multimodal sensing. *IEEE Robotics and Automation Letters*, 6(2):951–957, 2021.
- Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29, 2016.
- Minh Tuan Hua, Jan Tommy Gravdahl, and Filippo Sanfilippo. Adaptive cartesian position control with a switching strategy for robotic manipulator with mixed rigid-elastic joints. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7398–7405. IEEE, 2025a.
- Minh Tuan Hua, Jan Tommy Gravdahl, and Filippo Sanfilippo. Impedance control for robots with mixed rigid-elastic joints using sliding mode control. In *2025 European Control Conference (ECC)*, pages 2914–2920. IEEE, 2025b.
- Minh Tuan Hua, Mohammad Poursina, Viet-Hung Vu, and Filippo Sanfilippo. Dynamic modelling and analysis of robotic manipulators with mixed rigid-elastic joints. *IEEE Access*, 2026.
- Zhenyu Jiang, Yuqi Xie, Kevin Lin, Zhenjia Xu, Weikang Wan, Ajay Mandlekar, Linxi Jim Fan, and Yuke Zhu. Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 16923–16930. IEEE, 2025.
- Simar Kareer, Dhruv Patel, Ryan Punamiya, Pranay Mathur, Shuo Cheng, Chen Wang, Judy Hoffman, and Danfei Xu. Egomimic: Scaling imitation learning via egocentric video. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13226–13233. IEEE, 2025.
- Ji Woong Kim, Juo-Tung Chen, Pascal Hansen, Lucy Xiaoyang Shi, Antony Goldenberg, Samuel Schmidgall, Paul Maria Scheikl, Anton Deguet, Brandon M White, De Ru Tsai, et al. Srt-h: A hierarchical framework for autonomous surgery via language-conditioned imitation learning. *Science robotics*, 10(104):eadt5254, 2025.
- Weidong Li, Yudie Hu, Yong Zhou, and Duc Truong Pham. Safe human–robot collaboration for industrial settings: a survey. *Journal of Intelligent Manufacturing*, 35(5):2235–2261, 2024.
- Wenhai Liu, Junbo Wang, Yiming Wang, Weiming Wang, and Cewu Lu. Forcemimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1105–1112. IEEE, 2025.
- Dean A Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- Andrea Pozzi, Alessandro Incremona, and Daniele Toti. Neural network-based imitation learning for approximating stochastic battery management systems. *IEEE Access*, 2025.
- Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- Filippo Sanfilippo and Muhammad Hamza Zafar. Human-robot teaming: A universal controller with multi-modal feedback for emergency response robots. In *2024 9th International Conference on Robotics and Automation Engineering (ICRAE)*, pages 181–187. IEEE, 2024.

Filippo Sanfilippo, Muhammad Hamza Zafar, Timothy Wiley, and Fabio Zambetta. From caged robots to high-fives in robotics: Exploring the paradigm shift from human–robot interaction to human–robot teaming in human–machine interfaces. *Journal of Manufacturing Systems*, 78:1–25, 2025.

Ayumu Sasagawa, Kazuki Fujimoto, Sho Sakaino, and Toshiaki Tsuji. Imitation learning based on bilateral control for human–robot cooperation. *IEEE Robotics and Automation Letters*, 5(4):6169–6176, 2020.

Cecilia Scoccia, Barnaba Ubezio, Giacomo Palmieri, Michael Rathmair, and Michael Hofbaur. Experimental assessment of a vision-based obstacle avoidance strategy for robot manipulators: Off-line trajectory planning and on-line motion control. *Journal of Intelligent & Robotic Systems*, 110(3):107, 2024.

Tom Silver, Kelsey Allen, Josh Tenenbaum, and Leslie Kaelbling. Residual policy learning. *arXiv preprint arXiv:1812.06298*, 2018. URL <https://arxiv.org/abs/1812.06298>.

Toshiaki Tsuji, Yasuhiro Kato, Gokhan Solak, Heng Zhang, Tadej Petrič, Francesco Nori, and Arash Ajoudani. A survey on imitation learning for contact-rich tasks in robotics. *arXiv preprint arXiv:2506.13498*, 2025.

Xu Wang, Sen Wang, Xingxing Liang, Dawei Zhao, Jincan Huang, Xin Xu, Bin Dai, and Qiguang Miao. Deep reinforcement learning: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 35(4):5064–5078, 2022.

Muhammad Hamza Zafar, Syed Kumayl Raza Moosavi, and Filippo Sanfilippo. Applications of neuromorphic/event camera in robotics with human in loop: A systematic review, datasets, and challenges. *IEEE Transactions on Human-Machine Systems*, 2025a.

Muhammad Hamza Zafar, Syed Kumayl Raza Moosavi, and Filippo Sanfilippo. Federated learning-enhanced edge deep learning model for emg-based gesture recognition in real-time human–robot interaction. *IEEE Sensors Journal*, 25(5):9139–9151, 2025b.

Muhammad Hamza Zafar, Syed Kumayl Raza Moosavi, and Filippo Sanfilippo. Matrix-har: Lightweight temporal motion-guided feature network with adaptive temporal resolution based transfer learning for resource-constrained event camera-based human action recognition. *Expert Systems with Applications*, page 130628, 2025c.

Muhammad Hamza Zafar, Syed Kumayl Raza Moosavi, and Filippo Sanfilippo. Neuromorphic event camera-based object recognition and grasping position detection using a transfer learning-enhanced

multi-task model. *IEEE Transactions on Automation Science and Engineering*, 2025d.

Muhammad Hamza Zafar, Syed Kumayl Raza Moosavi, and Filippo Sanfilippo. Curfed-charm: Curriculum sequential federated learning with a channel-hierarchical action recognition model for non-iid heterogeneous data in neuromorphic/event cameras. *Computer Vision and Image Understanding*, page 104696, 2026.


Maryam Zare, Parham M Kebria, Abbas Khosravi, and Saeid Nahavandi. A survey of imitation learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics*, 54(12):7173–7186, 2024.

Boyuan Zheng, Sunny Verma, Jianlong Zhou, Ivor W Tsang, and Fang Chen. Imitation learning: Progress, taxonomies and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, 35(5):6322–6337, 2022.

AUTHOR BIOGRAPHIES



FILIPPO SANFILIPPO holds a PhD in Engineering Cybernetics from the Norwegian University of Science and Technology (NTNU), Norway. He is a Professor at the Faculty of Engineering and Science, University of Agder (UiA), Norway, and heads the Artificial Intelligence, Biomechatronics, and Collaborative Robotics research group. His research focuses on Human-Robot Teaming (HRT). He is an IEEE Senior Member, former Chair of the IEEE Norway Section, and current Chair of the IEEE Robotics and Automation, Control Systems, and Intelligent Transportation Systems Joint Chapter. He serves on several IEEE Region 8 committees and leads the Norway Section Life Members Affinity Group. He has published widely in leading journals and conferences and serves as a reviewer for international publications.



CECILIA SCOCCIA Cecilia Scoccia received the Ph.D. degree in Industrial Engineering from the Università Politecnica delle Marche, Ancona, Italy, where she is currently a Researcher with the Department of Industrial Engineering and Mathematical Sciences (DIISM). Her research interests encompass collaborative robotics, human–robot interaction, and the ergonomic evaluation of passive exoskeletons, with emphasis on human-centered and safe collaborative work environments. She also investigates bio-inspired underwater robotic systems for improved AUV maneuverability and contributes as a Guest Editor for international journals. Her current work includes the integration of digital twins and advanced sensing.