Artificial Intelligence in Collaborative and Industrial Robotics

Adam Burn[1], *Michael* Short[2], *Maher* Al-Greer[3], *Arockia Selvakumar Arockia* Doss[4]

[1]PhD Researcher, School of Computing, Engineering and Digital Technologies, Teesside University, UK.
[2]Professor, School of Computing, Engineering and Digital Technologies, Teesside University, UK
[3]Associate Professor, School of Computing, Engineering and Digital Technologies, Teesside University, UK.
[4] Professor, School of Mechanical Engineering, Vellore Institute of Technology, Chennai, India.

Abstract. Recent advances in artificial intelligence are reshaping collaborative and industrial robotics, enabling a transition from deterministic, pre-programmed automation toward adaptive, learning-enabled systems. This paper synthesises developments in imitation learning, diffusion-based visuomotor policies, and foundation models, and examines their integration within industrial robotic architectures. Particular attention is given to the convergence of language-based planning, multimodal perception, and digital twins for safe and flexible deployment. Electric vehicle battery recycling is considered as a representative high-variability and safety-critical case study, illustrating how contact-rich manipulation, sim-to-real transfer, and certified runtime supervision can be combined within a unified framework. It is argued that the same AI stack supporting flexible assembly in manufacturing can be extended to other related areas, such as disassembly-related circular-economy processes. Open challenges remain in safety certification, explainability, data scarcity, and multi-material interaction modelling. Future directions include cognitive digital twins, tactile foundation models, federated learning, and multi-robot coordination. The convergence of learning-based control and industrial digital infrastructures provides a pathway toward resilient and sustainable Industry 5.0 production systems.

1 Introduction

Industrial robots were historically designed for fixed, repetitive tasks using deterministic joint and trajectory control, with model-based sensor feedback. While highly effective in structured environments, traditional automation struggles with more advanced applications and uncertainties, e.g. due to product variability (requiring rapid reconfiguration), hazardous environments, complex manipulation tasks, mobile robotics, and sustainability-driven disassembly tasks.

---

*Corresponding Author: M.Short@tees.ac.uk

Simultaneously, manufacturing and recycling industries face labour shortages, stricter safety standards, and sustainability/circular-economy pressures. These constraints, along with breakthroughs in deep learning enabled perception and decision-making from raw data, motivate the integration of learning-based AI into robotic systems [1].

Recent developments in imitation learning, diffusion-based control policies, and foundation models for robotics enable robots to learn contact-rich skills from demonstration, interpret semantic task instructions, and generalise across related manipulation tasks. When combined with digital twins and certified safety layers, these approaches provide a viable framework for adaptive industrial robotics.

In contrast to conventional industrial robotic architectures, the framework proposed in this paper introduces a layered integration of learning-enabled modules while preserving standards-aligned industrial control principles. This allows adaptive task interpretation and skill generalisation without replacing the bounded and certifiable execution structures required in industrial settings. The proposed framework should therefore be understood not as a departure from established industrial robotic architectures, but as an extension of them toward greater flexibility, semantic awareness, and task adaptability.

This presented unified architecture links collaborative and industrial robotic tasks under a common learning-based robotic stack, using EV battery disassembly for recycling as a representative application. The contribution of this paper is threefold. First, it proposes a unified architectural perspective for integrating imitation learning, diffusion-based visuomotor policies, foundation-model-based task planning, digital twin validation, and deterministic safety supervision within industrial robotic systems. Second, it examines how this architecture may generalise across both manufacturing and disassembly domains, using EV battery recycling as a representative safety-critical application. Third, it identifies practical validation requirements, deployment constraints, and open research challenges that must be addressed before such systems can be robustly adopted in industrial environments. The purpose of this paper is not to exhaustively review all AI-enabled robotics paradigms, but to synthesise a coherent industrial systems perspective across those components most relevant to adaptive manipulation, safety, and deployment.

The remainder of this paper is structured as follows. Section 2 presents a short overview of related literature on learning-based robotics, and Section 3 discusses the role of digital twins and simulation, and outlines the proposed framework. Section 4 presents an EV battery case study as an illustrative example. Section 5 presents cross-sector approaches and application domains, along with relevant standards. Current challenges and future directions are presented in Sections 6 and 7 respectively.

Research Questions:

RQ1: How can learning-based robotic components be integrated with deterministic industrial control while preserving safety, traceability, and deployment realism?

RQ2: To what extent can a unified AI-enabled robotic architecture support both industrial assembly and disassembly domains, including EV battery recycling?

RQ3: What technical, safety, and data-related barriers most constrain the industrial deployment of such architectures?

2 Learning-driven industrial robotics

2.1 Imitation Learning in industrial settings

Imitation learning (IL) allows robots to acquire control policies from expert demonstrations rather than through trial-and-error reinforcement learning. Behaviour cloning and dataset aggregation approaches provide data-efficient skill acquisition suitable for hardware deployment [2].

In manufacturing, IL enables encoding of tacit human skills such as compliant insertion, torque-controlled fastening, and surface finishing. In disassembly tasks, similar methods support prying, peeling, and fastener removal operations where contact dynamics are difficult to model explicitly [3]. An example is shown in Fig. 1.

Fig. 1. Example of Imitation Learning. A human demonstrating task execution and a robot through learning from demonstration [4].

2.2 Diffusion-based visuomotor policies

Diffusion models have recently been adapted for robotic control, treating action sequences as a denoising generation process. These models produce smooth and multimodal trajectories, improving robustness in contact-rich manipulation [5]. Task decomposition is illustrated in Fig. 2.

Benchmark studies demonstrate improved performance over policy-gradient baselines in tasks involving tool-based interaction. Such properties are particularly relevant to contact-rich and multi-material disassembly in battery packs [6].

Fig. 2. Example of task decomposition through diffusion-based model policies[7]

2.3 Foundational models and language-enabled planning

Vision-language-action models such as RT-1/RT-2 and PaLM-E demonstrate that large-scale pretraining enables semantic reasoning grounded in perception and action. Language-conditioned policies allow robots to interpret high-level task instructions and generate executable plans [8, 9]. A general architecture is shown in Fig. 3.

In industrial contexts, this capability supports dynamic interpretation of work orders, retrieval of process parameters, and flexible sequencing across product variants. When combined with affordance filtering or feasibility scoring, LLM-based planners can operate within constrained industrial environments [10, 11].

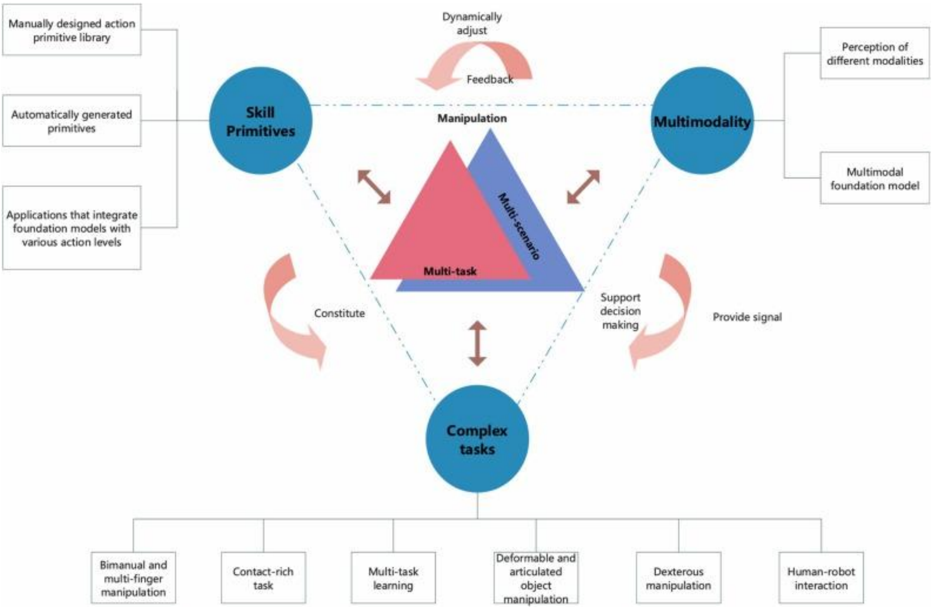

Fig. 3. Example of a general architecture for robot learning through foundational models [12].

2.4 Discussion and Comparison of Candidate Robot Control Approaches

To provide a clearer quantitative perspective, Table 1 summarises representative performance characteristics reported across the literature for imitation learning, diffusion-based visuomotor policies, and large-scale language-conditioned robotic models. Although not a direct benchmark under identical conditions, these values illustrate comparative tendencies amongst each approach.

Table 1. Representative performance characteristics of learning-based robotic approaches relevant to industrial manipulation

Approach	Typical strengths	Common quantitative indicators	Reported limitations for industrial use
Imitation Learning (Behaviour Cloning / DAgger)	Data-efficient skill acquisition from expert demonstrations; suitable for structured manipulation	Often evaluated using task success rate, trajectory error, intervention frequency, and number of demonstrations required	Performance may degrade under distribution shift, unseen task variants, or contact conditions not represented in training data
Diffusion-based visuomotor policies	Smooth multimodal action generation; strong performance in contact-rich manipulation	Common metrics include task success rate, trajectory smoothness, robustness to perturbation, and action consistency across trials	High data and compute requirements; limited interpretability; sim-to-real robustness still task dependent
Vision-Language-Action / Foundation Models	High-level semantic reasoning, instruction following, cross-task generalisation	Often evaluated using instruction completion rate, zero-shot generalisation, and task transfer performance	Requires large-scale multimodal datasets; difficult to validate in safety-critical execution without supervisory constraints

Hybrid industrial architecture (proposed framework)	Combines adaptive planning with deterministic safety and execution	Suggested metrics include task success rate, cycle time, safety intervention frequency, force/torque exceedance rate, and recovery from faults	Experimental validation across multiple industrial tasks remains future work

Taken together, these approaches occupy different practical roles within industrial robotics. Imitation learning remains comparatively attractive for structured manipulation tasks where demonstrations are available and deployment simplicity is important, but its robustness may degrade during tasks involving high amounts of product variance (such as heterogenous recycling feedstock). Diffusion-based policies offer stronger performance in contact-rich and multimodal settings, albeit with higher computational and data requirements. Foundation-model-based systems provide valuable semantic flexibility and task-level generalisation, but currently remain difficult to validate for direct safety-critical execution. As a result, their most practically viable near-term role in industrial environments is likely to be at the planning and supervisory layer rather than as fully autonomous low-level controllers.

3 Digital twins and robot safety

Beyond manipulation policies alone, the integration of LLM and agent-based AI architectures within digital twin infrastructures and collaborative robotics further reshapes the deployment of learning-enabled systems in industrial environments. Agentic robotic frameworks extend conventional control pipelines by embedding language-based reasoning modules capable of interpreting structured work flow instructions, querying production or relevant disassembly databases, and decomposing high-level objectives into executable skill sequences [11, 13, 14]. When coupled with digital twins, these reasoning processes can be validated in simulation, allowing candidate tool paths, force profiles, and task orderings to be assessed against geometric and process constraints prior to physical execution. In keeping with the international safety standards for industrial robots [15], this pre-execution validation reduces operational risk and supports traceability, particularly in the context of handling hazardous materials and safety-critical disassembly tasks. Within collaborative settings, where robots operate in shared workspaces with human operators, additional safeguards become essential. Modern safety standards require that adaptive behaviours remain verifiably bounded, leading to the adoption of layered supervision mechanisms such as control barrier functions [16], predictive safety filters, and monitored stop strategies [17].

From a systems perspective, the digital twin provides a structured intermediate layer between high-level AI reasoning and physical robot execution. Rather than allowing candidate actions generated by semantic planning or learned manipulation policies to be executed directly on hardware, the twin enables pre-execution checking against geometric, kinematic, and process-level constraints. This may include assessment of tool access, fixture interference, trajectory feasibility, and predicted contact behaviour under nominal operating conditions. In industrial contexts, such validation is particularly important because deployment failures can have direct consequences for safety, equipment integrity, and process reliability; the twin therefore functions not only as a simulation tool, but also as a mechanism for reducing deployment risk and improving traceability during iterative system development.

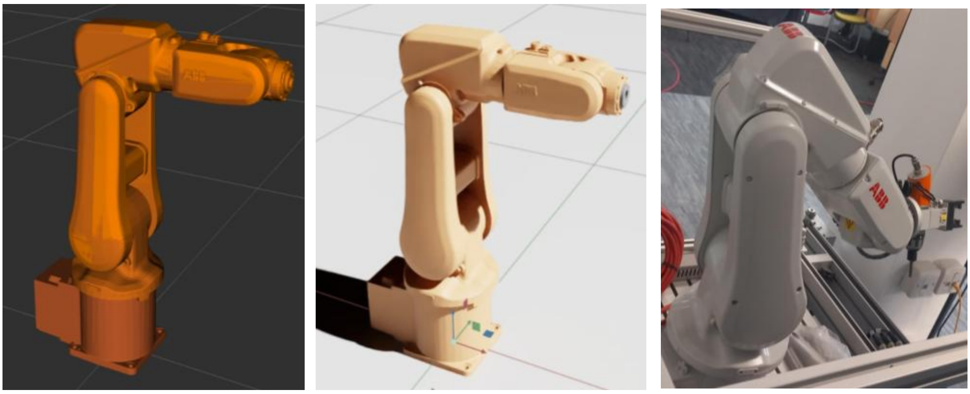

Robot Operation system built on the ROS2 framework

High Fidelity data generation in Isaac Sim

Training data validated with real robot in the lab

Fig. 4. High-level illustration of how robot control and manipulation data may be generated, refined, and validated using digital twin environments prior to physical deployment.

3.1 Proposed System Workflow

In this proposed architecture, the stochastic outputs of learning or language-based components are constrained by deterministic safety envelopes, preserving compliance while maintaining flexibility. The convergence of agentic planning, simulation-based validation, and certified collaborative control thus represents a pragmatic pathway for integrating advanced AI capabilities into real and generalisable systems without compromising safety or regulatory requirements. The proposed architecture is summarised in Fig. 5

Within the proposed architecture, task execution may be structured as a staged pipeline:
(1) A work instruction, operator request, or task objective is received by a semantic planning layer.
(2) Relevant contextual information is retrieved from structured databases, digital documentation, or metadata sources (such as a battery's BMS).
(3) A language-enabled planning module decomposes the task into candidate sub-actions and manipulation steps.
(4) Appropriate learned manipulation skills, such as imitation-learned or diffusion-based policies, are selected or conditioned based on the task context.
(5) Candidate trajectories and task sequences are evaluated within a digital twin environment to assess geometric feasibility, predicted force margins, and safety rule compliance.
(6) Validated motion plans are then passed to deterministic industrial control layers for execution.
(7) Runtime supervisory safety mechanisms monitor force, collision risk, and execution constraints during operation.
(8) Telemetry, execution logs, and task outcomes may then be stored to support traceability, adaptation, and future model improvement.

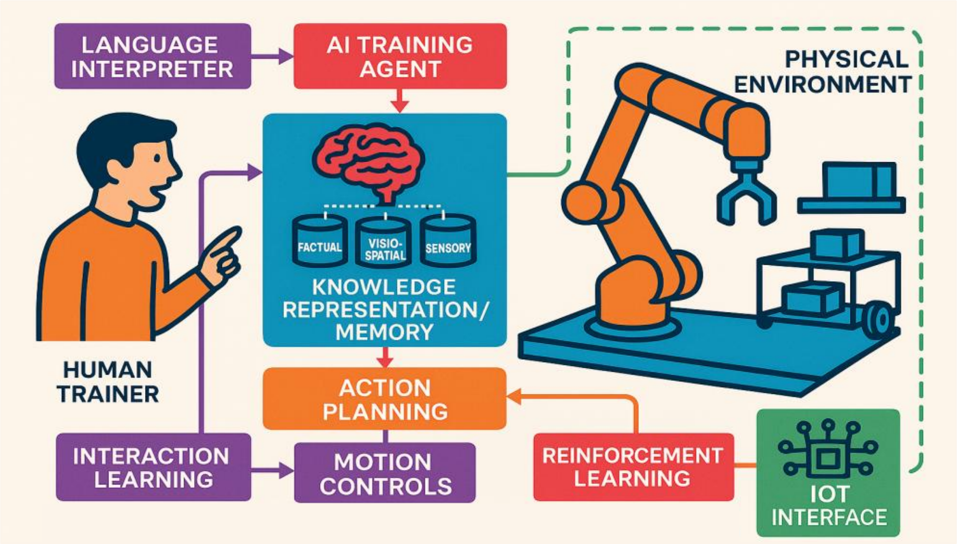

Fig. 5. Preliminary generalized robotic architecture.

Figure 5 is designed to be interpreted as a layered execution pipeline; the upper layers correspond to semantic reasoning, context retrieval, and candidate action generation; the middle layers to policy selection, simulation-based validation, and task sequencing; and the lower layers to deterministic motion execution, sensing, and safety supervision. Foundation-model-based components operate primarily at the planning level, while imitation-learned and diffusion-based policies provide adaptable manipulation skills for contact-rich tasks. Deterministic industrial controllers remain responsible for final execution, with supervisory safety layers enforcing bounded operational constraints; this layered separation preserves adaptability while maintaining industrially realistic execution and safety control.

Existing industrial robotic systems are generally characterised by high reliability, repeatability, and standards compliance, but often depend on extensive task-specific programming and limited adaptability to product variation, unstructured interactions, or contact-rich manipulation [15]. The framework proposed here differs by introducing a hierarchical integration strategy in which learning-based modules are used for perception, semantic task decomposition, and policy generation, while deterministic industrial control remains responsible for final execution and safety enforcement [10, 11, 16, 17]. In this sense, adaptive AI components do not replace industrial control logic, but instead augment it with improved generalisation and context awareness. This is particularly relevant in manufacturing and disassembly environments where workpiece variability, uncertain contact conditions, and incomplete process observability reduce the practicality of rigid pre-programmed workflows [5, 14, 18].

In addition, the architecture is designed to support context-aware robotic operation through links to structured databases, operational metadata, and digital twin environments [13, 19]. Such integration enables candidate actions to be checked against task constraints, process knowledge, and safety requirements before execution [13, 14, 18, 20]. This layered design also improves traceability, allowing robotic decisions to be linked to retrieved context, simulation validation, and supervisory constraints rather than relying solely on opaque end-to-end policy execution [19, 20].

While the core components of the proposed framework have each been studied individually in recent robotics literature, their integration within a unified industrially oriented architecture remains less systematically explored than the individual component technologies themselves [4, 5, 8–11, 16–18]. The contribution of the present work is therefore not the proposal of a new standalone algorithm, but the synthesis of these traditionally separate capabilities into a unified industrially oriented architectural perspective. With that in mind, this architecture is a conceptual systems framework intended to guide future integrated validation in simulation and physical deployment, rather than as a fully implemented and experimentally validated robotic system.

As summarised in Table 2, many existing robotics frameworks emphasise only a subset of the capabilities required for industrially deployable adaptive robotics. Language-conditioned planning frameworks provide strong semantic reasoning but are not typically coupled to bounded industrial execution; visuomotor learning approaches offer adaptable manipulation skills but often lack traceability, safety assurance, and deployment realism. Digital twin-based systems provide valuable simulation and validation capabilities, yet may not incorporate flexible semantic reasoning or learned skill adaptation.

Table 2. High-level comparison between representative AI-enabled robotic frameworks and the unified architecture proposed in this paper (compiled from [2, 5, 10, 11, 13, 15, 18])

Framework / approach	**Semantic task planning**	**Learned manipulation policy**	**Digital twin / simulation validation**	**Deterministic safety / execution layer**	**Industrial deployment orientation**	**Key limitation**
Language-conditioned robotic planning	Yes	Partial	No	Partial	Limited	Strong semantic planning, but weak industrial validation and bounded execution guarantees
Visuomotor policy learning	Partial	Yes	Partial	No	Limited	Strong skill learning, but limited traceability, safety assurance, and task-level reasoning
Digital twin-enabled robotic workflows	Partial	Partial	Yes	Partial	Moderate	Strong validation and simulation capability, but often limited adaptive semantic reasoning

Conventional industrial robotic architectures	No	No	Partial	Yes	Yes	High reliability and safety, but limited adaptability and dependence on task-specific programming

4 Case study: EV battery recycling

EV battery recycling constitutes a particularly demanding manipulation domain characterised by high variability and significant safety constraints. Battery packs (example shown in Fig.6) differ substantially in geometry, fastening strategy, adhesive usage, and internal architecture, and may arrive in varying states of degradation. In addition to mechanical complexity, residual electrical charge and the presence of flammable electrolytes introduce hazards that are not typically encountered in conventional assembly lines [14]. These factors make manual disassembly labour intensive and potentially unsafe, particularly at scale.

Fig. 6. Example of end-of-life Nissan Leaf battery modules and cells.

Learning-enabled robotic systems offer a structured response to this variability. By conditioning manipulation policies on pack configuration data, robotic platforms can adapt tool selection and motion strategies to different architectures without complete reprogramming [5, 9]. During separation tasks, force and torque signatures provide real-time feedback that could be used to detect fastener release, adhesive failure, or abnormal resistance [21]. Integration with battery management system data [13] further enables context-aware decision-making, for example by prioritising modules for second-life assessment based on state-of-health metrics, as illustrated in Fig. 7. Prior to physical execution, disassembly sequences may be validated within digital twin environments to estimate interaction forces and ensure compliance with safety constraints.

From a control and architectural perspective, these operations are closely related to assembly processes. The same imitation and diffusion-based policies used to achieve compliant insertion and torque-regulated fastening can be adapted to perform controlled extraction and separation [2, 5, 9]. This symmetry highlights the transferability of learning-based robotic architectures across production and circular-economy contexts, supporting the development of unified skill libraries and shared validation frameworks.

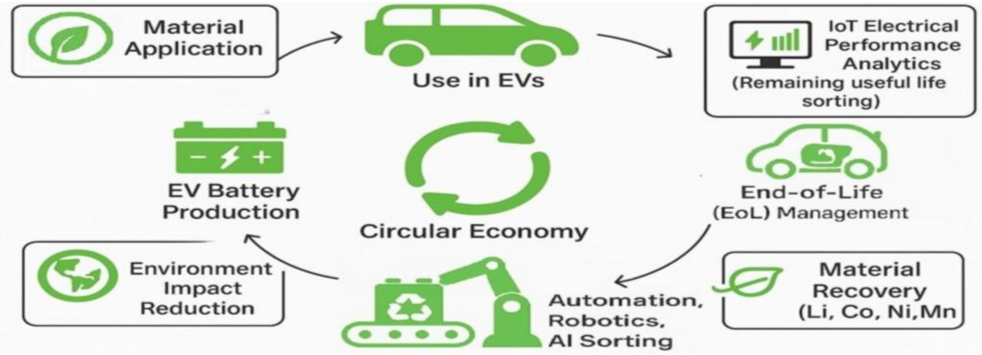

Fig. 7. Summarised pipeline for a closed-loop autonomous EV supply chain.

To strengthen the practical relevance, and enable future empirical assessment of the proposed framework, EV battery disassembly may be evaluated using a set of quantitative performance metrics aligned with industrial deployment requirements. Suitable measures include: (i) task success rate across different battery pack variants, (ii) average cycle time per disassembly stage, (iii) tool-change frequency, (iv) fastener or adhesive separation success rate, (v) force/torque exceedance frequency relative to safe operating thresholds, (vi) rate of abnormal event detection and recovery, and (vii) human intervention frequency per completed task sequence [14, 21].

Additional domain-specific indicators may include module identification accuracy, state-of-health informed prioritisation accuracy, and damage rate during extraction or separation operations. In a digital twin context, trajectory feasibility, predicted contact force margins, and safety rule violations may also be assessed prior to physical execution [13, 18]. Collectively, these metrics would provide a more rigorous basis for evaluating whether a unified learning-enabled architecture offers practical performance benefits over conventional manually programmed disassembly workflows. This case study is intended primarily as an illustrative application domain for the proposed architecture. However, future experimental work should explicitly quantify these metrics across multiple battery formats and degradation conditions to assess system robustness, transferability, and operational safety [14].

5 Cross Sector Architectures

One of the key findings of this review is that the same AI stack applies across sectors (summarized in Table 3). As discussed in sections 2-4, manufacturing and recycling both use imitation learning, diffusion policies, digital twins, and language-based planners; they simply operate on different tasks and data sources.

Table 3. Overlap in automatous robotic technologies in both industrial assembly tasks and EV battery recycling

Layer	Industrial Assembly	EV Recycling Cell
Perception	Vision, force, acoustic sensing	Vision, force, thermal/voltage sensors
Learning	IL / Diffusion policy for assembly primitives	IL / Diffusion policy for disassembly

Planner	LLM agent for interpreting work orders	LLM agent for interpreting pack meta data
Simulation	Digital twin for fixture and tool validation	Digital twin for pack geometry, safety analysis and tool validation
Safety	ISO 10218, ANSI/RIA compliance	ISO 10218 and handling rules
Database Link	Manufacturing execution system / enterprise resource planning	Battery Management system / recycling database
Outcome	Flexible production	Adaptive and sustainable material recovery

This structural similarity supports a unified learning-based robotics paradigm, capable of deployment across both production and recovery systems.

6 Challenges

Despite substantial progress, several limitations continue to constrain the industrial deployment of learning-enabled robotic systems. These limitations suggest that the remaining barriers to widespread adoption are increasingly linked to validation methodology, system integration, and governance considerations rather than being solely attributable to fundamental shortcomings in learning algorithms themselves.

6.1 Safety Considerations

While deterministic control systems can be analysed using established verification and certification methodologies, stochastic neural controllers and diffusion-based policies are considerably more difficult to validate within current industrial safety frameworks [15–17]. In practice, safe deployment therefore depends on architectural separation between adaptive policy generation and bounded execution. Within the proposed framework, learning-based components are used to generate candidate actions or task plans, while deterministic supervisory layers remain responsible for constraint enforcement, motion feasibility, collision avoidance, and emergency interruption [16, 17]. Runtime safety filters, control barrier functions, and monitored stop strategies provide practical mechanisms for ensuring that adaptive behaviours remain within acceptable operational envelopes. Nevertheless, a broadly accepted certification methodology for learning-enabled industrial robotics has yet to emerge, particularly for systems involving contact-rich manipulation and semantic planning.

In practice, compliance-oriented deployment would likely require architectural partitioning between certifiable deterministic execution layers and non-certifiable adaptive modules. Candidate plans or actions generated by learning-based systems could be validated offline or pre-execution against workspace, force, and collision constraints, while runtime monitors enforce operational boundaries during execution. Under this model, certification would apply primarily to bounded execution, supervision, and fail-safe behaviour, rather than to unrestricted learning outputs themselves. This reflects the current reality that fully general certification pathways for stochastic planning and control remain immature.

6.2 LLM transparency

Language-based planning modules introduce a semantic reasoning layer whose internal decision processes are not inherently transparent [10, 11, 19, 20]. In safety-critical applications, this opacity presents challenges for traceability, accountability, and operator trust, particularly when task sequencing decisions or tool selections must be justified retrospectively [19, 20]. Within the proposed framework, explainability is addressed procedurally through structured task decomposition, retrieval-grounded reasoning, explicit simulation validation, and logged supervisory interventions [10, 11, 13, 17]. Such mechanisms can improve auditability by linking robotic actions to retrieved context and validated execution paths. However, these measures do not fully resolve the broader interpretability limitations of large-scale foundation models, and explainability therefore remains a key open challenge for industrial AI deployment [19, 20].

6.3 Computational Complexity, Data Scarcity and Foundation Model Requirements

Industrial datasets are typically proprietary, fragmented, and expensive to acquire, restricting large-scale training and cross-site generalisation [14, 19, 21]. This challenge is especially pronounced for foundation-model-style robotic systems, which may require diverse multimodal datasets spanning visual observations, robot state trajectories, force/torque signals, language-conditioned task descriptions, and outcome labels across many task variants [8, 9, 12, 19]. In industrial settings, the challenge is therefore not only one of data volume, but also of annotation cost, confidentiality, operational safety during collection, and deliberate representational coverage of rare but safety-relevant events [14, 19]. While simulation, digital twins, transfer learning, and synthetic data generation may alleviate some of these constraints [13, 21], the data requirements for robust industrial foundation models remain substantially higher than those associated with traditional fixed-task automation [8, 9]. In practical deployment, computational requirements also impose an important constraint on system design. Large-scale foundation models and diffusion-based planners may be unsuitable for direct high-frequency control loops because of inference latency and resource demands. Their most practical role is therefore likely to be in asynchronous planning, skill selection, or pre-execution reasoning, while deterministic industrial controllers continue to operate at real-time control rates.

6.4 Simulation Fidelity and Sim-to-Real Transfer

Accurate sim-to-real transfer remains a major barrier to deploying learning-enabled robotic systems in industrial environments. Differences between simulated and real systems may arise from sensor noise, actuator latency, calibration mismatch, fixture tolerances, frictional uncertainty, and incomplete modelling of environmental disturbances. These issues are amplified in contact-rich tasks where task success depends on subtle force interactions rather than purely geometric motion execution. Although digital twins and domain randomisation can reduce this gap, learned policies remain vulnerable to distribution shifts when exposed to previously unseen physical conditions. In industrial contexts, this challenge is especially significant because deployment failures may have direct implications for safety, equipment integrity, and process reliability [14, 21].

A further challenge lies in uncertainty-aware decision-making. In industrial environments, uncertainty may arise from incomplete perception, hidden fasteners, degraded materials, inconsistent pack conditions, sensor noise, or unforeseen contact responses. Robust deployment therefore requires not only learned task competence, but also mechanisms for confidence estimation, anomaly detection, fallback behaviours, and safe recovery when predicted execution confidence is low. Such uncertainty-aware supervisory behaviour remains comparatively underexplored in many current AI-enabled robotic systems.

6.5 Multi-Material Interaction Modelling

Robust modelling of multi-material interaction remains particularly difficult for both simulation and learning-based control [5, 18, 22]. Industrial disassembly tasks frequently involve adhesives, polymers, foils, elastomers, coatings, and heterogeneous metal assemblies whose interaction behaviour is nonlinear, discontinuous, and often poorly observable [14, 18, 22]. Such systems are not easily captured by conventional rigid-body simulation, and even force-conditioned learning policies may struggle to generalise when material properties vary across product variants or degradation states [5, 18, 21, 22]. These limitations are especially relevant in EV battery disassembly, where layered structures and uncertain fastening conditions can significantly affect contact behaviour, force signatures, and extraction dynamics [16]. As a result, multi-material interaction modelling remains one of the clearest technical limitations of current AI-enabled industrial robotic systems.

7 Future directions

Several research directions appear particularly relevant for advancing learning-enabled industrial robotics. These developments suggest a gradual progression toward increasingly adaptive and self-optimising production cells. However, fully autonomous industrial ecosystems remain a long-term objective, dependent not only on algorithmic advances but also on validation frameworks, regulatory evolution, and carefully designed human–AI collaboration models [19, 23].

7.1 Adaptive digital twins and multimodal representation learning

Digital twins may pivot from static geometric or kinematic replicas toward adaptive models that are continuously updated using real-time telemetry [18]. Such systems could incorporate residual error learning, parameter adaptation, and predictive modelling, allowing tighter coupling between symbolic planning and physical system behaviour. In parallel, growing interest in tactile and force-based representation learning suggests that foundation models trained on large-scale contact data may significantly improve robustness in manipulation tasks where visual information alone is insufficient. This is especially pertinent for processes involving deformable or multi-layered materials, such as battery module handling and adhesive separation [8, 21].

Future validation should assess whether simulation-based pre-execution checking measurably improves deployment outcomes, for example through reductions in failed trajectories, unsafe contact events, cycle time variability, or unplanned interventions. Such evaluation would help determine whether digital twin integration offers measurable practical benefit beyond conventional offline programming or static fixture validation.

7.2 Distributed learning and coordinated robotic systems

Federated learning frameworks offer a potential solution to the confidentiality constraints that limit cross-site data sharing in industrial environments. By enabling distributed model improvement without centralised dataset aggregation, such approaches may support collaborative performance gains while preserving proprietary information. Complementary approaches such as transfer learning and modular policy adaptation may further reduce the need for extensive retraining when deploying robotic systems across different product variants, facilities, or hardware platforms, thereby improving scalability.

At the same time, structured multi-robot coordination and communication protocols are likely to become increasingly important as factories and recycling facilities adopt teams of collaborative and industrial manipulators operating within shared workspaces [24, 25]. The proposed unified architecture is conceptually extensible to such multi-robot environments, where coordinated agents may share task workloads, tooling resources, and workspace regions. In these settings, scalability depends not only on policy transferability, but also on communication latency, scheduling efficiency, collision-aware coordination, and consistency of shared task representations across agents [16, 24, 26]. Future work should therefore evaluate how the framework performs when extended from single-cell robotic operation to coordinated multi-robot assembly or disassembly systems, particularly under dynamic task allocation and shared safety constraints [26, 27].

8 Conclusions

Artificial intelligence has broadened the capabilities of industrial robotics beyond rigid automation toward adaptive, multimodal, and semantically informed systems which can dynamically reconfigure in unstructured and uncertain environments. The integration of imitation learning, diffusion-based control, language-enabled planning, digital twins, and certified safety supervision reveals a coherent architectural paradigm which appears applicable across advanced manufacturing domains, including EV battery recycling. This paper has argued that imitation learning, diffusion-based control, language-enabled planning, digital twins, and deterministic safety supervision can be understood as complementary components of unified industrial robotics architecture. The principal contribution is therefore not the proposal of a new standalone learning algorithm, but a systems-level architectural synthesis intended to support adaptive yet bounded robotic deployment across manufacturing and disassembly contexts.

The present work is itself subject to several limitations. Most notably, it remains conceptual and does not include integrated experimental validation of the proposed architecture. The discussion synthesises multiple rapidly evolving AI subfields and therefore necessarily abstracts over some implementation-specific details. The paper should therefore be understood as a systems-level conceptual contribution intended to structure future validation, rather than as a finalised deployment-ready methodology. Future work will focus on integrated simulation and physical testing of the proposed concepts.

In addition to its conceptual contribution, the proposed framework highlights several practical requirements for future validation, including comparative benchmarking of learning paradigms, quantitative case-study evaluation, explainability under safety-critical constraints, and assessment of scalability to coordinated multi-robot systems. Addressing these challenges will be essential in determining whether unified AI-enabled architectures can transition from promising research prototypes toward robust industrial deployment.

The convergence of these technologies suggests a gradual transition toward more flexible and sustainable production ecosystems. However, continued progress will depend on advances in safety validation, interpretability, simulation fidelity, and human–AI integration. Learning-enabled robotics therefore represents not a completed transformation, but an evolving framework shaping the future of industrial and circular-economy automation.

References

1. Mnih V, Kavukcuoglu K, Silver D, et al (2015) Human-level control through deep reinforcement learning. Nature 2015 518:7540 518:529–533. https://doi.org/10.1038/nature14236
2. Argall BD, Chernova S, Veloso M, Browning B (2009) A survey of robot learning from demonstration. Rob Auton Syst 57:469–483. https://doi.org/10.1016/J.ROBOT.2008.10.024
3. Ross S, Gordon GJ, Bagnell JA (2010) A Reduction of Imitation Learning and Structured Prediction to No-Regret Online Learning. Journal of Machine Learning Research 15:627–635
4. Fang B, Jia S, Guo D, et al (2019) Survey of imitation learning for robotic manipulation. International Journal of Intelligent Robotics and Applications 2019 3:4 3:362–369. https://doi.org/10.1007/s41315-019-00103-5
5. Chi C, Xu Z, Feng S, et al (2023) Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. International Journal of Robotics Research 44:1684–1704. https://doi.org/10.1177/02783649241273668
6. Kang JH, Joshi S, Huang R, Gupta SK (2025) Robotic Compliant Object Prying Using Diffusion Policy Guided by Vision and Force Observations arXiv preprint arXiv:2503.03998. https://doi.org/10.48550/arXiv.2503.03998
7. Wolf R, Shi Y, Liu S, Rayyes R (2025) Diffusion models for robotic manipulation: a survey. Front Robot AI 12:1606247. https://doi.org/10.3389/frobt.2025.1606247
8. Driess D, Xia F, Sajjadi MSM, et al (2023) PaLM-E: An Embodied Multimodal Language Model. arXiv preprint arXiv:2303.03378
9. Brohan A, Brown N, Carbajal J, et al (2023) RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control. arXiv preprint arXiv:2307.15818
10. Liang J, Huang W, Xia F, et al (2022) Code as Policies: Language Model Programs for Embodied Control. Proc IEEE Int Conf Robot Autom 2023-May:9493–9500. https://doi.org/10.1109/ICRA48891.2023.10160591
11. Ahn M, Brohan A, Brown N, et al (2022) Do As I Can, Not As I Say: Grounding Language in Robotic Affordances. Proc Mach Learn Res 205:287–318
12. Xiao X, Liu J, Wang Z, et al (2023) Robot Learning in the Era of Foundation Models: A Survey. arXiv preprint arXiv:2311.14379
13. Tao F, Zhang M (2017) Digital Twin Shop-Floor: A New Shop-Floor Paradigm Towards Smart Manufacturing. IEEE Access 5:20418–20427. https://doi.org/10.1109/ACCESS.2017.2756069
14. Rehman S, Al-Greer M, Burn AS, et al (2025) High-Volume Battery Recycling: Technical Review of Challenges and Future Directions. Batteries 2025, Vol 11, Page 94 11:94. https://doi.org/10.3390/BATTERIES11030094
15. ISO 10218-2:2025 - Robotics — Safety requirements — Part 2: Industrial robot applications and robot cells. https://www.iso.org/standard/73934.html. Accessed 13 Jan 2026
16. Ames AD, Xu X, Grizzle JW, Tabuada P (2016) Control Barrier Function Based Quadratic Programs for Safety Critical Systems. IEEE Trans Automat Contr 62:3861–3876. https://doi.org/10.1109/TAC.2016.2638961

17. Wabersich KP, Zeilinger MN (2021) A predictive safety filter for learning-based control of constrained nonlinear dynamical systems. Automatica 129:109597. https://doi.org/10.1016/J.AUTOMATICA.2021.109597
18. Fuller A, Fan Z, Day C, Barlow C (2020) Digital Twin: Enabling Technologies, Challenges and Open Research. IEEE Access 8:108952–108971. https://doi.org/10.1109/ACCESS.2020.2998358
19. Lee J, Bagheri B, Kao HA (2015) A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. Manuf Lett 3:18–23. https://doi.org/10.1016/j.mfglet.2014.12.001
20. Bommasani R, Hudson DA, Adeli E, et al (2022) On the Opportunities and Risks of Foundation Models
21. Ablett T, Limoyo O, Sigal A, et al (2025) Multimodal and Force-Matched Imitation Learning with a See-Through Visuotactile Sensor. IEEE Transactions on Robotics 41:946–959. https://doi.org/10.1109/TRO.2024.3521864
22. Brunke L, Greeff M, Hall AW, et al (2021) Safe Learning in Robotics: From Learning-Based Control to Safe Reinforcement Learning
23. Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D (2016) Concrete Problems in AI Safety. arXiv preprint arXiv:1606.06565. https://doi.org/10.48550/arXiv.1606.06565
24. Yan Z, Jouandeau N, Cherif AA (2013) A survey and analysis of multi-robot coordination. Int J Adv Robot Syst 10:. https://doi.org/10.5772/57313
25. Gielis J, Shankar A, Prorok A (2022) A Critical Review of Communications in Multi-robot Systems. Current Robotics Reports 3:213–225. https://doi.org/10.1007/S43154-022-00090-9
26. Lerman K, Jones C, Galstyan A, Mataríc MJ (2006) Analysis of Dynamic Task Allocation in Multi-Robot Systems. International Journal of Robotics Research 25:225–241. https://doi.org/10.1177/0278364906063426
27. Erdogan C, Contreras CA, Stolkin R, Rastegarpanah A (2024) Multi-Robot Task Planning for Efficient Battery Disassembly in Electric Vehicles. Robotics 2024, Vol 13, Page 75 13:75. https://doi.org/10.3390/ROBOTICS13050075