Journal of
Manufacturing and Materials Processing

MDPI

Review

Learning from Demonstration for Robotic Deburring and Polishing: A Systematic Mapping Study

Ercan Düzgün

Department of Mechanical Engineering, Faculty of Engineering, Bursa Uludağ University, 16059 Bursa, Türkiye; eduzgun@uludag.edu.tr

Abstract

Contact-rich manufacturing processes, such as surface cleaning, deburring, and polishing, require precise force regulation and complex trajectory tracking that are challenging to automate using conventional robot programming methods. Learning from Demonstration (LfD) offers a powerful alternative to transfer these expert skills from human operators to robotic systems. The objective of this study is to systematically map academic publications addressing LfD applications in robotic deburring and polishing between 2016 and 2026, classify the algorithmic structures, sensory modalities, and control configurations employed, and identify key industrial integration challenges. In accordance with the PRISMA 2020 guidelines, a systematic search was conducted across Scopus, Web of Science, IEEE Xplore, and Google Scholar databases. Out of the 288 initially retrieved records, duplicate removal and a two-stage screening process (Title/Abstract review, followed by full-text review) resulted in a final corpus of 24 primary studies included for qualitative synthesis. The included studies were classified into five algorithmic clusters: Dynamic Movement Primitives (DMPs) and variants (9 out of 24 studies, 38%), probabilistic and statistical models (8 out of 24 studies, 33%), deep learning and generative AI architectures (4 out of 24 studies, 17%), autonomous dynamical systems (2 out of 24 studies, 8%), and direct impedance control (1 out of 24 studies, 4%). Force/torque sensing remains the dominant modality; it was utilized exclusively in 71%—17 out of 24—of studies and in 87.5% of studies as any configuration (either as a sole modality or in multimodal setups). However, recent years have documented a trend toward multimodal perception and generative action policies (e.g., Diffusion Policies). The findings suggest that while LfD offers potential cost-reduction and flexibility benefits for small- and medium-sized enterprises (SMEs), technical barriers, such as the sim-to-real transfer gap, high-frequency impact dynamics in deburring, and the autonomous identification of local non-polishing areas (LNP areas), continue to limit widespread industrial deployment.

Keywords: robotic deburring; robotic polishing; robotic surface finishing; Learning from Demonstration; imitation learning; force control; systematic mapping study; PRISMA 2020

Check for updates

Received: 6 July 2026
Revised: 7 August 2026
Accepted: 9 August 2026
Published: 12 August 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Deburring and polishing represent critical, high-precision surface-finishing stages in the manufacturing of components across key industries, including aerospace, automotive, medical devices, and consumer electronics. These operations are characterized by complex physical interactions, demanding the simultaneous satisfaction of competing constraints: maintaining a controlled, continuous contact force normal to the workpiece surface, executing smooth spatial trajectories over freeform geometries, and dynamically adapting to

material variations, tool wear, and evolving surface topologies. In traditional industrial practice, these requirements are met through the tacit expertise of highly skilled human craftsmen. These operators possess an intuitive understanding of optimal contact angles, material-specific feed rates, and corrective force responses—aspects that are highly challenging to formalize mathematically or encode using conventional control models. However, rising labor costs and a declining availability of skilled manual labor in industrialized economies have driven a pressing demand to automate these processes by transferring expert human skills to robotic systems in a reliable and flexible manner.

The automation of surface finishing using conventional robotic programming paradigms—such as CAD/CAM-based offline trajectory generation, teach-pendant programming, and model-based force control—has achieved only limited success in high-mix, low-volume production. These approaches exhibit several fundamental limitations. First, they require high-fidelity three-dimensional CAD models of the workpiece, which are frequently unavailable, inaccurate, or geometrically altered due to manufacturing tolerances in cast, forged, or individually worn components. Second, generating force-compliant trajectories for freeform surfaces requires highly specialized robotic engineering expertise, making the programming process prohibitively expensive and time-consuming for Small- and Medium-sized Enterprises (SMEs). Third, conventional setups display poor adaptability; any change in part geometry or material properties requires the robotic path and control parameters to be extensively reprogrammed, causing substantial changeover downtime. These challenges are particularly pronounced in deburring, where burr size, location, and hardness are inherently stochastic, and in polishing, where uniform surface roughness demands real-time compliance far exceeding the capabilities of open-loop position controllers.

To overcome these limitations, the Learning-from-Demonstration (LfD) paradigm—also known as Programming by Demonstration (PbD) or imitation learning—has emerged as a compelling methodology. LfD enables robots to acquire complex manipulation skills directly from expert demonstrations rather than through explicit mathematical programming. In an LfD framework, a human operator guides the robot kinesthetically via haptic teleoperation interfaces or using instrumented tools, during which the system records motion trajectories, contact forces, and multimodal sensory data. Machine learning algorithms are then deployed to extract generalizable motion primitives or control policies from the demonstration data, allowing the robot to reproduce the task autonomously under varying environmental conditions. Over the past decade, LfD research has generated a rich family of algorithmic representations for surface finishing. These range from mathematically elegant Dynamic Movement Primitives (DMPs) and probabilistic representations—such as Gaussian Mixture Models (GMMs) and Probabilistic Movement Primitives (ProMPs)—to recent deep generative architectures, including Diffusion Policies and Neural Ordinary Differential Equations (Neural ODEs). The maturation of collaborative robots (cobots) equipped with joint torque sensors, such as the Franka Emika Panda, has further accelerated the practical feasibility of LfD-based finishing on factory floors.

Although the broader domain of robotic Learning from Demonstration has been surveyed in several comprehensive reviews, these works address LfD in general manipulation contexts and do not provide a focused, systematic analysis of its application to contact-rich surface finishing. The domains of robotic deburring and polishing present unique physical and mechatronic challenges that are absent from general LfD tasks: the necessity of coupling spatial trajectory imitation with contact force profile learning, the management of tool–workpiece interaction dynamics under freeform shapes, the integration of multimodal sensory feedback, and the strict constraints imposed by industrial deployability. To the best of the authors' knowledge, no systematic mapping study has yet been conducted

that exclusively addresses LfD-based methodologies in robotic deburring and polishing, synthesizes the algorithmic and sensory configurations employed, characterizes the performance evaluation metrics used across studies, and identifies the open research challenges that hinder transition from laboratory settings to industrial practice. This absence of a dedicated structured synthesis represents a significant gap in the literature, limiting the ability of both researchers and practitioners to identify the current State of the Art, understand methodological trade-offs, and prioritize future research directions in this rapidly evolving field.

To address this gap, the present work conducts a systematic mapping study of the literature on LfD-based robotic deburring and polishing, covering peer-reviewed publications from 2016 to 2026. A systematic mapping study is a form of secondary research that aims to provide a broad overview of a research area through the classification and thematic aggregation of primary studies, and is particularly well-suited to emerging fields where the body of evidence is growing but remains insufficiently synthesized. The study protocol was pre-registered on the Open Science Framework (OSF) prior to the literature search to ensure methodological transparency and minimize reporting bias (https://doi.org/10.17605/OSF.IO/5Y6BR).

The principal contributions of this work are as follows:

1. Structured Taxonomy: A structured taxonomy of LfD architectures applied in robotic deburring and polishing, categorizing 24 primary studies into five methodological clusters: DMP-based methods, probabilistic and statistical models, deep learning and generative AI architectures, autonomous dynamical systems, and direct adaptive control approaches.
2. Mechatronic Analysis: A systematic analysis of sensory and control configurations, documenting how robotic platforms, force/torque sensors, vision systems, and haptic interfaces are integrated to support LfD-based surface finishing.
3. Evaluation Metrics Synthesis: A comprehensive characterization of evaluation metrics employed in the field, spanning kinematic trajectory accuracy, dynamic force-tracking performance, and physical surface quality measures.
4. Industrial Challenges Roadmap: An evidence-based synthesis of industrial deployment challenges, identifying the key technical barriers that currently impede the adoption of LfD-based surface finishing in production environments.

To guide this systematic mapping study, the following four Research Questions (RQs) were defined:

- RQ1—Methodological Landscape: Which Learning-from-Demonstration (LfD) architectures, trajectory learning algorithms, and mathematical representations are most commonly used in robotic deburring and polishing tasks?
- RQ2—System and Sensory Configuration: How are robotic systems, control architectures, and sensory modalities (e.g., force sensing, vision, multimodal perception) designed and integrated to support LfD-based surface-finishing operations?
- RQ3—Evaluation and Performance Metrics: What performance evaluation criteria and validation metrics are used to assess the effectiveness of LfD-based robotic deburring and polishing methods?
- RQ4—Industrial Deployment Challenges: What are the key technical challenges, limitations, and research gaps that hinder the transition of LfD-based robotic surface-finishing methods from laboratory settings to industrial applications?

The remainder of this paper is organized as follows. Section 2 describes the methodology of the systematic mapping study, including the protocol and registration, eligibility criteria, search strategy, quality assessment framework, and screening and selection process

in accordance with PRISMA 2020 guidelines [1]. Section 3 presents the results of the mapping, beginning with the PRISMA flow diagram, study characteristics, bibliometric analysis, and quality assessment (Sections 3.1–3.4), followed by responses to the Research Questions: Methodological Landscape (RQ1, Section 3.5); System and Sensory Configuration (RQ2, Section 3.6); Evaluation Metrics (RQ3, Section 3.7); Industrial Deployment Challenges (RQ4, Section 3.8). Section 4 provides an integrated discussion of the findings. Section 5 acknowledges the limitations of this study, Section 6 outlines priority directions for future research, and Section 7 concludes the paper.

2. Methodology

2.1. Protocol and Registration

This systematic mapping study was conducted in strict accordance with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement, with the completed checklist provided as Supplementary Materials. To ensure methodological transparency, minimize reporting bias, and guarantee reproducibility, a detailed research protocol was drafted and registered prior to the commencement of the literature search and screening process. The protocol is publicly accessible on the Open Science Framework (OSF) registry under the following persistent identifier: https://doi.org/10.17605/OSF.IO/5Y6BR. No amendments or deviations were made to the registered protocol during the conduct of this study.

2.2. Eligibility Criteria

To identify relevant studies that address the defined Research Questions (RQ1–RQ4), a set of formal inclusion and exclusion criteria was established. The scope of this mapping study was restricted to primary studies focusing on the intersection of Learning from Demonstration and robotic surface-finishing processes. To ensure mechatronic and methodological consistency, we define 'robotic surface finishing' as any contact-rich robotic operation designed to modify or condition a workpiece surface, including polishing and deburring.

Studies were deemed eligible for inclusion if they met all of the following conditions:

- Thematic Relevance: The study must explicitly investigate robotic surface-finishing operations, specifically deburring or polishing.
- Methodological Approach: The proposed robotic control or programming framework must incorporate an LfD-based method (e.g., kinesthetic teaching, teleoperated demonstration, or haptic imitation learning) where motion or force policies are extracted from human demonstrations.
- Technical Depth: The study must provide experimental validation or detailed theoretical formulations regarding trajectory generation, force regulation, control architectures, or sensor configurations.
- Publication Type and Language: Only peer-reviewed academic journal articles or conference proceedings published in the English language were included.

Studies were excluded if they exhibited any of the following characteristics:

- Out-of-Scope Operations: Studies addressing generic robotic manipulation tasks (e.g., pick-and-place, assembly, peg-in-hole insertion, or trajectory tracking in free space) without a specific surface-finishing context.
- Non-Demonstration Control: Research utilizing traditional robotic programming (e.g., offline CAD/CAM programming, manual teach-pendant jogging) or model-based adaptive controllers that do not learn or adapt based on human demonstration data.
- Insufficient Quality or Length: Extended Abstracts, editorial prefaces, technical reports, white papers, book reviews, or unpublished Master's/Doctoral theses.

- Accessibility Barriers: Articles for which the full-text version was unavailable or lacked sufficient data to extract parameters necessary for answering the Research Questions.

A summary of the eligibility criteria is presented in Table 1.

Table 1. Summary of the Eligibility Criteria.

Criteria Category	Inclusion Criteria	Exclusion Criteria
Application Domain	Robotic deburring and polishing operations	Generic assembly, pick-and-place, free-space trajectory tracking, or heavy rough grinding/milling (bulk shaping)
Methodology	Learning from Demonstration (LfD, PbD, imitation learning, DMP, GMM, etc.)	Standard CNC programming, offline CAD/CAM paths, non-learning adaptive control
Evidence Type	Experimental validation or detailed simulation of finishing tasks	Abstract concepts without task validation, technical sheets, patents
Publication Medium	Peer-reviewed journal articles and conference papers	Abstract concepts without task validation, technical sheets, patents
Language	English language only	Abstract concepts without task validation, technical sheets, patents

2.3. Search Strategy and Information Sources

A comprehensive, systematic search of the literature was executed to retrieve all peer-reviewed studies published between 1 January 2016 and 24 June 2026. This decade-long timeframe was selected to capture the most recent and significant advancements in LfD architectures, collaborative robotic hardware, and deep learning algorithms.

The search was conducted across four major electronic databases:

1. IEEE Xplore: Targeted for its high concentration of robotics, automation, and control engineering publications.
2. Scopus: Utilized for broad, multidisciplinary coverage of the engineering and physical sciences literature.
3. Web of Science (Core Collection): Employed to capture high-impact journals and citation indexes.
4. Google Scholar: Used as a supplementary resource to capture early-access articles, conference papers, and preprints, thereby minimizing potential publication bias.

To ensure a high recall rate, the search queries were constructed by combining keywords from two main thematic blocks using Boolean operators:

- Block A (Learning Paradigm): "Learning from Demonstration", "LfD", "Imitation Learning", "Programming by Demonstration", "DMP", "Dynamic Movement Primitives", "ProMP", "Probabilistic Movement Primitives", "Skill learning", "Task learning".
- Block B (Application Context): "Deburring", "Polishing", "Surface finishing", "Surface-finishing", "Finishing", "Robotic finish*".

The search strings were tailored to the syntax requirements of each database. For the IEEE Xplore query, the date restriction was applied manually via the search interface filters (2016–2026) because the query editor does not support date syntax directly inside the search string. The queries were executed on 24 June 2026. The specific search strings applied to each database are summarized in Table 2.

Table 2. Search Queries Applied to the Electronic Databases.

Database	Search Query
Scopus	TITLE-ABS-KEY ((“Learning from Demonstration” OR “LfD” OR “Imitation Learning” OR “Programming by Demonstration” OR “DMP” OR “Dynamic Movement Primitives” OR “ProMP” OR “Probabilistic Movement Primitives” OR “Skill learning” OR “Task learning”) AND (“Deburring” OR “Polishing” OR “Surface finishing” OR “Surface-finishing” OR “Finishing” OR “Robotic finish*”)) AND PUBYEAR > 2015 AND PUBYEAR < 2027 AND (LIMIT-TO(DOCTYPE, “ar”) OR LIMIT-TO (DOCTYPE, “cp”)) AND (LIMIT-TO (LANGUAGE,“English”))
Web of Science	TS=((“Learning from Demonstration” OR “LfD” OR “Imitation Learning” OR “Programming by Demonstration” OR “DMP” OR “Dynamic Movement Primitives” OR “ProMP” OR “Probabilistic Movement Primitives” OR “Skill learning” OR “Task learning”) AND (“Deburring” OR “Polishing” OR “Surface finishing” OR “Surface-finishing” OR “Finishing” OR “Robotic finish*”)) AND PY=(2016-2026) AND DT=(ARTICLE OR PROCEEDINGS PAPER) AND LA=(ENGLISH)
IEEE Xplore	(“Learning from Demonstration” OR “LfD” OR “Imitation Learning” OR “Programming by Demonstration” OR “DMP” OR “Dynamic Movement Primitives” OR “ProMP” OR “Probabilistic Movement Primitives” OR “Skill learning” OR “Task learning”) AND (“Deburring” OR “Polishing” OR “Surface finishing” OR “Surface-finishing” OR “Finishing” OR “Robotic finish*”)
Google Scholar	(“Learning from Demonstration” OR “Imitation Learning” OR “Programming by Demonstration” OR “Dynamic Movement Primitives” OR “Probabilistic Movement Primitives”) AND (Deburring OR Polishing OR “Surface finishing” OR “Surface-finishing” OR Finishing OR “Robotic finish*”)

2.4. Screening and Selection Process

Upon completing the database searches, all identified records were imported into reference management software, and duplicate entries were systematically removed. The study selection process strictly followed the PRISMA 2020 guidelines [1], executing three sequential screening phases:

1. Identification: The initial database searches yielded a total of 288 records (Scopus: 57; Web of Science: 43; IEEE Xplore: 29; Google Scholar: 159). Following the elimination of 65 duplicate entries, 223 unique records were compiled for screening.
2. Screening (Title and Abstract Review): A Title-and-Abstract review was conducted on the 223 unique records against the predefined inclusion and exclusion criteria. This phase resulted in the exclusion of 197 records that lacked relevance to LfD-based robotic deburring or polishing, leaving 26 records for full-text eligibility assessment.
3. Eligibility (Full-Text Review): The full texts of the 26 remaining studies were retrieved and analyzed. During this phase, 2 studies were excluded due to out-of-scope applications: one study focused exclusively on grinding-based material removal [2], and another study addressed sanding without a polishing or deburring component [3].

Following this rigorous process, exactly 24 primary studies were deemed eligible and included in the final corpus for systematic mapping and thematic synthesis. The screening and selection process was conducted by a single primary researcher, with borderline or ambiguous cases reviewed multiple times to verify strict adherence to the thematic scope. To address potential selection bias and align with standard methodological recommendations, a retrospective inter-rater reliability check was subsequently performed by a second independent researcher. A random subset (20% of the initial records; 45 out of 223 articles) was selected for Title-and-Abstract screening, and 100% of the full-text eligible articles (26 articles) were evaluated independently. The inter-rater agreement was high: Cohen’s

Kappa κ was 0.88 for the Title/Abstract screening stage (95.6% agreement) and 1.00 for the full-text eligibility stage (100% agreement). Any minor differences in classification during the Title/Abstract stage were resolved through consensus discussion. A standardized data extraction form was utilized to ensure reliability.

2.5. Methodological Quality Assessment

To evaluate the methodological rigor of the included primary studies, a domain-specific quality assessment (QA) framework was developed. Standard clinical risk-of-bias tools are not directly applicable to experimental robotics research. Therefore, a rubric was constructed based on established methodological quality criteria for software and systems engineering systematic reviews and mapping studies and grounded in experimental design principles for software and engineering systems. This rubric was adapted to the specific experimental and mechatronic characteristics of Learning-from-Demonstration (LfD)-based robotic surface finishing.

The framework comprises six key criteria (QA1–QA6), each evaluated on a three-point scale: 2 (fully met), 1 (partially met), and 0 (not met or not reported):

- QA1—Experimental Validation: Does the study validate the proposed LfD framework on a physical robotic manipulator performing a surface-finishing task?
- QA2—Reproducibility and Statistical Reporting: Are the experiments repeated over multiple trials? Are quantitative statistical measures (e.g., mean and standard deviation) reported?
- QA3—Baseline Comparison: Is the proposed method compared against a baseline (e.g., standard model-based control, human expert performance, or alternative LfD algorithms)?
- QA4—Generalizability: Is the learned skill tested on varying workpiece geometries, orientations, or materials?
- QA5—Quantitative Performance Metrics: Does the study report quantitative evaluation metrics for both kinematic trajectory accuracy and dynamic force tracking?
- QA6—Industrial Relevance: Is the task validated using an industrial workpiece? Are real-world industrial deployment constraints (e.g., tool wear, cycle time, worker safety) discussed?

The full results of the methodological quality assessment are reported in Section 3.4. To ensure objective evaluation and prevent assessor bias, the scoring criteria are operationalized as follows:

- QA1—Experimental Validation: Does the study validate the proposed LfD framework on a physical robotic manipulator performing a surface-finishing task?
    - Score 2 (Fully Met): The proposed LfD algorithm is validated on physical robotic hardware (i.e., a physical manipulator) performing a real finishing operation (e.g., polishing, deburring) to account for real-world high-frequency contact forces. The majority of the included studies (22 out of 24, 92%; e.g., [4,5]) fully met this criterion.
    - Score 1 (Partially Met): The framework is evaluated only in high-fidelity robotic simulators (e.g., CoppeliaSim, Gazebo, MuJoCo, PyBullet) with representative physical models, but it lacks hardware execution.
    - Score 0 (Not Met): The framework is evaluated only in basic mathematical simulation (e.g., MATLAB plotting paths without robot dynamics/collision models) or not validated at all.

- QA2—Reproducibility and Statistical Reporting: Are the experiments repeated over multiple trials? Are quantitative statistical measures (e.g., mean and standard deviation) reported?
  - Score 2 (Fully Met): Multiple experimental trials (typically $N \geq 5$) are conducted, and quantitative statistical descriptors (e.g., mean and standard deviation, confidence intervals, box plots) are reported for the performance metrics (e.g., [5–7]).
  - Score 1 (Partially Met): Experimental repetition is mentioned, or multiple trials are plotted, but without formal statistical aggregation (e.g., only plotting individual force/trajectory profiles or reporting a single "representative" or "best" run; e.g., [4,8]).
  - Score 0 (Not Met): The study reports only a single test trial or provides no information regarding repetition and statistical metrics.
- QA3—Baseline Comparison: Is the proposed method compared against a baseline (e.g., standard model-based control, human expert performance, or alternative LfD algorithms)?
  - Score 2 (Fully Met): The proposed method is quantitatively compared against one or more external baselines (e.g., standard admittance control, classical DMP, pure Reinforcement Learning, or human expert performance) under identical task conditions (e.g., [5,6,9,10]).
  - Score 1 (Partially Met): The proposed method is compared to a baseline only qualitatively (e.g., visual comparison of force plots) or is compared only against simplified or un-optimized versions of itself (ablation study only; e.g., [4,8,9]).
  - Score 0 (Not Met): The proposed method's performance is reported in isolation with no baseline comparison or ablation study.
- QA4—Generalizability: Is the learned skill tested on varying workpiece geometries, orientations, or materials?
  - Score 2 (Fully Met): The learned skill is evaluated on workpieces with different physical geometries (e.g., transitioning from flat to curved surfaces), different spatial orientations (e.g., rotated workpiece setups), or different material properties without manual retraining (e.g., [11–13]).
  - Score 1 (Partially Met): Generalizability is evaluated under limited perturbations, such as small changes in starting position/orientation or velocity scaling, but on the exact same workpiece geometry and material (e.g., [4,5]).
  - Score 0 (Not Met): The learned skill is evaluated only on the exact same workpiece geometry and trajectory configuration used in the demonstration phase, without any perturbation.
- QA5—Quantitative Performance Metrics: Does the study report quantitative evaluation metrics for both kinematic trajectory accuracy and dynamic force tracking?
  - Score 2 (Fully Met): The study reports quantitative metrics (e.g., RMSE, MAE, correlation coefficients) for both kinematic tracking (position/velocity) and dynamic interaction forces (e.g., [4,5,14]).
  - Score 1 (Partially Met): The study reports quantitative metrics for only one domain (e.g., reporting trajectory RMSE but presenting force profiles only qualitatively via plots, or vice versa; e.g., [6,8,15]).
  - Score 0 (Not Met): The evaluation is purely qualitative (e.g., visual surface finish inspection) with no quantitative metrics reported for either force or trajectory.
- QA6—Industrial Relevance: Is the task validated using an industrial workpiece? Are real-world industrial deployment constraints (e.g., tool wear, cycle time, worker safety) discussed?

- Score 2 (Fully Met): The validation is performed on a real industrial workpiece (e.g., turbine blades, engine components, weld beads) and/or explicitly incorporates practical industrial constraints, such as tool wear, processing cycle times, or system safety parameters (e.g., [4,12,16,17]).
- Score 1 (Partially Met): Validation is performed on simple mock-up workpieces (e.g., flat plates, acrylic blocks), followed by a structured, dedicated discussion of industrial deployment issues, system calibration, or scalability constraints (e.g., [5,9]).
- Score 0 (Not Met): Validation is restricted to simple mock-up workpieces with no discussion of industrial deployment, cycle times, tool wear, or real-world deployment challenges.

Based on the cumulative score (ranging from 0 to 12), studies were classified into three quality tiers:

- High Quality (H): Cumulative score of 9 to 12.
- Moderate Quality (M): Cumulative score of 5 to 8.
- Low Quality (L): Cumulative score of 0 to 4.

The quality assessment rubric was customized for robotic LfD. Since physical validation is a core eligibility inclusion requirement to address real-world surface-finishing force dynamics, all 24 included studies met the QA1 criterion (Score 2) for physical hardware, while simulation-only works (which would receive Score 1) were excluded during the selection phase.

2.6. Data Extraction and Synthesis

A standardized data extraction template was designed to systematically capture the technical characteristics of each included study. The template captured data across five primary dimensions:

1. Bibliographic Metadata: Authors, publication year, publication venue, and the geographic location of the corresponding author.
2. LfD Algorithmic Architecture: Trajectory representation model, learning algorithms, and mathematical formulations (e.g., DMPs, GMM, GMR, Diffusion Policy, Neural ODEs).
3. Sensory Modalities: Sensors integrated during the human demonstration phase and the robot execution phase (e.g., 6-DOF force/torque sensors, joint torque sensors, RGB/RGB-D cameras, haptic interfaces).
4. Hardware and Control Platform: The specific robotic manipulator model, degrees of freedom (DOF), and low-level force/position control strategies (e.g., impedance control, admittance control, variable stiffness control).
5. Application Environment: The specific surface-finishing task (deburring, polishing, cleaning, rust removal) and the shape/material of the test workpiece.

Data extraction was completed by the primary researcher in a continuous session to ensure coding consistency. For studies where specific hardware parameters were not reported or were irrelevant (e.g., studies validating skills using custom instrumented tools unattached to a robotic arm), the corresponding fields were coded with a '-' symbol to maintain data integrity. The extracted data were then qualitatively synthesized to answer the Research Questions, and quantitative summaries (frequencies and percentages) were calculated for the algorithmic and sensory distributions. For the purpose of this systematic mapping study, the primary outcomes of interest are defined as the qualitative LfD algorithmic architectures, mechatronic sensory configurations, low-level control systems,

and the corresponding quantitative evaluation metrics (e.g., force and trajectory RMSEs, surface roughness Ra) reported across the corpus.

2.7. Effect Measures and Synthesis Rationale

Due to the qualitative and taxonomic nature of this systematic mapping study, standard statistical effect measures (such as risk ratios, relative risks, or mean differences) commonly employed in clinical systematic reviews are not applicable. Instead, the synthesis is structured around descriptive and frequency-based thematic aggregation of algorithmic, mechatronic, and control features. Quantitative baseline performance metrics (e.g., trajectory-tracking root-mean-square errors and surface roughness measures) are summarized qualitatively to establish baseline benchmarks without performing statistical meta-analyses.

3. Results

3.1. Study Selection and PRISMA Flow Diagram

The initial systematic search across the four electronic databases yielded 288 records. After importing these records into reference management software, 65 duplicate entries were identified and removed, leaving 223 unique records for the screening phase.

During the Title-and-Abstract screening, 197 records were excluded as they did not meet the predefined eligibility criteria (primarily due to lack of LfD methodology or relevance to robotic surface finishing). The remaining 26 studies were retrieved for full-text eligibility assessment.

Following a detailed full-text review, 2 studies were excluded: one focused exclusively on grinding-based material removal [2], and the other addressed sanding without a polishing or deburring component [3]. This process resulted in a final corpus of 24 primary studies included in the systematic mapping and thematic synthesis. The study selection process is illustrated in Figure 1.

To ensure transparency and reproducibility regarding our data collection and screening process, the platform settings and procedures are detailed below:

- Search Query and Filters: The specific search keywords used are presented in Table 2. The publication date filter was manually set to the 2016–2026 range, and the document type filter was manually restricted to "Articles only".
- Retrieval Limits and Pages: This search yielded a total of 159 results, which were distributed across nine pages on Google Scholar. Since the total number of results was well below Google Scholar's standard pagination limit (approx. 1000 results), all nine pages of results were fully retrieved.
- Export Procedure: Due to Google Scholar's lack of native batch-export features for large datasets, the search results from these nine pages were manually exported into Excel.
- Duplicate Detection and Reference Management: Records retrieved from Scopus, Web of Science, and IEEE Xplore were first automatically imported into Mendeley. Duplicate detection and removal were performed automatically within Mendeley, and the cleaned list was subsequently exported to Excel. Finally, manual cross-checking and duplicate elimination between the database records and the Google Scholar results were performed within Excel.

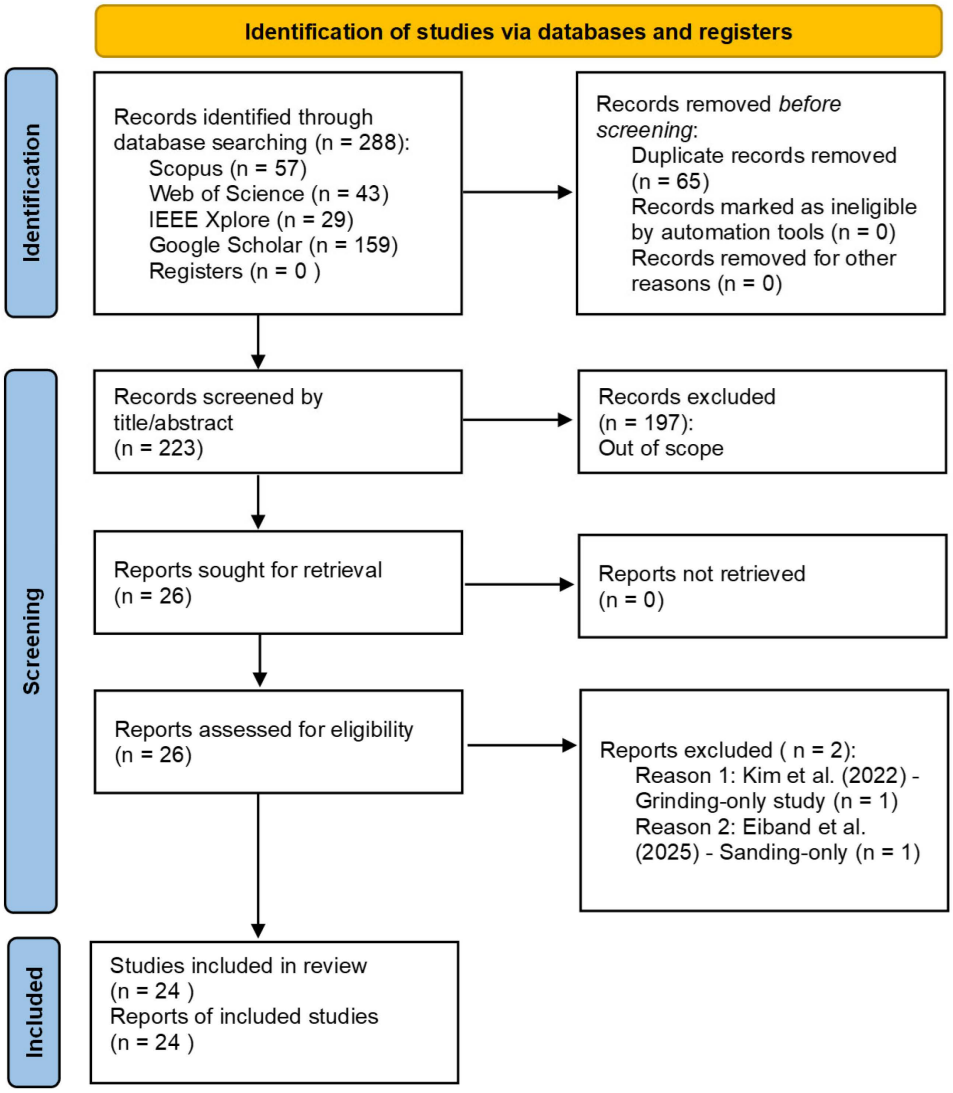

Figure 1. PRISMA 2020 Flow Diagram of the Study Selection Process [2,3].

3.2. Quantitative Study Characteristics

A summary of the extracted data from the 24 included primary studies is presented in Table 3. The corpus represents a decade of research (2016–2026) and documents a diverse range of LfD algorithms, sensory modalities, robotic systems, and surface-finishing applications.

Table 3. Data Extraction and Technical Characteristics of the Included Studies (n = 24).

No.	Authors	Year	Robot Type	Application Area	Algorithm Used
1	Min et al. 2024 [4]	2024	Franka Emika Panda	Polishing	Computed-torque impedance control
2	Li et al. 2025 [5]	2025	xArm 7	Polishing	Diffusion Policy + Residual RL
3	Fischer et al. 2025 [6]	2025	KUKA LBR iiwa 14	Polishing	ProMPs (Few-Shot)
4	Xu et al. 2025 [7]	2025	Universal Robots UR5	Polishing	Neural ODEs (Hyper-NODEs)
5	Si et al. 2024 [8]	2024	Franka Emika Panda	Polishing	DS-based imitation learning

Table 3. *Cont.*

No.	Authors	Year	Robot Type	Application Area	Algorithm Used
6	Wang et al. 2023a [9]	2023a	Universal Robots UR16e	Polishing	PMDRNN + DMPs
7	Ke et al. 2025 [10]	2025	KUKA LBR iiwa 14	Polishing	Vision-Diffusion+ Mesh-DMP
8	Wu et al. 2023 [11]	2023	Franka Emika Panda	Polishing	GMM-GMR, Var. Impedance
9	Möhl et al. 2025 [12]	2025	None (Point cloud morphing)	Polishing	Neural network morphing
10	Wang et al. 2025 [13]	2025	Universal Robots UR16e	Polishing	PI2-BDMPs
11	Wang et al. 2023b [14]	2023b	Universal Robots UR16e	Polishing	AL_ProMP
12	Kulak et al. 2020 [15]	2020	Franka Emika Panda	Polishing	Fourier Movement Primitives (FMPs)
13	Wu et al. 2025 [16]	2025	Franka Emika Panda	Polishing	PC-GMM-DS, Var. Impedance
14	Parvizi et al. 2017 [17]	2017	Phantom haptic device	Deburring	Modified DMPs (sDMP)
6	Acikgoz et al. 2017 [18]	2017	PI 6-DOF Hexapod	Deburring	DMPs
7	Zhai et al. 2022 [19]	2022	Franka Emika Panda	Polishing	GMM-GMR, Var. Impedance
10	Zhang et al. 2024 [20]	2024	Rethink Robotics Sawyer	Polishing	DTW-ILC + GMM
13	Haninger et al. 2023 [21]	2023	Franka Emika Panda	Polishing	MPC + Gaussian Processes
18	Shen et al. 2024 [22]	2024	KUKA KR6 R900	Polishing	FDC-DMP
19	Hamdan et al. 2024 [23]	2024	Universal Robots UR5	Polishing	MLP-based force learning
20	Liao et al. 2024 [24]	2024	Franka Emika Panda	Polishing	Riemannian DMP + QP
21	Li et al. 2020 [25]	2020	Rethink Robotics Baxter	Polishing	DMPs + Vision
23	Nemec et al. 2018 [26]	2018	Yaskawa Motoman MH-6	Polishing	Virtual mechanism + ILC
24	Duarte et al. 2024 [27]	2024	KUKA LBR iiwa 7 and Kinova Gen3	Polishing	Dynamical system (Limit cycle)

While this temporal concentration highlights the rapid acceleration and the current State of the Art in learning-based paradigms for contact-rich manufacturing, it also introduces a potential recency bias, as many included papers have had limited time for citation impact, widespread replication, or validation by the broader scientific community.

3.3. Bibliometric Analysis

3.3.1. Temporal Distribution

The temporal distribution of the 24 selected studies between 2016 and 2026 reveals a significant growth in research activity. As summarized in Table 4, the annual publication volume remained low but stable from 2017 to 2022, with a sharp increase beginning in 2023. The years 2024 and 2025 represent the peak period of academic output, accounting for 14 out of 24 studies (58.3%) of the entire corpus. This trend reflects the growing interest in utilizing learning-based paradigms to automate contact-rich manufacturing tasks.

Table 4. Temporal Distribution of the Included Studies (2016–2026).

Publication Year	Frequency	Percentage (%)	Cum. Percentage (%)
2017	2	8.3%	8.3%
2018	1	4.2%	12.5%
2019	0	0%	12.5%
2020	2	8.3%	20.8%
2021	0	0%	20.8%
2022	1	4.2%	25.0%
2023	4	16.7%	41.7%
2024	7	29.2%	70.8%
2025	7	29.2%	100%
Total	24	100%	100%

3.3.2. Geographical Distribution

The geographical analysis, determined by the affiliation of the corresponding author, indicates a highly centralized research landscape. As shown in Table 5, East Asian institutions—specifically from China—account for the vast majority of the research, contributing 15 out of 24 papers (62.5%). Europe forms the second-largest research cluster, led by Austria (2 out of 24, 8.3%) and supported by Germany, Slovenia, and Portugal (each contributing 1 paper, 4%). Türkiye contributes 3 out of 24 publications (13%), representing two distinct research groups: one group at Middle East Technical University (METU) contributes two papers with overlapping authors [17,18], and a separate collaborative group led by Koç University contributes one paper [23].

Table 5. Geographical Distribution of the Corresponding Author Affiliations.

Country/ Region	Frequency	Percentage (%)	Key Institutions
China	15	62.5%	Huazhong University of Science and Technology, Harbin Institute of Technology
Türkiye	3	12.5%	Middle East Technical University, Koç University
Austria	2	8.3%	PROFACTOR GmbH, Johannes Kepler University Linz
Germany	1	4.2%	Fraunhofer Institute for Manufacturing Engineering and Automation
Slovenia	1	4.2%	Jožef Stefan Institute
Switzerland	1	4.2%	Idiap Research Institute/EPFL
Portugal	1	4.2%	Instituto Superior Técnico, University of Lisbon
Total	24	100%	-

3.3.3. Quantitative Methodological and Sensory Distributions

To provide a quantitative overview of the technological configurations, the 24 included studies were classified based on their primary LfD algorithmic approach (RQ1) and sensory modalities (RQ2).

Dynamic Movement Primitives (DMPs) and their force-coupled extensions represent the most prevalent mathematical framework, utilized in 9 studies (37.5%). Next, probabilistic models (GMM-GMR, ProMPs, FMPs, GPs) follow closely with 8 studies (33.3%). Modern deep learning and generative AI architectures represent an emerging cluster with 4 studies (16.7%), as summarized in Table 6.

Table 6. Algorithmic Architecture Frequency Distribution (n = 24).

Algorithmic Cluster	Frequency	Percentage (%)	Representative Methods
Dynamic Movement Primitives (DMPs) and Variants	9	37.5%	FDC-DMP, B-Spline DMP, Riemannian DMP
Probabilistic and Statistical Models	8	33.3%	GMM-GMR, AL-ProMP, Fourier MP, Gaussian Processes
Deep Learning and Generative AI	4	16.7%	Diffusion Policy, Residual RL, Hyper-NODEs, MLPs
Autonomous Dynamical Systems (DSs)	2	8.3%	Stable Limit Cycles, DS-based Imitation
Direct Impedance Control and Parameter Estimation	1	4.2%	Computed-Torque Impedance Control
Total	24	100%	-

Regarding sensory configurations, force and torque sensing remains the dominant modality. As detailed in Table 7, 17 studies (70.8%) rely exclusively on force/torque feedback (either via external F/T sensors or joint torque measurements) to regulate tool–workpiece contact. Vision-only systems account for 3 studies (12.5%), while multimodal configurations combining force, vision, and haptic feedback represent 4 studies (16.7%).

Table 7. Sensory Modality Frequency Distribution (n = 24).

Sensory Modality	Frequency	Percentage (%)	Key Hardware Elements
Force/Torque Sensing Only	17	70.8%	6-DOF F/T Sensors, Joint Torque Sensors
Vision Only (RGB/RGB-D)	3	12.5%	Depth Cameras, Point Clouds
Multimodal (Force + Vision/Haptic)	4	16.7%	RGB-D + F/T Sensor, Haptic Interface + Force Feedback
Total	24	100%	-

3.4. Methodological Quality Assessment Results

The methodological quality of the 24 included studies was evaluated using the six-criteria rubric (QA1–QA6) defined in Section 2.5. The results are presented in Table 8.

Overall, the corpus demonstrates a high–moderate level of methodological rigor: 10 studies (41.7%) were rated High Quality (score 9–12), and 14 studies (58.3%) were rated Moderate Quality (score 5–8). No studies were rated Low Quality, reflecting the peer-reviewed selection process.

Regarding the quality assessment (QA) process, criterion QA1 yielded a uniform score of 2 across all included studies, resulting in limited discriminatory power. However, QA1 was deliberately retained within our evaluation framework not as an elimination or differentiation metric, but as a foundational methodological threshold. Specifically, this criterion ensures that every included study meets the minimum mandatory standards of peer-reviewed scientific rigor, clear objective definition, and methodological transparency required for systematic synthesis. While criteria with uniform scores do not contribute to the relative ranking or weighting of the studies, retaining QA1 serves as an essential quality-control gatekeeper, confirming that all analyzed studies fulfill baseline academic criteria before proceeding to the final PRISMA extraction phase. The most common limitation was QA2 (Reproducibility and Statistical Reporting), reflecting a general lack of statistical

replication (e.g., reporting mean ± SD over multiple trials) across the experimental setups, with only 2 studies (8.3%) fully meeting this criterion.

Table 8. Methodological Quality Assessment of the Included Studies (n = 24).

Source	QA1	QA2	QA3	QA4	QA5	QA6	Total	Rating
Min et al. [4]	2	1	1	1	2	2	9	High
Li et al. [5]	2	2	2	1	2	1	10	High
Fischer et al. [6]	2	2	2	2	1	1	10	High
Xu et al. [7]	2	2	2	1	2	1	10	High
Si et al. [8]	2	1	1	1	1	1	7	Moderate
Wang et al. [9]	2	1	1	1	2	1	8	Moderate
Ke et al. [10]	2	1	2	2	1	2	10	High
Wu et al. [11]	2	1	1	2	2	1	9	High
Möhl et al. [12]	2	1	2	2	1	2	10	High
Wang et al. [13]	2	1	2	2	2	1	10	High
Wang et al. [14]	2	1	1	1	2	1	8	Moderate
Kulak et al. [15]	2	1	1	1	1	1	7	Moderate
Wu et al. [16]	2	1	1	1	2	2	9	High
Parvizi et al. [17]	2	1	1	1	1	2	8	Moderate
Acikgoz et al. [18]	2	1	1	1	1	2	8	Moderate
Zhai et al. [19]	2	1	1	1	2	1	8	Moderate
Zhang et al. [20]	2	1	1	1	2	1	8	Moderate
Haninger et al. [21]	2	1	1	1	2	1	8	Moderate
Shen et al. [22]	2	1	1	2	2	2	10	High
Hamdan et al. [23]	2	1	1	1	2	1	8	Moderate
Liao et al. [24]	2	1	1	1	2	1	8	Moderate
Li et al. [25]	2	1	1	2	1	1	8	Moderate
Nemec et al. [26]	2	1	1	1	2	1	8	Moderate
Duarte et al. [27]	2	1	1	1	1	1	7	Moderate
Fully Met (score = 2)	24	2	6	6	14	6	-	-
% Fully Met	100%	8.3%	25.0%	25.0%	58.3%	25.0%	-	-

3.5. RQ1: Methodological Landscape

Based on the systematic mapping of the literature, LfD architectures in robotic deburring and polishing are categorized into five primary methodological clusters. These approaches range from mathematical trajectory parameterization to advanced deep generative artificial intelligence models:

3.5.1. Dynamic Movement Primitives (DMPs) and Variants

DMPs represent the most prevalent framework (37.5% of studies), utilizing a system of second-order differential equations (spring-damper systems) augmented with a nonlinear forcing term to encode and reproduce demonstrated trajectories. The transformation system is mathematically formulated as follows:

$$\tau\dot{v} = K_p(g - x) - D_p v + (g - x_0)f(s) \quad (1)$$

$$\tau\dot{x} = v \quad (2)$$

where x is the position, v is the velocity, g is the goal position, x_0 is the starting position, K_p and D_p are the stiffness and damping gains, and τ is a time scaling parameter. The forcing term $f(s)$ is parameterized using Gaussian basis functions, $\psi_i(s)$, as follows:

$$f(s) = \frac{\sum_i w_i \psi_i(s)}{\sum_i \psi_i(s)} s \quad (3)$$

where w_i are the weights learned from demonstration, and s is the phase variable governed by the canonical system $\tau\dot{s} = -\alpha_s s$ (with decay rate α_s). While standard DMPs guarantee spatial and temporal scaling, they are kinematically restricted and cannot adapt to control force variations. To address this, studies have introduced several domain-specific modifications:

- Forced-Controlled Dynamic Coupling DMPs (FDC-DMP): Introduced by Shen et al. [22], this framework adds a coupling term driven by interaction forces directly into the DMP acceleration equation, enabling the robot to dynamically modify its path (e.g., avoiding obstacles during bus body polishing) without altering the global target.
- B-Spline DMPs (BDMPs): Wang et al. [13] proposed replacing the standard Gaussian basis functions in DMPs with B-splines. This modification significantly improves trajectory modeling accuracy with fewer basis functions. The forcing term is parameterized as $f(s) = \sum_i B_j(s)w_j$, where $B_j(s)$ represents the B-spline basis functions. When optimized using Policy Improvement with Path Integrals (PI^2) Reinforcement Learning, BDMPs demonstrate high generalization capability for polishing trajectories and force profiles under unseen workpiece positions.
- Riemannian DMPs: Liao et al. [24] extended DMPs to Riemannian manifolds (e.g., Cartesian space and 2D sphere manifolds) to simultaneously model human motion, 3D endpoint stiffness, and contact forces from a one-shot demonstration, solved via Quadratic Programming (QP). Trajectories on the manifold M are generated by mapping the states to the tangent space, T_xM, maintaining geometrical properties of robot orientations.
- Neural Network-Augmented DMPs: Wang et al. [9] integrated a Phase-Modulated Diagonal Recurrent Neural Network (PMDRNN) with DMPs to adaptively predict trajectory offsets based on real-time force-tracking deviations, mitigating environmental uncertainties.

3.5.2. Probabilistic and Statistical Models

Probabilistic models (33.3% of studies) represent demonstrations as joint probability distributions, capturing task variations, correlations, and rhythmic patterns across multiple demonstrations:

- Gaussian Mixture Models and Regression (GMM-GMR): Used to model the joint distribution of time, space, and force parameters. Wu et al. [11] utilized GMMs to encode human polishing dynamics, combining GMR with a variable-impedance controller to regulate contact compliance. Zhai et al. [19] integrated GMM-GMR with a vector-valued Gaussian Process (GP) to online-modulate robotic trajectories when subjected to human external forces.
- Probabilistic Movement Primitives (ProMPs): Unlike DMPs, ProMPs capture the statistical variance of demonstrations. ProMPs represent a trajectory as a linear combination of basis functions: $y(t) = \Phi(t)w + \epsilon$, where $w \sim \mathrm{N}(\mu_w, \sum_w)$ captures the statistical variance across multiple demonstrations. Wang et al. [14] developed Arc Length ProMPs (AL-ProMP), mapping the probability distribution of contact forces to spatial coordinates (arc-length s_l) rather than time t, formulating the trajectory as $y(s_l) = \Phi(s_l)w + \epsilon$. This formulation prevents trajectory distortions during non-linear speed scaling.
- Fourier Movement Primitives (FMPs): Grounded in signal processing, Kulak et al. [15] proposed FMPs using Fourier series as basis functions: $y(t) = a_0 + \sum_{k=1}^{K}(a_k\cos(k\omega t) + b_k\sin(k\omega t))$ FMPs approximate periodic, multi-frequency signals (e.g., circular polishing patterns) from unaligned demonstrations without requiring phase alignment or frequency extraction.

3.5.3. Deep Learning and Generative AI Architectures

Representing 16.7% of the studies, these approaches leverage deep neural networks to directly map high-dimensional visual or proprioceptive observations to continuous control actions:

- Diffusion Policies: Ke et al. [10] and Li et al. [5] utilized diffusion models to generate continuous, expert-like motion–force trajectories. In the DP-RRL framework [5], the Diffusion Policy generates a trajectory distribution by iteratively denoising a random sequence $x^K, x^{K-1}, \ldots, x^0$ using a noise predictor $\epsilon_\theta\left(x^k,\ k,\ O\right)$ conditioned on observation O. The residual RL agent then predicts a displacement, ΔF_t, to correct the reference force based on contact dynamics.
- Neural Ordinary Differential Equations (Hyper-NODEs): Xu et al. [7] developed a Hyper-NODE architecture to generate continuous position and orientation (quaternion) trajectories. The system dynamics are modeled as follows:

$$\frac{dh(t)}{dt} = f(h(t), t, \theta) \tag{4}$$

  where $h(t)$ represents the continuous hidden state, and θ are the weights generated by the hypernetwork. Paired with Control Barrier Functions (CBFs), the system guarantees obstacle avoidance in local non-polishing areas (LNP areas) while estimating admittance control parameters.
- Neural Network Morphing: Möhl et al. [12] designed a keypoint-driven non-linear morphing network to transfer demonstrated trajectories between 3D point clouds of geometrically similar objects without CAD models.

3.5.4. Autonomous Dynamical Systems (DSs)

Grounded in control theory, DS-based approaches (8.3% of studies) model robot motion as time-invariant differential equations to ensure global asymptotic stability and immediate reactivity to physical perturbations:

- Stable Limit Cycles: Duarte et al. [27] represented periodic human polishing movements (e.g., circular or elliptical motions) using a time-invariant DS with a stable limit cycle attractor. The system is formulated as a second-order non-linear dynamical system, $\dot{x} = f(x)$, where the trajectories are forced to converge asymptotically to a closed orbit C. This formulation guarantees that the robot converges back to the demonstrated polishing pattern even after being physically displaced.
- DS-based Imitation Learning: Si et al. [8] proposed a dynamically stable energy-field framework to provide virtual haptic guidance during teleoperated human demonstrations, reducing operator physical workload.

3.5.5. Direct Adaptive Control and Parameter Estimation

Representing 4.2% of the studies [4], direct adaptive control and parameter estimation focus on decoupling demonstrations into pure "motion skills" (discrete pose sequences) and "force skills" (desired normal forces) to achieve accurate skill transfer through a computed-torque impedance control law on complex surfaces.

3.6. RQ2: System and Sensory Configuration

To support LfD-based robotic surface finishing, mechatronic architectures must be carefully designed to capture high-fidelity human demonstrations, perceive unstructured environments, and execute contact-rich tasks safely.

3.6.1. Sensory Modalities and Multimodal Perception

- Force and Torque Sensing: As the primary driver of closed-loop execution, force feedback is integrated in 87.5% of the studies (either as the sole sensor or in multimodal setups). While end-effector 6-DOF F/T sensors are standard, Hamdan et al. [23] proposed a dual-force sensor configuration. One sensor measures the human operator's guiding force (F_h), while the second measures the tool–workpiece interaction force (F_{int}). By calculating the environmental reaction force as follows, the system isolates the environmental dynamics (stiffness, friction) from human guidance inputs:

$$F_{env} = F_h - F_{int} \tag{5}$$

- Vision and Spatial Perception: To handle geometrically complex surfaces, depth sensors (e.g., overhead or wrist-mounted RGB-D cameras) are integrated. These sensors capture raw 3D point clouds, which are processed via PointNet++ or keypoint-based neural networks to reconstruct surface meshes [10] or guide trajectory morphing [12].
- Kinematic and Biometric Tracking: Teleoperation and kinesthetic demonstration interfaces utilize haptic devices (e.g., Geomagic Touch) or wearable inertial measurement units (IMUs). More advanced setups incorporate surface electromyography (sEMG) sensors on the human arm to capture synergistic muscle activity, translating muscle co-contraction directly into robot joint stiffness parameters.
- Instrumented Tools: To facilitate platform-independent demonstrations, Fischer et al. [6] designed custom instrumented tools and mechanical alignment plates to record high-quality contact forces and orientations directly on the workpiece.

3.6.2. Control Architectures

Due to the rigid nature of physical contact during surface finishing, pure position control is unsafe. Consequently, studies rely on indirect force control strategies:

- Variable-Impedance and Admittance Control: These strategies model the robot–workpiece interface as a mass-spring-damper system. The low-level dynamic behavior is governed by the admittance control law:

$$M_d(\ddot{x} - \ddot{x}_d) + D_d(\dot{x} - \dot{x}_d) + K_d(x - x_d) = F_{ext} - F_{ref} \tag{6}$$

  Here, M_d, D_d and K_d are the desired mass, damping, and stiffness matrices. x_d is the reference trajectory, x is the actual position, F_{ext} is the external interaction force, and F_{ref} is the target reference force. Variable-impedance controllers [11,16] modulate stiffness K_d and damping D_d online. For instance, stiffness is reduced when transitioning onto hard, brittle materials (e.g., iron) to prevent impact chatter and increased when transitioning onto soft materials (e.g., wood) to ensure uniform material removal.
- Iterative Learning Control (ILC): To compensate for repetitive tracking errors, ILC is integrated with impedance control. Zhang et al. [20] combined GMM trajectory models with Dynamic Time Warping ILC (DTW-ILC), iteratively updating the robot's reference path based on stiffness estimation to achieve fast force convergence over multiple polishing passes.

3.6.3. Robotic Systems and Hardware Integration

The Franka Emika Panda (7-DOF) and Universal Robots (UR5, UR10e) cobots are highly preferred for their backdrivability, joint torque sensing, and active gravity compensation, which are crucial for kinesthetic teaching. For high-precision micro-machining (e.g., deburring)—where cobots lack sufficient stiffness—specialized setups are used, such as a

6-DOF hexapod paired with high-speed piezoelectric actuators controlled via teleoperated haptic interfaces [18]. Control loops typically separate high-frequency force control (running at 1000 Hz) from low-frequency visual perception and policy planning (running at 10–30 Hz) to maintain real-time stability.

3.7. RQ3: Evaluation and Performance Metrics

The effectiveness of LfD-based robotic finishing is evaluated using a multi-tiered framework, spanning kinematic accuracy, dynamic force tracking, and physical surface quality.

3.7.1. Kinematic and Trajectory Accuracy Metrics

- Dynamic Time Warping (DTW) Distance: Quantifies the spatiotemporal similarity between the demonstrated human trajectory and the robot's executed path, especially when feed rates vary.
- Root Mean Square Error (RMSE): Calculates the spatial deviation (in millimeters) between the executed end-effector path x_i and the demonstrated trajectory $\hat{x}_i$ over N samples:

$$\text{RMSE}_p = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \|x_i - \hat{x}_i\|^2} \tag{7}$$

- Pearson Correlation Coefficient (r): Measures the shape similarity of the trajectories:

$$r = \frac{\sum_{i=1}^{N} (x_i - \overline{x})(\hat{x}_i - \overline{\hat{x}})}{\sqrt{\sum_{i=1}^{N}(x_i - \overline{x})^2}\sqrt{\sum_{i=1}^{N}(\hat{x}_i - \overline{\hat{x}})^2}} \tag{8}$$

  with values closer to 1.0 indicating high imitation fidelity.
- Relative Smoothness (r_{c_p} for position, r_{c_q} for orientation): Evaluates the jerk of the generated trajectory to ensure smooth robotic motion.

3.7.2. Force Tracking and Dynamic Interaction Metrics

- Force Root Mean Square Error (Force RMSE): The primary metric to evaluate force-tracking performance, measuring the deviation between the executed contact force F_i and the demonstrated reference force profile $\hat{F}_i$:

$$\text{RMSE}_F = \sqrt{\frac{1}{N}\sum_{i=1}^{N} (F_i - \hat{F}_i)^2} \tag{9}$$

- Maximum Impact Force (F_{max}): Evaluates system compliance and safety during initial tool contact or material transitions.
- Mean Force Deviation (ΔF): Quantifies the stability of the normal force during continuous polishing.

3.7.3. Surface Quality and Process-Specific Metrics

- Surface Roughness (Ra, Rq, Rz): Measured using contact profilometers or white-light interferometers. A reduction in average roughness (Ra) verifies successful surface smoothing.
- Material Removal Rate (MRR): Replicating the expert's material removal strategy is evaluated based on Preston's equation:

$$\Delta z = k_p P v \Delta t \tag{10}$$

where Δz is the thickness of the removed material, k_p is Preston's coefficient (depending on tool and material properties), P is the contact pressure (directly proportional to normal force F_n), and v is the relative tool speed. Previous studies have estimated these parameters online to adapt forces dynamically on curved surfaces [4].

- Remaining Stain Ratio (RSR): In cleaning applications, image segmentation is used to calculate the percentage of stains remaining on the surface post-execution.

3.8. RQ4: Industrial Deployment Challenges

Despite laboratory advancements, several critical technical barriers hinder the transition of LfD-based finishing to industrial assembly lines.

3.8.1. The Sim-to-Real Gap and Complex Contact Dynamics

Deep Reinforcement Learning (DRL) and generative models (e.g., Diffusion Policies) require thousands of training episodes. However, physics engines struggle to simulate contact dynamics (friction hysteresis, tool deformation, workpiece stiffness, and material removal profiles) with high fidelity. Consequently, policies trained in simulation exhibit degraded performance when deployed on physical hardware.

3.8.2. Demonstration Quality and Hardware Constraints

Kinesthetic teaching is prone to human error and kinematic limitations:

- Kinematic Interference: The physical weight and joint limits of the robot arm restrict the operator's natural movement, leading to distorted demonstrations.
- Sensor Noise: High-speed spindle rotation and pneumatic tool vibrations generate significant mechanical noise, degrading force/torque sensor readings during the teaching phase.
- Cognitive Overload: Controlling the robot's spatial path, tool orientation, and contact force simultaneously in real time places a high cognitive demand on the human expert.

3.8.3. Generalization to Complex Geometries and LNP Areas

- Local Non-Polishing (LNP) Areas: Industrial components often feature functional geometry (e.g., threaded holes, slots, and ribs) that must remain untouched. Autonomously detecting and avoiding these LNP areas while maintaining a constant normal force on the surrounding freeform surface remains an open control problem.
- Geometric Generalization: Trajectory generalization models (e.g., DMPs) often distort orientations when scaling trajectories to highly curved, non-planar workpieces, risking collision or uneven polishing.

3.8.4. Multimodal Perception and Computational Complexity

Fusing high-dimensional point clouds, haptic data, and force feedback in real time requires substantial computational resources. Traditional LfD models like Gaussian Processes or GMMs suffer from poor scalability, making high-frequency closed-loop control (>500 Hz) difficult to achieve on standard industrial controllers.

3.8.5. Need for Robust Human-in-the-Loop (HITL) Systems

One-shot offline learning is vulnerable to environmental changes. Industrial deployment requires active online correction mechanisms (HITL), allowing human operators to intuitively intervene, adjust control parameters (e.g., stiffness or feed rate), and correct localized trajectory segments in real time without restarting the programming process.

4. Discussion

4.1. Methodological and Academic Perspectives: Temporal Evolution (2016–2026)

The systematic mapping of the literature reveals a significant temporal evolution in the methodological landscape of robotic Learning from Demonstration for surface-finishing tasks. In the early phase of the reviewed period (2016–2020), research was heavily dominated by Dynamic Movement Primitives (DMPs) and classical statistical representations, such as Gaussian Mixture Models with Gaussian Mixture Regression (GMM-GMR). These approaches focused primarily on the mathematical parameterization of human hand trajectories, relying on second-order spring-damper equations to guarantee spatial scaling and temporal robustness. However, these early frameworks faced critical limitations in contact-rich tasks due to their inability to dynamically adapt contact force profiles under unmodeled surface variations.

From 2021 onward, a distinct paradigm shift toward adaptive mechatronic integration and hybrid control is observed. Researchers increasingly coupled movement primitives with variable-impedance and admittance control, transitioning LfD from a purely kinematic imitation tool into a dynamic, multimodal learning framework. The latest period (2024–2026) is characterized by the emergence of deep generative AI models, such as Diffusion Policies and Neural Ordinary Differential Equations (Neural ODEs). These frameworks process high-dimensional spatial point clouds directly to generate obstacle-aware continuous trajectories in real time, eliminating the need for rigid pre-computed workpiece geometry models. This progression highlights a transition from simple trajectory replication to complex, sensory-driven reactive behaviors capable of generalizing to entirely new surface topologies.

Beyond the algorithmic representations, the systematic mapping reveals notable quantitative benchmark trends across the primary studies (as summarized in Table 9). In terms of force-tracking accuracy, Dynamic Movement Primitives (DMPs) coupled with adaptive variable-impedance or admittance controllers consistently report force-tracking root-mean-square errors (RMSEs) between 0.5 N and 1.5 N under laboratory conditions (e.g., [9,24]). Regarding surface quality, polishing frameworks report substantial improvements, with average surface roughness (Ra) typically reduced from initial post-machining values of 1.5–3.0 µm down to finished tolerances of 0.1–0.3 µm (e.g., [4,20]). For spatial trajectory reproduction, deep learning and generative AI models (such as Diffusion Policies and Hyper-NODEs) demonstrate high spatial fidelity, achieving Pearson's correlation coefficients (r) greater than 0.95 and endpoint trajectory RMSE values under 2.0 mm relative to the expert human demonstrations (e.g., [5,7,10]). These values provide a quantitative baseline for evaluating future LfD-based robotic finishing frameworks.

Table 9. Study-Level Experimental Contexts and Quantitative Performance Benchmarks.

Study (Ref.)	Robot Platform and Sensors	Workpiece Material and Geometry	Tooling and Normal Force Range	Specific Performance Metrics
[4]	Franka Emika Panda (7-DOF), joint torque sensors	Wooden violin body (highly curved)	Pneumatic sanding tool; $F_z = 6.2$ N	Successful skill transfer; roughness R_a reduced below manual craft baseline
[5]	7-DOF robot arm, overhead/wrist RGB-D cameras, F/T sensor	Basin (unstructured, variable geometry)	Cleaning brush / polishing	Cleaning success rate: 93% ± 3%; trajectory RMSE < 2.0 mm; correlation r > 0.95

Table 9. *Cont.*

Study (Ref.)	Robot Platform and Sensors	Workpiece Material and Geometry	Tooling and Normal Force Range	Specific Performance Metrics
[7]	Franka Emika Panda (7-DOF), joint torque sensors	Curved workpieces (HT200 Cast Iron)	Polishing disk	Force fluctuation: mean = 0.0171 N; endpoint trajectory RMSE: 0.07046 mm (WP1)/0.0015 mm (sim); roughness R_a: 14.39 µm to 5.29 µm (WP1, 63.2% reduction) and 2.63 µm to 0.459 µm (WP2, 56.1% reduction)
[9]	Franka Emika Panda (7-DOF), joint torque sensors	Curved mold surface	Polishing tool with force feedback	Force RMSE: 1.51 N (PMDRNN-DMP) vs. 2.14 N (PMNN) and 1.83 N (baseline)
[10]	Collaborative robot, RGB-D camera	Freeform curved workpiece	Polishing pad	High trajectory generalization on freeform meshes via Visual-Guided Diffusion Policy and Mesh-DMP
[14]	Collaborative robot, end-effector F/T sensor	Aluminum alloy plate	Polishing pad; variable speed-scaling	Roughness R_a: reduced from 0.909 µm to 0.106 µm after 10 repetitions (AL-ProMP)
[20]	Collaborative robot, end-effector F/T sensor	Curved metal plate	Polishing tool; target polishing force	Force-tracking error: mean decreased to 0.18 N (SD: 0.12 N) via DTW-ILC + PBIC
[24]	Franka Emika Panda (7-DOF), joint torque + F/T sensor	Aluminum plate (height step) and brick	Polishing wool pad; variable target forces	Force RMSE: 0.8588 N (with QP) vs. 3.6290 N (without QP) on brick; 2.2619 N on aluminum

4.2. Task-Specific Synthesis: Polishing vs. Deburring

A key finding of this systematic mapping is the severe methodological imbalance between polishing and deburring applications. Out of the 24 included primary studies, polishing and cleaning operations represent the vast majority, while only 3 studies [17,18,25] focus primarily on robotic deburring. This asymmetry is driven by the fundamentally different physical dynamics of the two processes. Polishing is characterized by continuous, low-frequency contact forces over smooth surfaces, which align well with the spatial smoothness assumptions of primitives like DMPs or Gaussian Processes. In contrast, deburring involves highly non-linear, high-frequency impact forces encountered when the tool contacts rigid, variable-sized burrs. These transient contact dynamics introduce severe chattering, tool-wear uncertainties, and risk of mechanical failure, which are exceptionally difficult to capture via standard imitation learning.

Consequently, the three deburring studies in the corpus relied on specialized teleoperated haptic interfaces or specified DMPs (sDMPs) optimized via heuristic algorithms (e.g., Particle Swarm Optimization) to learn safety-critical force reflexes. To address the high-frequency impact dynamics unique to deburring, these studies utilized specialized mechatronic and algorithmic control strategies:

- High-Bandwidth Micro–Macro Control: Ref. [18] implemented a dual-stage macro-micro control scheme where a slow hexapod robot (macro) tracks the tool path, while a high-bandwidth piezoelectric actuator (micro-manipulator with a 1-mm stroke operating at 1000 Hz in a haptic loop) executes micro-adjustments to maintain a constant normal force.

- Force-Coupled Specified DMPs (sDMPs): Ref. [17] developed an sDMP architecture that replaces the standard spatial DMP forcing function with a Gaussian function of the contact force. This force-coupling term enables the robot to adaptively alter its trajectory and feed rate during physical interaction.
- Meta-Heuristic Optimization of Compliance: To tune the coupling parameters of force-sensitive primitives, meta-heuristic algorithms like Particle Swarm Optimization (PSO) are used to minimize trajectory error under oscillatory contact forces [17]. The underrepresentation of deburring represents a critical research gap.

4.3. Economic and Operational Implications for SMEs

For Small- and Medium-sized Enterprises (SMEs) operating in high-mix, low-volume (small batch) production environments, traditional robotic automation is often economically challenging due to the high setup and programming costs of CAD/CAM systems. Although no formal economic or production-scale cost–benefit analysis was synthesized in the reviewed literature, qualitative analysis suggests that LfD could potentially transform this financial landscape by providing three hypothesized operational advantages:

- Reduced Reliance on Programming Expertise: LfD has the potential to lower the reliance on highly specialized robotic programming engineers. By using intuitive kinesthetic teaching or teleoperation, existing shop-floor operators can demonstrate finishing strategies. However, specialized engineering oversight is still required for low-level controller tuning, safety configuration, and system integration.
- Setup and Changeover Time Reduction: LfD is projected to shorten setup and changeover times. While traditional CAD/CAM programming requires watertight 3D models and offline path planning, an LfD system can theoretically adapt to new geometries in laboratory trials within minutes through a single expert demonstration or neural morphing. Production-scale validation is necessary to confirm these time savings under factory conditions where part tolerances drift.
- Potential Hardware Cost Reductions: LfD may lower hardware barriers by utilizing consumer-grade sensors (e.g., depth cameras) and avoiding expensive force-accurate simulators. Nevertheless, the capital cost of collaborative robots and industrial force/torque sensors remains a significant financial barrier for many SMEs. By consolidating these hypothesized mechatronic, personnel, and temporal savings, LfD represents a promising and potentially viable pathway for SMEs to automate complex finishing operations under highly variable production requirements. However, empirical studies on cycle time, reliability, and economics are required to confirm these benefits.

4.4. Comparative Synthesis of the Included Studies

To systematically map the LfD literature in robotic surface finishing, Table 10 synthesizes the core commonalities (similarities) and key distinctions (differences and unique contributions) of all 24 primary studies included in the corpus.

To contextualize the strengths and limitations of Learning from Demonstration (LfD) in robotic surface finishing, it must be critically compared with alternative paradigms: traditional computer-aided-design/manufacturing (CAD/CAM)-based trajectory generation and pure Reinforcement Learning (RL). Table 8 summarizes these paradigms across key technical and operational dimensions.

Traditional CAD/CAM-based finishing has long been the industry standard. It utilizes nominal CAD models to calculate tool paths, relying on active force compensation (e.g., force-control end-effectors) to handle surface irregularities. While offering exceptional geometric precision for high-volume production, its dependency on precise CAD templates

makes it unsuitable for high-mix, low-volume production or parts with high dimensional variance. Furthermore, reprogramming CAD/CAM systems is time-consuming and requires specialized engineering skills. In contrast, LfD platforms allow shop-floor operators to teach robots new paths in minutes via kinesthetic guidance or teleoperation, bypassing CAD models entirely and accelerating changeover times. This represents a significant advantage in SME settings, where flexibility is paramount.

Table 10. Comparative Analysis of Similarities and Differences Across the Included Studies.

No.	Study	LfD Category	Core Commonalities with Corpus	Unique Differences and Key Contributions
1	Min et al. [4]	Direct Impedance	Focuses on polishing; utilizes force control and collaborative robot platforms.	Decouples human demonstrations into separate "motion skills" (discrete pose sequences) and "force skills," validating on complex violin surfaces.
2	Li et al. [5]	Deep Learning	Focuses on cleaning/polishing; utilizes force control and collaborative robot platforms.	Combines a generative Diffusion Policy (for motion–force generation) with a Residual RL agent (for online force updates) using point clouds.
3	Fischer et al. [6]	Probabilistic	Focuses on cleaning; utilizes haptic interfaces and movement primitives.	Developed a location-invariant few-shot cleaning framework using an instrumented manual tool to capture expert data independent of the robot platform.
4	Xu et al. [7]	Deep Learning	Focuses on polishing; utilizes admittance control and collaborative robot platforms.	Uses Hyper-NODEs to generate smooth position–quaternion trajectories, combined with CLF/CBF for obstacle avoidance in LNP areas.
5	Si et al. [8]	Dynamical Systems	Focuses on polishing; utilizes force control and haptic teleoperation interfaces.	Introduces a dynamically stable energy-field-based virtual haptic guidance force that decays iteratively to reduce operator workload.
6	Wang et al. [9]	DMP	Focuses on polishing; utilizes force control and collaborative robot platforms.	Integrates a Phase-Modulated Diagonal Recurrent Neural Network (PMDRNN) to adaptively predict trajectory offsets based on force errors.
7	Ke et al. [10]	DMP	Focuses on polishing; utilizes force control and depth cameras.	Combines a Diffusion Policy to generate continuous spatial actions from RGB-D images and embeds them on freeform meshes using Mesh-DMP
8	Wu et al. [11]	Probabilistic	Focuses on polishing; utilizes force control and variable-impedance architectures	Learns a non-parametric, globally stable GMM for rhythmic motions, optimizing variable impedance via GMR to minimize control torque.
9	Möhl et al. [12]	Deep Learning	Focuses on trajectory transfer; utilizes depth cameras and vision data	Direct trajectory transfer between scan point clouds of similar objects using keypoint-driven neural network morphing, without CAD models.
10	Wang et al. [13]	DMP	Focuses on polishing; utilizes force control and collaborative robot platforms.	Introduces B-spline DMPs (BDMPs) requiring fewer basis functions, optimized via Policy Improvement with Path Integrals (PI^2) for generalization.

Table 10. *Cont.*

No.	Study	LfD Category	Core Commonalities with Corpus	Unique Differences and Key Contributions
11	Wang et al. [14]	Probabilistic	Focuses on polishing; utilizes force control and collaborative robot platforms.	Proposes Arc-Length ProMPs (AL-ProMP) to decouple force scaling and speed scaling in the spatial coordinate (arc-length) domain.
12	Kulak et al. [15]	Probabilistic	Focuses on polishing; utilizes collaborative robots and joint torque sensing.	Uses Fourier series basis functions (FMP) to learn periodic tasks from unaligned demonstrations without temporal or phase registration.
13	Wu et al. [16]	Probabilistic	Focuses on polishing; utilizes force control and variable-impedance architectures.	Tailored for Heterogeneous Material Components (HMCs); uses PC-GMM-DS and SMoGP to regulate rapid transitions across wood–iron splicing.
14	Parvizi et al. [17]	DMP	Focuses on deburring; utilizes haptic teleoperation interfaces	Uses Particle Swarm Optimization (PSO) to parameterize sDMPs to capture expert force responses under sharp corner and circular geometries.
15	Acikgoz et al. [18]	DMP	Focuses on deburring; utilizes haptic teleoperation interfaces.	Developed for deburring; human guides a 1-DOF haptic knob, and a high-speed piezoelectric actuator executes micro-adjustments on the workpiece.
16	Zhai et al. [19]	Probabilistic	Focuses on polishing; utilizes force control and collaborative robot platforms.	Couples GMM-GMR with a vector-valued Gaussian Process to enable online trajectory deformation under human physical intervention.
17	Zhang et al. [20]	Probabilistic	Focuses on polishing; utilizes force control and collaborative robot platforms.	Combines GMM with Dynamic Time Warping Iterative Learning Control (DTW-ILC) to estimate environment stiffness and update reference paths.
18	Haninger et al. [21]	Probabilistic	Focuses on co-manipulation polishing; utilizes force control and collaborative robot platforms	Captures task uncertainty using Gaussian Processes (GPs) and solves trajectory and impedance planning online using a non-linear MPC.
19	Shen et al. [22]	DMP	Focuses on polishing; utilizes force control and collaborative robot platforms	Introduces force-controlled dynamic coupling terms (FDC-DMP) using virtual coupling forces to dynamically alter local paths.
20	Hamdan et al. [23]	Deep Learning	Focuses on polishing; utilizes force control and admittance control	Employs a dual-force sensor configuration to isolate human guide forces (F_h) from environmental reaction forces (F_{int}), training an MLP.
21	Liao et al. [24]	DMP	Focuses on polishing; utilizes force control and collaborative robot platforms.	Uses Riemannian DMPs and QP optimization to simultaneously learn motion, 3D endpoint stiffness, and applied forces from a one-shot demonstration.
22	Li et al. [25]	DMP	Focuses on finishing; utilizes collaborative robot platforms and movement primitives.	Pairs DMPs with machine vision object detection to automatically recognize workpiece locations and generalize trajectory paths.

Table 10. *Cont.*

No.	Study	LfD Category	Core Commonalities with Corpus	Unique Differences and Key Contributions
23	Nemec et al. [26]	DMP	Focuses on polishing; utilizes force control and haptic teleoperation interfaces	Models the tool as a "virtual mechanism" (augmented kinematic chain) for redundancy resolution, refining trajectories via Iterative Learning Control.
24	Duarte et al. [27]	Dynamical Systems	Focuses on polishing; utilizes collaborative robot platforms and human data.	Models all circular/ellipse polishing motions as a time-invariant dynamical system with a stable limit cycle attractor, mapping non-verbal human cues.

On the other hand, pure RL approaches optimize finishing policies by maximizing a reward function (e.g., material removal rate and force-tracking error) through trial-and-error. Although pure RL can theoretically exceed human finishing performance, its sample complexity is a massive barrier. Generating millions of interaction steps on physical hardware is impractical due to tool wear, workpiece damage and safety risks. Bridging this gap via physics simulators is challenging, as existing simulators cannot model contact dynamics (e.g., friction hysteresis, compliance, tool degradation) with high fidelity. While some hybrid architectures combine LfD with residual RL to bootstrap learning and ensure safe exploration (e.g., [5]), pure RL remains restricted. LfD remains the most viable paradigm for complex, contact-rich tasks on physical hardware, providing safe, sample-efficient initialization by copying human haptic and motion skills. However, it remains bounded by the performance ceiling of the human demonstrator.

4.5. Demonstration Efficiency Across LfD Algorithms

Demonstration efficiency is defined as the volume of human training data (i.e., the number of physical demonstrations) required to learn a stable, generalizable control policy. However, high number of physical demonstration requirements create significant challenges in industrial settings. They directly lead to increased programming downtime, operator fatigue, and higher setup costs. Ultimately, these drawbacks can negate the flexible automation benefits of LfD.

The empirical evidence in the corpus reveals distinct trade-offs across these families:

- Dynamic Movement Primitives (DMPs) and Variants: DMPs are highly demonstration-efficient, typically generalizing from a single human demonstration (one-shot learning). By leveraging a second-order spring-damper system as a structural prior, DMPs guarantee global asymptotic convergence to the goal, allowing the learnable forcing term to be fitted using standard regression techniques (e.g., Locally Weighted Regression or Least Squares) on one trajectory. In surface finishing, DMPs can adapt a single learned polishing path to varying coordinate frames (e.g., [22]) or learn force–velocity profiles from a one-shot demonstration (e.g., [17,24]). This makes them highly suitable for small-batch industrial lines where changeover time must be kept under a few minutes.
- Probabilistic and Statistical Models: Unlike DMPs, probabilistic models like Gaussian Mixture Models/Regression (GMM-GMR) or Probabilistic Movement Primitives (ProMPs) require multiple demonstrations (typically 5 to 15) to capture the statistical covariance, temporal correlations, and task variability across different trials. Zhang et al. [20] evaluated path learning errors using 5, 10, and 15 demonstrations, showing that the path error converges to a baseline (under 1.0 mm) after 10 demonstrations,

and that additional demonstrations capture human variability without significantly improving spatial precision. This suggests a practical limit of approximately 10 demonstrations for statistical learning in contact-rich tasks. ProMP models like AL-ProMP map these probabilities to spatial arc-lengths rather than time, preserving the statistical correlation of contact forces along the path from a handful of demonstrations. These models are highly feasible for industrial use, requiring less than 30 min of manual demonstration.

- Autonomous Dynamical Systems (DSs): DS-based models guarantee immediate, time-invariant reactivity to physical perturbations. However, learning a stable, high-dimensional vector field over the entire state space requires demonstrations initiated from multiple starting points to define the attraction basin. Refs. [8,11] utilized five physical demonstrations to train stable guidance fields and non-linear dynamics, respectively, while Duarte et al. [27] noted that multiple demonstrations are required to extrapolate periodic polishing limit cycles.
- Deep Learning and Generative AI Architectures: Deep neural networks (e.g., Diffusion Policies, Neural ODEs) have high capacity for visual–haptic integration but suffer from extreme sample inefficiency because they lack physical structural priors (e.g., second-order attractor dynamics). Ref. [5] collected 300 expert demonstrations on physical hardware to train a multimodal Diffusion Policy for basin cleaning, demonstrating the high data-collection overhead. Fischer et al. [6] explicitly addressed this bottleneck in industrial cleaning, noting that traditional imitation learning requires hundreds of demonstrations (e.g., 650 trials), and proposed a few-shot learning approach utilizing approximately 10 demonstrations by using instrumented tool-based segmentation. Thus, deep generative architectures remain restricted to high-volume manufacturing lines due to the days required for data collection and model training.
- Direct Adaptive Control and Parameter Estimation: Decoupling the task into a spatial path (motion skill) and a normal force profile (force skill) allows direct transfer from a single demonstration [4], as the active controller (e.g., adaptive admittance or computed-torque impedance control) compensates for surface irregularities online without requiring a statistical model. This achieves high demonstration efficiency, though it is limited to tasks that can be represented by 1D normal force profiles on simple geometries.

4.6. Identification of Research Gaps in the LfD Literature

The synthesis of the 24 primary studies reveals four critical research gaps that currently limit the adoption of LfD-based robotic finishing in industrial environments:

1. Tactile Complexity and Mechanical Hazards of Deburring:

There is a severe imbalance in task representation, with 21 studies focusing on polishing/cleaning and only 3 addressing deburring. Polishing involves low-frequency contact normal to smooth surfaces, which aligns well with standard spatial primitives. In contrast, deburring involves high-frequency, non-linear impact forces when contacting rigid, stochastic burrs. These dynamics introduce chattering, high tool wear, and risk of mechanical failure, which standard LfD frameworks fail to model or stabilize.

2. Generalization to Freeform Geometries without CAD Templates:

Most current LfD frameworks assume flat or simple curved workspaces. Generalizing a demonstrated trajectory to highly curved, freeform 3D objects frequently leads to spatial trajectory distortion and orientation misalignment (quaternion errors). While keypoint morphing [12] and Mesh-DMPs [10] have emerged to address this, they require

high computational power and struggle to maintain constant contact forces on highly nonplanar surfaces.

3. The Sim-to-Real Gap in Physical Interaction Physics:

Modern deep learning models (e.g., Diffusion Policies, Residual RL) offer excellent trajectory generation capabilities but require thousands of training episodes. Running these on physical hardware is impractical due to safety risks and tool wear, necessitating training in simulation. However, physics engines cannot model contact physics (friction hysteresis, tool compliance, material removal profiles) with high fidelity. Consequently, policies trained in simulation degrade when deployed on physical factory floors.

4. Lack of Intuitive Online Human-in-the-Loop (HITL) Correction:

The majority of LfD frameworks rely on offline, one-shot demonstrations. If the demonstration is noisy or if the workshop environment drifts (e.g., due to tool wear), the robot cannot adapt. There is a lack of real-time, online interactive correction systems that allow human operators to physically intervene, correct localized trajectory segments, and adapt impedance parameters on the fly.

4.7. How the Current Study Addresses These Research Gaps

The present systematic mapping study contributes directly to addressing these literature gaps by establishing a structured, evidence-based foundation for future LfD research:

- Addressing the Task Imbalance: By exposing the severe deficit in deburring research (only 12.5% of the corpus), this study provides a clear technical analysis of why deburring is exceptionally challenging (high-frequency transient contact, mechanical chattering) and catalogues the haptic and PSO-based control strategies used by the few successful deburring studies. This guides future researchers toward the control and haptic configurations necessary to tackle deburring.
- Providing a Mechatronic and Sensory Roadmap: This study maps the exact mechatronic configurations required to support LfD (such as the dual-force sensor configuration of [23] to isolate human forces from environmental reaction forces). By documenting the frequency of force-only (70.8%) vs. multimodal (16.7%) setups, it provides a design guideline for building hardware platforms capable of perceiving both geometric and tactile environments.
- Synthesizing Quantitative Performance Benchmarks: To assist researchers in bridging the sim-to-real gap, this work extracted and synthesized concrete quantitative baseline performance metrics reported from physical experiments (force-tracking RMSEs of 0.5–1.5 N, trajectory RMSEs under 2.0 mm, and finished surface roughness *Ra* of 0.1–0.3 µm). These benchmarks provide a standard against which simulated policies can be verified and validated.
- Structuring Algorithmic Trade-offs: By categorizing the LfD algorithms into five methodological families (Table 6) and comparing their characteristics (Table 10), this study maps which algorithms are best suited for specific challenges. For instance, it highlights that while DMPs excel at spatial scaling, probabilistic models (like GMM-GMR) are better suited for online trajectory deformation during human intervention (HITL), and generative AI is best for visual point-cloud mapping, allowing practitioners to select the optimal control architecture.

5. Limitations of This Study

The findings and sectoral implications of this systematic mapping study should be interpreted in light of several methodological limitations:

- Database Coverage and Search Strategy: The literature search was restricted to Scopus, Web of Science, IEEE Xplore, and Google Scholar. While these are the primary repositories for engineering and robotics research, some publications indexed in regional or specialized databases may have been omitted. Additionally, keyword-based search queries centered around terms like "Learning from Demonstration" and "DMP" might have missed relevant studies that utilize alternative terminology, such as "skill transfer" or "human–robot co-manipulation," despite sharing the same underlying architecture.
- Language Bias: In accordance with the screening protocol, only articles published in the English language were included. Given that 62.5% of the included literature originates from China, and that countries like Japan and Germany possess strong academic and industrial backgrounds in robotic manufacturing, excluding non-English publications likely introduced a language bias. Highly innovative papers published in Chinese, Japanese, or German may have been overlooked.
- Exclusion of the Grey Literature: To ensure scientific rigor, only peer-reviewed journal articles and conference proceedings were included, while patents, technical white papers, and corporate reports were excluded. Because real-world industrial implementations of LfD are frequently protected as commercial trade secrets, the exclusion of the grey literature may have limited our ability to map the exact degree of current commercial adoption.
- Selection Bias and Single Screener Limitations: Since the screening, eligibility selection, and data extraction processes were conducted by a single primary researcher—rather than by two independent reviewers as recommended by the PRISMA 2020 guidelines [1]—there is an inherent risk of selection bias. To mitigate this limitation, a retrospective inter-rater reliability check was subsequently conducted by a second independent reviewer on a random 20% subset of records at the Title/Abstract stage and 100% of the full-text articles. The analysis demonstrated strong agreement (Cohen's Kappa $\kappa = 0.88$ for Title/Abstract screening, and $\kappa = 1.00$ for full-text eligibility), indicating that the selection process was highly reliable. However, the lack of a second independent auditor during the initial, live extraction phase remains a minor methodological limitation. Although borderline or ambiguous cases were re-evaluated multiple times to ensure coding consistency and adherence to the eligibility protocol, the lack of a second independent auditor represents a methodological limitation.
- Temporal Concentration and Recency Bias: A notable limitation of this systematic mapping study is the high concentration of the literature in the most recent publication years, with 58.3% of the included primary studies (14 out of 24) published between 2024 and 2025. This temporal clustering reflects a surge in research interest driven by the maturation of collaborative robot hardware and the emergence of deep generative control policies. However, this introduces a potential recency bias, over-representing very recent and popular algorithmic paradigms (e.g., Diffusion Policies, Neural ODEs) that have not yet undergone long-term industrial testing. Furthermore, because these studies are very recent, they have had limited time to establish significant citation impact or undergo independent replication and verification by the wider research community. Consequently, some reported performance benchmarks must be treated as early, laboratory-validated indicators rather than long-term, industry-proven baselines.

6. Future Directions

Based on the synthesized evidence, future research in LfD-based robotic surface finishing should prioritize three key technological integrations:

6.1. Multimodal Perception and High-Dimensional Datasets

Future LfD frameworks must transition from low-dimensional haptic representations to domain-specific multimodal perception that links visual geometry with mechatronic contact physics. To accurately estimate surface-finishing quality (e.g., surface roughness parameter Ra) and material removal rates (MRRs) in real time, visual feedback from 3D point clouds must be tightly coupled with high-frequency force/torque (F/T) profiles. For instance, rather than relying on simple 1-D force templates, future models should employ geometry-aware trajectory morphing to transfer expert surface-finishing trajectories onto complex, unstructured workpieces registered from optical point clouds [12]. Furthermore, database collection must address mechatronic coupling constraints by separately capturing human-applied guidance forces and tool–workpiece interaction forces, which can be accomplished using dual-force-sensor mechatronic configurations [23]. To monitor material-specific machining states, future LfD frameworks should directly integrate high-frequency vibration and acoustic emission (AE) sensors into the state representation of movement primitives (e.g., ProMPs or Neural ODEs) to adaptively regulate feed rates and spindle speeds, suppressing tool chatter and compensating for dynamic tool wear. Scaling these perception models requires the development of public, high-dimensional demonstration datasets that systematically capture the coupled dynamics of feed rate, spindle speed, tool pose, and contact force across diverse workpiece materials and geometries.

6.2. Bridging the Sim-to-Real Gap and Safe Exploration

While deep generative models, such as Diffusion Policies and Neural Ordinary Differential Equations (Neural ODEs), have shown promise for generating complex finishing policies, bridging the simulation-to-real (sim-to-real) gap remains a major mechatronic bottleneck. Simulating contact-rich physics is notoriously difficult due to non-linear friction, workpiece stiffness uncertainties, and tool wear. Future research should focus on embedding physics-informed contact models (such as mechatronic models of tool–workpiece compliance and material removal equations) directly into LfD policies as structural priors. To ensure stability and safety during execution, future architectures should combine offline imitation learning (which models the global task structure) with online residual control loops (which handle local contact dynamics). For example, a hybrid policy could utilize a Diffusion Policy to generate nominal tool trajectories, while a high-frequency residual Reinforcement Learning (RL) or an adaptive admittance control loop adjusts the commanded force online to adapt to unmodeled contact stiffness [5]. Furthermore, machining heterogeneous material components (HMCs) under geometric variations requires learning mechatronically stable non-linear dynamics—such as contraction-stable or Lyapunov-stable dynamical systems [11,16]. These stable LfD models should be coupled with Control Barrier Functions (CBFs) to guarantee state and force safety, preventing tool or workpiece damage when navigating near Local Non-Polishing (LNP) zones on freeform surfaces [7].

6.3. Interactive Human-in-the-Loop (HITL) Systems

Traditional LfD frameworks are limited by the high physical and mental workload imposed on human experts during the demonstration of contact-rich tasks. To improve demonstration quality and safety, future research must introduce active haptic guidance systems. For example, during teleoperated or kinesthetic demonstration, the robot can provide virtual mechatronic guidance forces to assist the operator in maintaining consistent tool–workpiece contact [8]. Moreover, to resolve the compliance and workspace mismatch between rigid robotic manipulators and the human hand, future LfD architectures must transition from one-shot training to interactive online refinement. Non-expert operators should be able to refine learned skills (such as B-spline DMPs or ProMPs [13]) using intu-

itive, online haptic correction interfaces. Specifically, future frameworks should integrate Dynamic Time Warping Iterative Learning Control (DTW-ILC) with probabilistic skill models, enabling the robot to iteratively estimate workpiece contact stiffness and update the path-updating laws from sparse, localized human physical corrections [20]. This dynamic stiffness adaptation allows the robot to generalize learned trajectories across differing object scales and geometries without requiring full task re-demonstrations.

6.4. Control and Mechatronic Architectures for High-Frequency Deburring Dynamics

To address the polishing–deburring imbalance, future research must move beyond standard imitation learning of spatial paths and focus on specific control architectures designed for high-frequency, non-linear contact dynamics. We propose the following mechatronic and control strategies as concrete pathways for robust robotic deburring:

- Dual-Stage (Macro–Micro) Mechatronics: Standard robotic manipulators are limited by low joint control bandwidth (under 50 Hz) due to gear elasticity and link inertia, which is insufficient to suppress high-frequency chatter. Future systems should integrate high-bandwidth auxiliary actuators—such as piezoelectric end-effectors or active voice-coil-driven compliant tool holders operating at or above 1 kHz haptic loops [18]—to execute fast normal force adjustments, leaving coarse spatial path tracking to the macro-manipulator.
- Force-Coupled Specified Movement Primitives: Classical movement primitives (DMPs, ProMPs) must transition to force-coupled formulations, such as Specified DMPs (sDMPs) [17]. Directly embedding contact force feedback as a coupling term within the canonical and transformation equations allows the robot to dynamically scale feed rates or deform paths when encountering sudden force spikes.
- Adaptive Variable-Impedance Control (AVIC): During rigid workpiece–tool collisions, static impedance parameters lead to excessive force peaks or chatter. Control systems must adaptively modulate stiffness and damping coefficients online (e.g., via evolutionary optimization such as PSO [17] or residual Reinforcement Learning [5]) to increase damping at the moment of burr impact, absorbing kinetic energy and stabilizing the physical interaction.

7. Conclusions

This systematic mapping study provides a comprehensive, PRISMA-2020 [1]-compliant synthesis of Learning-from-Demonstration (LfD) methodologies applied to robotic deburring and polishing between 2016 and 2026. By analyzing 24 primary peer-reviewed studies, this work established a structured taxonomy categorizing LfD approaches into five methodological clusters, with Dynamic Movement Primitives (DMPs) and probabilistic models representing the dominant frameworks, and deep generative policies (e.g., Diffusion Policies, Neural ODEs) emerging as the State of the Art.

The synthesis of mechatronic configurations highlights that force/torque sensing remains the foundational sensory modality, utilized exclusively in 70.8% of studies and in 87.5% of studies in any configuration (including multimodal setups), and is integrated primarily with variable-impedance and admittance control architectures to guarantee physical interaction compliance. Quantitatively, the State of the Art demonstrates high performance in laboratory environments, achieving force-tracking errors (RMSEs) between 0.5 N and 1.5 N, trajectory-reproduction errors under 2.0 mm ($r > 0.95$), and surface roughness (*Ra*) reductions to 0.1–0.3 µm. However, these values are highly dependent on specific experimental setups and should not be generalized across all industrial processes.

Several technical barriers must be resolved before LfD can be widely adopted on factory floors. These include the sim-to-real gap in contact physics, the high-frequency im-

pact dynamics of deburring, the autonomous avoidance of local non-polishing areas (LNP areas), and the lack of intuitive online human-in-the-loop correction interfaces. Addressing these gaps will facilitate the transition of LfD from a laboratory concept to a promising trend that shows potential as a flexible, cost-effective automation tool, particularly for SMEs operating in high-mix, low-volume manufacturing environments, provided that future research establishes its long-term economic, cycle-time, and reliability feasibility.

Supplementary Materials: The following are available online at https://www.mdpi.com/article/10.3390/jmmp10080293/s1, PRISMA 2020 checklist.

Funding: This research received no external funding.

Data Availability Statement: The data extraction sheets, search strategy logs, and quality assessment sheets are publicly available in the Open Science Framework (OSF) repository: https://doi.org/10.17605/OSF.IO/5Y6BR.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Page, M.J.; McKenzie, J.E.; Bossuyt, P.M.; Boutron, I.; Hoffmann, T.C.; Mulrow, C.D.; Shamseer, L.; Tetzlaff, J.M.; Akl, E.A.; Brennan, S.E.; et al. The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ* **2021**, *372*, n71. [CrossRef] [PubMed]
2. Kim, Y.; Sloth, C.; Kramberger, A. Skill transfer for surface finishing tasks based on estimation of key parameters. In *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*; IEEE: New York, NY, USA, 2022; pp. 2148–2153.
3. Eiband, T.; Leimbach, L.; Nottensteiner, K.; Albu-Schäffer, A. Extraction of Robotic Surface Processing Strategies from Human Demonstrations. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: New York, NY, USA, 2025; pp. 20494–20500.
4. Min, K.; Ni, F.; Chen, Z.; Liu, H. A Force Control Method Integrating Human Skills for Complex Surface Finishing. *Machines* **2024**, *12*, 756. [CrossRef]
5. Li, Y.; Lyu, Q.; Yang, J.; Salam, Y.; Wang, W. A Hybrid Framework Using Diffusion Policy and Residual RL for Force-Sensitive Robotic Manipulation. *IEEE Robot. Autom. Lett.* **2025**, *10*, 10266–10273. [CrossRef]
6. Fischer, A.; Unger, C.; Kugi, A.; Hartl-Nesic, C. Few-Shot Learning of a Force-Based Industrial Cleaning Process using an Instrumented Tool. *IFAC-PapersOnLine* **2025**, *59*, 103–108. [CrossRef]
7. Xu, X.; Qian, K.; Liu, A.; Yue, Z.; Huang, W. Polishing via ODEs: Adaptive admittance control for robot polishing based on Neural ODEs. *J. Manuf. Process.* **2025**, *155*, 428–442. [CrossRef]
8. Si, W.; Jin, Z.; Lu, Z.; Wang, N.; Yang, C. A Stable Guidance Method for Teleoperation-based Robot Learning from Demonstration. In *IEEE International Conference on Automation Science and Engineering (CASE)*; IEEE: New York, NY, USA, 2024; pp. 2376–2381. [CrossRef]
9. Wang, Y.; Zheng, Z.; Chen, C.; Wang, Z.; Gao, Z.; Peng, F.; Tang, X.; Yan, R. Adaptive Tuning of Robotic Polishing Skills based on Force Feedback Model. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)*; IEEE: New York, NY, USA, 2023; pp. 1–7. [CrossRef]
10. Ke, S.; Zhang, J.; Zhao, H.; Guo, Y.; Wei, Z.; Pan, J.; Ding, H. Visual-Guided Diffusion Policy and Mesh-DMP Integration for Robotic Freeform Surface Polishing. In *International Conference on Intelligent Robotics and Applications*; Springer: Singapore, 2025; pp. 77–90. [CrossRef]
11. Wu, H.; Zhai, X.; Wu, X.; Gu, S.; Liao, Z.; Xu, Z.; Zhou, X. Learning Stable Nonlinear Dynamics and Interactive Force-Aware Variable Impedance Control for Robotic Contact Tasks. *Procedia Comput. Sci.* **2023**, *226*, 127–133. [CrossRef]
12. Möhl, P.; Pratheepkumar, A.; Ikeda, M.; Pichler, A. Morphing based transfer of demonstrated surface finishing trajectories to point clouds of similar objects. *Procedia Comput. Sci.* **2025**, *253*, 1002–1011. [CrossRef]
13. Wang, Y.; Chen, C.; Hong, Y.; Zheng, Z.; Gao, Z.; Peng, F.; Yan, R.; Tang, X. PI2-BDMPs in combination with contact force model: A robotic polishing skill learning and generalization approach. *IEEE/ASME Trans. Mechatron.* **2025**, *30*, 978–988. [CrossRef]
14. Wang, Y.; Chen, C.; Peng, F.; Zheng, Z.; Gao, Z.; Yan, R.; Tang, X. AL-ProMP: Force-relevant skills learning and generalization method for robotic polishing. *Robot. Comput.-Integr. Manuf.* **2023**, *82*, 102538. [CrossRef]
15. Kulak, T.; Silvério, J.; Calinon, S. Fourier movement primitives: An approach for learning rhythmic robot skills from demonstrations. In Proceedings of the Robotics: Science and Systems (RSS), Virtual, 12–16 July 2020. [CrossRef]
16. Wu, H.; Zhai, X.; Zheng, H.; Liao, Z.; Xu, Z.; Zhou, X. Learning stability-guaranteed skill and adaptive control strategies from demonstrations for heterogeneous component robotic machining. *J. Manuf. Process.* **2025**, *151*, 506–520. [CrossRef]

17. Parvizi, P.; Ugurlu, M.C.; Acikgoz, K.; Konukseven, E.I. Parametrization of robotic deburring process with motor skills from motion primitives of human skill model. In *International Conference on Methods and Models in Automation and Robotics (MMAR)*; IEEE: New York, NY, USA, 2017; pp. 373–378. [CrossRef]
18. Acikgoz, K.; Parvizi, P.; Donder, A.; Ugurlu, M.C.; Konukseven, E.I. Dynamic movement primitives and force feedback: Teleoperation in precision grinding process. In *10th International Conference on Electrical and Electronics Engineering (ELECO), Bursa, Turkey, 30 November–2 December 2017*; IEEE: New York, NY, USA, 2017; pp. 722–726, ISBN 978-605-01-1134-7.
19. Zhai, X.; Ou, Y.; Xu, Z.; Jiang, L.; Zhou, X.; Wu, H. Effective learning and online modulation for robotic variable impedance skills. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)*; IEEE: New York, NY, USA, 2022; pp. 1–6. [CrossRef]
20. Zhang, R.; Xia, J.; Ma, J.; Huang, D.; Zhang, X.; Li, Y. Human-robot interactive skill learning and correction for polishing based on dynamic time warping iterative learning control. *IEEE Trans. Control Syst. Technol.* **2024**, *32*, 2310–2320. [CrossRef]
21. Haninger, K.; Hegeler, C.; Peternel, L. Model predictive impedance control with Gaussian processes for human and environment interaction. *Robot. Auton. Syst.* **2023**, *165*, 104431. [CrossRef]
22. Shen, N.; Mao, J.; Li, J.; Mao, Z. Research on trajectory learning and modification method based on improved dynamic movement primitives. *Robot. Comput.-Integr. Manuf.* **2024**, *89*, 102748. [CrossRef]
23. Hamdan, S.; Aydin, Y.; Oztop, E.; Basdogan, C. Robotic Learning of Haptic Skills from Expert Demonstration for Contact-Rich Manufacturing Tasks. In *IEEE International Conference on Automation Science and Engineering (CASE)*; IEEE: New York, NY, USA, 2024; pp. 2334–2341. [CrossRef]
24. Liao, Z.; Tassi, F.; Gong, C.; Leonori, M.; Zhao, F.; Jiang, G.; Ajoudani, A. Simultaneously learning of motion, stiffness, and force from human demonstration based on Riemannian DMP and QP optimization. *IEEE Trans. Autom. Sci. Eng.* **2024**, *22*, 7773–7785. [CrossRef]
25. Li, X.; Yang, C.; Feng, Y. The Generalization of Robot Skills Based on Dynamic Movement Primitives. *IFAC-Pap. OnLine* **2020**, *53*, 265–270. [CrossRef]
26. Nemec, B.; Yasuda, K.; Mullennix, N.; Likar, N.; Ude, A. Learning by demonstration and adaptation of finishing operations using virtual mechanism approach. In *IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: New York, NY, USA, 2018; pp. 7219–7225. [CrossRef]
27. Duarte, N.F.; Santos-Victor, J. Robot Imitation of Polishing Motions by Observing Humans: From Human Non-Verbal Cues to Stable Limit Cycles. In *IEEE International Conference on Development and Learning (ICDL)*; IEEE: New York, NY, USA, 2024; pp. 1–7. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.