Control Strategies for Contact Force Regulation on Unknown Surfaces in Assistive Human–Robot Co-Manipulation

A Systematic Literature Review

Peter Ralph Desir   Arno Stienen

Delft University of Technology

Preprint, July 2026

Abstract

Musculoskeletal disorders from repetitive surface-finishing tasks remain a leading cause of occupational disability. Robotic assistive devices can reduce operator muscular load, yet existing control strategies typically assume known surface geometry, decouple force regulation from human guidance, or treat all contact directions identically. This systematic review investigates how a control strategy for a mechanically coupled assistive device can maintain desired contact force along the surface normal while providing power-assist in tangential directions, when surface geometry is unknown and must be estimated online from force-motion interaction. The research question is decomposed into five components: online surface estimation (C1), physical human–robot interaction (C2), haptic shared control (C3), anisotropic force control (C4), and ergonomic load reduction (C5).

Following PRISMA 2020 guidelines, Scopus and IEEE Xplore were searched on 8 January 2026 using seven structured queries, supplemented by citation chaining from 23 seed papers. Each paper was scored on a five-component rubric (ABSENT / IMPLICIT / EXPLICIT) and assigned to quality tiers (1 DIAMOND, 27 GOLD, 31 SILVER, 5 RIVAL, 7 FOUNDATIONAL).

From approximately 12,600 identified records, title-and-abstract screening retained 72 reports for full-text assessment, of which 71 studies were included. Component coverage is uneven: C2 is addressed by 49 papers (69%), C3 by 44 (62%), C4 by 38 (54%), C5 by 30 (42%), and C1 by 25 (35%). A progressive funnel analysis reveals that only 7 papers combine C1+C2+C4 at IMPLICIT level or above, and **0 of 71** combine all five components with genuine online surface estimation from force-motion interaction. The gap is structural: surface estimation methods assume autonomous operation, pHRI methods assume known geometry, and ergonomic validations use simplified surfaces. A search update run on 12 July 2026, immediately prior to posting, screened 113 newly indexed records against the same rubric, with every sought report retrieved and assessed in full text; none combines these components, and the gap stands.

The review confirms a specific integration gap: the five required capabilities have each been demonstrated individually, but the assumptions underlying each are mutually incompatible. Limitations include single-reviewer screening and coverage restricted to two databases, partially mitigated by citation chaining. The protocol was not prospectively registered. No external funding was received.

Contents

1 Introduction

1.1. Context and Motivation

Industrial surface finishing (grinding, sanding, polishing, and deburring) requires operators to apply sustained contact forces against workpiece surfaces while simultaneously guiding tools along complex geometries, often across full working shifts. The combination of forceful exertion, repetitive motion, and constrained posture places operators at significant risk of musculoskeletal disorders (MSDs). This section examines the problem from four perspectives (occupational health, economic, social, and technical), each of which independently motivates the need for assistive technology in contact tasks.

From an occupational health perspective, MSDs represent the single largest category of work-related illness worldwide. The World Health Organization estimates that 1.71 billion people globally live with musculoskeletal conditions, which account for approximately 149 million years lived with disability, accounting for 17% of all YLDs worldwide [1]. A secondary analysis of the Global Burden of Disease 2019 study reported 322.75 million incident MSD cases globally, with occupational ergonomic factors identified as the highest attributable risk factor [2]. The burden peaks among workers aged 50–54 [2], an age band strongly represented in skilled manufacturing trades.

In Europe, the scale is equally stark. A comprehensive EU-OSHA report drawing on the European Working Conditions Survey (N = 31,612) found that 58% of EU-28 workers report at least one MSD, with backache affecting 43% and upper-limb pain 41% [3]. MSDs account for 60% of all work-related health problems in the EU [3]. Craft and trades workers, the occupational group most directly involved in surface finishing, report an MSD prevalence of 65%, among the highest of any occupational group (exceeded only by skilled agricultural workers at 69% and plant and machine operators at 66%) [3]. More recent data from the 2025 OSH Pulse survey (N = 25,688) confirm that 28% of EU-27 workers report bone, joint, or muscle problems caused or made worse by work in the preceding 12 months, with the figure rising to 29% in manufacturing and 34% among skilled and semi-skilled workers [4]. The epidemiological evidence for upper-extremity MSDs in tool-handling occupations has been established since the landmark NIOSH review by Bernard [5]. Bernard's review identified forceful exertion, the defining physical demand of surface finishing, as a risk factor for upper-extremity disorders, with the strongest evidence for combined force and repetition [5], yet three decades later, MSD prevalence among manual workers remains persistently high: the 2025 OSH Pulse survey still reports rates of 28% across all workers and 34% among skilled and semi-skilled occupations [4].

The economic burden of MSDs provides a second, independent justification for investing in assistive technology. Bevan [6] estimated that MSDs cost approximately EUR 240 billion annually across the EU (up to 2% of GDP) through a combination of healthcare expenditure, productivity losses, and social welfare payments. MSDs account for 53% of all work-related diseases in the EU-15 and are responsible for 50% of absences exceeding three working days and 60% of cases of permanent work incapacity [6]. The broader cost of work-related injuries and illness in the EU has been estimated at 3.3% of GDP, or approximately EUR 476 billion per year [3]. In the United States, the National Research Council and Institute of Medicine estimated that MSDs impose USD 13–20 billion annually in workers' compensation costs alone, with total economic burden, including indirect costs such as lost productivity, retraining, and reduced quality, reaching USD 45–54 billion [7]. For industries that depend on manual surface preparation (petrochemical maintenance, shipbuilding, and metal fabrication), these costs manifest directly as workers' compensation claims, production delays from absent skilled workers, and the expense of retraining replacements.

Demographic trends compound this challenge rather than alleviating it. The manufacturing workforce in industrialized nations is ageing. Stedmon et al. [8] reported that, in the ten UK manufacturing occupations most affected by impending retirements, workers over 45 made up 35–57% of the workforce between 1998 and 2008, a proportion that demographic trends suggest has continued to grow. Ageing is directly relevant to MSD risk because muscle strength declines 15–50% from its peak in the third decade of life, and peak oxygen uptake decreases 20–30% between ages 20 and 49 [8]. For tasks demanding sustained grip force and upper-limb endurance (such as grinding, sanding, and polishing), this decline in capacity is particularly consequential. As physiological capacity diminishes while task demands remain constant, ergonomic interventions become increasingly important for sustaining an ageing workforce in physically demanding roles [8]; robotic assistance is one such intervention.

Collaborative robotics offers a technological path forward, but current implementations fall short of what contact tasks demand. The global cobot market has grown rapidly: from 11,107 collaborative robot installations in 2017 (2.8% of all industrial robots) to 64,542 in 2024 (11.9%), representing nearly a six-fold increase in seven years [9]. This growth reflects industrial demand for robots that augment human capability rather than replace it. However, cobot deployments to date are dominated by non-contact tasks such as pick-and-place, assembly, packaging, and machine tending [9], where the robot operates in free space and physical interaction with the environment is minimal. The safety frameworks that enable human-robot coexistence, codified in ISO 10218 [10] and ISO/TS 15066 [11], govern the safety of human-robot contact itself; they do not address the regulation of sustained robot–workpiece interaction forces during contact tasks.

Surface finishing tasks present fundamentally different requirements. The robot must not only move cooperatively with the human operator but must simultaneously manage physical interaction with a workpiece whose geometry may be unknown, variable, or partially occluded. This demands *anisotropic* control behavior: the system must resist motion along the surface normal to maintain contact force while assisting motion in tangential directions to reduce operator effort. To the best of the authors' knowledge, no widely deployed commercial cobot platform currently provides this capability, and the research literature, as this review will demonstrate, has addressed the constituent subproblems largely in isolation rather than as an integrated system.

1.2. Problem Statement

The core technical challenge is *geometric uncertainty*. Unlike assembly or pick-and-place tasks, where workpiece geometry is defined by CAD models or fixture design, surface finishing operates on workpieces that are often:

- **Unknown a priori:** the workpiece may be a freeform surface with no available digital model.
- **Varying during operation:** the surface normal direction changes continuously as the tool moves across the surface.
- **Sensed only through contact:** vision-based surface estimation requires line-of-sight, which is frequently occluded by the tool, the operator's hands, or debris from the finishing process itself.

Standard admittance control, the dominant paradigm for power-assist in physical human-robot interaction (pHRI), amplifies operator motion intent but treats all Cartesian directions equally. This isotropy is acceptable for free-space tasks but fundamentally inadequate for contact tasks: the control should *resist* operator motion along the surface normal (to maintain the desired contact force for material removal) while *assisting* motion in tangential directions (to reduce the muscular effort required for surface coverage). Achieving this anisotropic behavior requires real-time knowledge of the surface normal at the contact point, which, given the constraints above, must be estimated online from force-motion sensor data rather than from prior geometric knowledge or vision.

The problem is therefore one of *coupled estimation and control*: the surface geometry must be inferred from the very interaction that the controller is regulating, and the controller must adapt its directional behavior based on an estimate that may be noisy, delayed, or temporarily unavailable.

1.3. Research Question

This systematic literature review investigates the following research question:

> *"How can a control strategy for a mechanically coupled assistive device maintain desired contact force along the surface normal while providing power-assist in tangential directions, when surface geometry is unknown and must be estimated online from force-motion interaction, to reduce operator muscular load during tool-handling tasks?"*

Component Decomposition

The research question decomposes into five components, each corresponding to a clause in the question:

1. **C1: Online Surface Estimation.** *"...surface geometry is unknown and must be estimated online from force-motion interaction..."*

   How can the surface normal be estimated in real time from force and motion sensor data, without prior geometric knowledge or vision?
2. **C2: Power-Assist / pHRI.** *"...mechanically coupled assistive device...providing power-assist..."*

   What constitutes effective power augmentation in a physically coupled human-robot system for tool handling?
3. **C3: Haptic Shared Control.** Implicit in the coupling between human intent and robot assistance.

   How should human and robot negotiate control authority during contact tasks with uncertain goals?
4. **C4: Anisotropic Force Control.** *"...maintain desired contact force along the surface normal while providing power-assist in tangential directions..."*

   How can admittance or impedance control be made direction-dependent, distinguishing normal force maintenance from tangential motion assistance?
5. **C5: Ergonomic Load Reduction.** *"...to reduce operator muscular load during tool-handling tasks..."*

   What evidence exists that such control strategies measurably reduce operator muscular effort, and how is this validated?

This five-component framework structures the entire review: the search strategy targets each component, the screening rubric scores papers against them, and the synthesis analyzes the literature along these five dimensions.

1.4. Scope and Boundaries

Included

- Continuous contact tasks: grinding, sanding, polishing, deburring, and other surface finishing operations.
- Physical human-robot interaction with haptic (mechanical) coupling between operator and robot.
- Control-based force anisotropy, meaning active directional modulation rather than passive mechanical properties.
- Shared control architectures with variable or negotiated authority allocation.
- Literature from 2015–2025, with foundational exceptions for seminal works predating this window.

Excluded

- Fully autonomous robotic machining with no human in the loop.
- Vision-only surface reconstruction without force-based estimation.
- Passive exoskeletons without active control.

- Medical and rehabilitation exoskeletons, unless the control logic is directly generalisable to industrial tool handling.
- Discrete contact tasks such as pick-and-place, peg-in-hole insertion, or spot welding.

1.5. Contribution and Structure

This review contributes a structured, critical analysis of the state of the art across all five components of the research question. Rather than surveying each component in isolation, it maps the *intersections*, identifying which combinations of components have been addressed in the literature and, critically, which have not. The review applies a tier-based quality framework (DIAMOND, GOLD, SILVER, RIVAL, FOUNDATIONAL) to weight the evidence according to relevance, and uses quantified gap claims (e.g., "0 of 71 papers combine all five components with genuine online force-motion estimation") rather than vague assertions of insufficient research. Where the literature reveals competing approaches, notably between force-based and vision-based surface estimation, the review presents these as alternative paradigms with different assumptions and trade-offs, rather than dismissing either.

The remainder of this review is organized as follows. Chapter 2 describes the systematic review methodology, including the PRISMA protocol, search strategy, screening criteria, and data extraction approach. Chapter 3 presents the descriptive results: corpus demographics, component coverage statistics, and the PRISMA flow diagram. Chapter 4 synthesizes the literature thematically along the five components, comparing papers against each other and forming critical arguments about the current state of the art. Chapter 5 discusses the findings in the context of the research question, integrates them with broader technical knowledge, and identifies specific research gaps with quantified evidence. Chapter 6 draws conclusions, formulates recommendations for the field grounded in the identified gaps, and provides a genuine assessment of this review's own limitations.

2

Methodology

This review applies the PRISMA 2020 protocol [12] to identify research at the intersection of five technical components derived from the research question.

2.1. Research Question Decomposition

The research question guiding this review is:

> *"How can a control strategy for a mechanically coupled assistive device maintain desired contact force along the surface normal while providing power-assist in tangential directions, when surface geometry is unknown and must be estimated online from force-motion interaction, to reduce operator muscular load during tool-handling tasks?"*

This question was systematically decomposed into five technical components, each representing a necessary capability for the envisioned system. Table 2.1 presents the decomposition with corresponding screening questions used during paper selection.

Table 2.1: Research question decomposition into five screening components.

Code	Component	RQ Clause	Screening Question
C1	Online Surface Estimation	"surface geometry is unknown and must be estimated online from force-motion interaction"	Does the paper estimate surface geometry online from force-motion data?
C2	Power-Assist / pHRI	"mechanically coupled assistive device... providing power-assist"	Is the system a physically coupled human-robot device providing force amplification?
C3	Haptic Shared Control	Implicit: human-robot coupling requires intent negotiation	Does the paper address how human and robot negotiate control authority?
C4	Anisotropic Force Control	"maintain desired contact force along the surface normal... in tangential directions"	Does the controller differentiate normal force maintenance from tangential motion?
C5	Ergonomic Load Reduction	"to reduce operator muscular load during tool-handling tasks"	Does the paper measure or target operator muscular load reduction?

2.2. Search Strategy

Database Selection

Searches were conducted in the following electronic databases:

- **Scopus**: Selected as the primary database due to its comprehensive coverage of engineering literature, including robotics, control systems, and human factors research.
- **IEEE Xplore**: Selected for its specialized coverage of robotics, automation, and control systems publications, including major conferences (ICRA, IROS, RO-MAN, Haptics Symposium).

Additionally, forward and backward citation chaining (snowballing) was performed on foundational works and high-relevance papers to mitigate keyword mismatch and capture influential works that may use different terminology.

Query Design Rationale

Seven queries (Q1–Q7) were designed, each targeting specific component combinations while maintaining robot-domain specificity through anchor terms (`robot`, `manipulator`, `pHRI`, `human-robot`). A global exclusion filter was applied to remove rehabilitation, exoskeleton, and gait-related literature that falls outside the scope of industrial tool-handling tasks. Table 2.2 presents the complete query strings as executed in Scopus. Equivalent queries were adapted for IEEE Xplore syntax.

Table 2.2: Search queries aligned to RQ components (Scopus syntax).

ID	Query String
Q1	`TITLE-ABS-KEY(("co-manipulation" OR "comanipulation" OR "manual guidance" OR "physical human-robot" OR "pHRI" OR "human-robot collaboration") AND ("contact force" OR "surface contact" OR "contact task" OR "environmental contact") AND ("admittance control" OR "impedance control")) AND NOT TITLE-ABS-KEY("exoskeleton" OR "rehabilitation" OR "lower limb" OR "gait" OR "walking")`
Q2	`TITLE-ABS-KEY(("force control" OR "hybrid force" OR "force/position" OR "force-position" OR "contact force control") AND ("polishing" OR "grinding" OR "finishing" OR "deburring" OR "sanding" OR "surface treatment" OR "fettling") AND ("robot" OR "manipulator" OR "robotic"))`
Q3	`TITLE-ABS-KEY(("online estimation" OR "real-time estimation" OR "contact-based estimation" OR "force-based estimation") AND ("surface normal" OR "surface geometry" OR "curvature" OR "constraint direction" OR "contact geometry")) AND NOT TITLE-ABS-KEY("vision" OR "camera" OR "point cloud" OR "3D reconstruction" OR "RGB-D")`
Q4	`TITLE-ABS-KEY(("anisotropic" OR "selective compliance" OR "task-frame" OR "task frame" OR "normal-tangential" OR "direction-dependent" OR "directional stiffness" OR "directional impedance") AND ("force control" OR "impedance control" OR "admittance control" OR "compliance control") AND ("contact" OR "surface" OR "constrained"))`
Q5	`TITLE-ABS-KEY(("unknown surface" OR "unstructured environment" OR "unknown geometry" OR "unmodeled surface" OR "uncertain environment" OR "unknown workpiece") AND ("force control" OR "contact force" OR "interaction force" OR "impedance" OR "admittance") AND ("robot" OR "manipulator" OR "robotic"))`
Q6	`TITLE-ABS-KEY(("power assist" OR "force amplification" OR "effort reduction" OR "load sharing" OR "intelligent assist device") AND ("contact" OR "surface" OR "tool" OR "polishing" OR "grinding" OR "finishing" OR "manipulation task")) AND TITLE-ABS-KEY(("human-robot" OR "pHRI" OR "co-manipulation" OR "comanipulation" OR "collaborative")) AND NOT TITLE-ABS-KEY("exoskeleton" OR "walking" OR "gait" OR "rehabilitation" OR "prosthesis")`
Q7	`TITLE-ABS-KEY(("task frame" OR "task-frame" OR "compliance frame" OR "constraint frame") AND ("force control" OR "hybrid control" OR "impedance" OR "admittance") AND ("robot" OR "manipulator"))`

Citation Chaining and Verification

Forward citation chaining was performed on 23 foundational and highly-cited works, archived as 24 exports (the Hogan citation search was exported both unrestricted and restricted to 2018–2026, with the overlap removed at deduplication), to capture influential papers that may use different terminology than the database queries. Table 2.3 lists the ten largest citation sources; the exports for the three most-cited classical works (Hogan, Raibert, and Mason) were restricted to 2018–2026, as their earlier citing literature was judged less likely to yield novel methodological contributions, while the remaining sources were retrieved unrestricted.

All retrieved records were archived in structured files (BibTeX for Scopus and citation exports, CSV for IEEE Xplore, Markdown for backward citation lists) and verified against source databases. The complete record counts by source are reported in Chapter 3, Table 3.1.

Table 2.3: Forward citation sources (partial list, top 10 by citation count).

Source Paper	Topic	Citations Found
Hogan (1985)	Impedance Control	745
Ajoudani et al. (2018)	pHRI Survey	660
Raibert & Craig (1981)	Hybrid Force/Position	567
Whitney (1987)	Force Control Framework	433
Peshkin et al. (2001)	Cobots	323
Kazerooni (1990)	Human Extenders	309
Bowyer & Rodriguez y Baena (2014)	Active Constraints	287
Colgate et al. (1996)	Intelligent Assist Devices	249
De Schutter & Van Brussel (1988)	Compliant Motion	235
Mason (1981)	Compliance and Force Control	228
Total from all sources		**7,150**

2.3. Selection Criteria

Inclusion Criteria

Papers were included if they met the following criteria:

1. **Publication type**: Peer-reviewed journal articles or major conference proceedings (ICRA, IROS, RO-MAN, Haptics Symposium, CASE).
2. **Date range**: Published between 2015 and 2026, with explicit exceptions for foundational works establishing theoretical frameworks.
3. **Task domain**: Focus on continuous contact tasks (surface following, grinding, sanding, polishing, deburring) or control methods applicable to such tasks.
4. **Human-robot coupling**: Physical coupling between human and robot where the human provides intentional input during task execution.
5. **Component coverage**: Addresses at least two of the five RQ components (C1–C5) with explicit methodology or results.

Criterion 4 defines the review's target system but was not applied as a blanket filter: fully autonomous studies whose surface-estimation or force-control methodology bears directly on C1 or C4 were retained as methodology references, and vision-based systems as alternative-paradigm comparators, rather than excluded under E1; the 22 included studies with C2 ABSENT in Table 3.4 enter on this basis, while E1's exclusions are records offering no such methodological contribution. The July 2026 search update (Section 2.7) applies the same convention.

Exclusion Criteria

Papers were excluded according to the codes defined in Table 2.4. Each excluded paper was assigned at least one exclusion code to maintain an audit trail.

Table 2.4: Exclusion codes with definitions and counts at title/abstract (T/A) screening.

Code	Reason for Exclusion	Count
E1	No physical human-robot coupling (fully autonomous systems)	150
E2	No contact force task (free-space motion only)	70
E3	Known or pre-scanned surface geometry (CAD-based control)	15
E4	Pure simulation without hardware validation, or review paper	13
E5	Medical, surgical, or rehabilitation domain (non-transferable methods)	15
E6	Non-English publication	0
E7	Duplicate record	1
E8	Walking, gait, or mobile navigation focus	14
Total excluded at T/A		**278**

Relevance Scoring

To manage the large volume of unique records (~9,534), an automated keyword-based relevance filter was applied before title/abstract screening. A screening script matched each record's title, abstract, and keywords against five keyword clusters, one per RQ component, and computed a weighted relevance score:

- **Core terms**: 2 points per matched core keyword of a component cluster
- **Associated terms**: 1 point per matched supporting keyword
- **Contact-task bonus**: 3 points if a contact-task term (grinding, sanding, polishing, deburring, or similar) was present
- **Advancement rule**: records advanced if they matched at least two component clusters, or one component cluster together with a contact-task term

The remaining 9,184 records were excluded as unlikely to address the multi-component research question, leaving 350 candidates for title/abstract screening with retained relevance scores between 5 and 8. This pre-screening step concentrated screening effort on records matching multiple RQ components; its sensitivity was not formally validated (Section 2.9).

2.4. Screening Process

A two-stage screening process was employed, with explicit documentation of decisions at each stage.

Stage 1: Title/Abstract Screening

After deduplication and keyword-based relevance filtering, 350 candidate papers were screened by title and abstract. Each record was assessed against the five component filters (Table 2.1) and assigned a preliminary relevance score.

Papers were classified as:

- **INCLUDE**: Clear evidence of addressing ≥ 2 components
- **EXCLUDE**: Assigned an exclusion code (E1–E8)
- **UNCERTAIN**: Flagged for full-text review

Stage 2: Full-Text Assessment

Papers passing Stage 1 (72 records) were retrieved and assessed using the structured data extraction template (Section 2.5). One paper was excluded at this stage: a Chinese-language publication for which the full text could not be retrieved through institutional access, precluding methodological assessment. This resulted in 71 papers for final inclusion.

Cynical Assessment Questions

To guard against superficial keyword matching and methodological overclaiming, five "cynical questions" were applied to borderline candidates during screening and, as a mandatory assessment block, to every paper reaching full-text extraction. These questions test whether papers genuinely address the RQ components or merely use related terminology without substantive contribution.

1. **The "Old News" Test**: *Is this 1990s Hybrid Force/Position control with modern buzzwords?* Papers repackaging classical Raibert & Craig (1981) or Mason (1981) approaches without genuine methodological advancement were flagged.
2. **The "CAD Trap"**: *Does "unknown surface" actually mean "unknown position on a known mesh"?* Many papers claim to handle "unknown surfaces" but rely on pre-scanned CAD models for geometry. Only position along the known surface is estimated online.
3. **The "Passive Anisotropy" Trap**: *Is direction-dependence from hardware design, not active control?* Some systems exhibit anisotropic behavior due to mechanical structure (e.g., SCARA kinematics) rather than control-based direction differentiation.
4. **The "Autonomous" Trap**: *Is the human actually in the control loop?* Papers claiming "human-robot collaboration" but running fully autonomous experiments with no real-time human input were excluded (E1).
5. **The "No Validation" Trap**: *Is muscular load reduction claimed but never measured?* Papers claiming ergonomic benefits without electromyography (EMG) measurement, subjective assessment, or force reduction metrics were noted for weak C5 evidence.

These questions were applied systematically during both title/abstract screening and full-text assessment, with explicit documentation when a paper "failed" a cynical test.

2.5. Data Extraction

For each included study, information was extracted using a structured template covering all five RQ components. Table 2.5 presents the extraction fields organized by component.

Each component was scored as ABSENT, IMPLICIT (mentioned but not central), or EXPLICIT (directly addressed with methodology or results). This three-level coding enables nuanced analysis of component coverage across the corpus.

2.6. Quality Assessment and Tiering

Papers passing full-text review were assigned to one of five tiers based on their relevance to the research question and methodological contribution. Table 2.6 defines the tiering criteria. Papers addressing fewer than two components were excluded during title/abstract screening (codes E1–E8) rather than being tiered, ensuring only substantively relevant work entered the final corpus.

Rationale for RIVAL Classification

Papers classified as RIVAL represent methodologically valid alternative approaches that rely on vision-based surface estimation or pre-scanned geometry rather than force-motion estimation. These papers were retained in the corpus for explicit comparison in the Discussion chapter, acknowledging that:

- Vision-based approaches require surface visibility and cannot operate during contact or in occluded conditions
- Pre-scanned approaches require prior knowledge of workpiece geometry, limiting applicability to known parts
- Force-motion estimation (the focus of this review) operates during contact and handles truly unknown surfaces

The RIVAL tier prevents dismissal of valid alternatives while maintaining clear scope boundaries.

Rationale for FOUNDATIONAL Classification

Seven papers published before 2015 were included as FOUNDATIONAL works despite the date restriction, as they establish theoretical frameworks essential for understanding the field:

- Stepien et al. [13]: Tool/workpiece contact force control for robotic deburring

Table 2.5: Data extraction fields by RQ component.

Component	Field	Options / Values
C1: Surface Estimation	Surface Assumption	Known / Unknown / Estimated / N/A
	Estimation Method	Force-based / Vision / Hybrid / None
	Online/Offline	Online / Offline / Both
	Geometry Validated	Flat / Curved / Complex / Sim only
	Estimation Details	Free text
C2: Power-Assist / pHRI	Human Role	Co-manipulation / Teleoperation / Supervisory
	Coupling Type	Direct (hand-on) / Master-slave / Wearable
	Assistance Type	Force amplification / Gravity comp / Motion guidance
	Hardware Platform	Industrial arm / Cobot / Custom / Simulation
C3: Haptic Shared Control	Control Authority	Fixed human / Fixed robot / Variable / Negotiated
	Intent Detection	None / Force-based / Motion-based / Learned
	Conflict Resolution	None / Human override / Blending / Explicit
	Force Feedback	Yes / No / Partial
C4: Force Control	Control Anisotropy	Isotropic / Anisotropic / Adaptive
	Normal Force Strategy	Constant / Adaptive / Tracking
	Tangential Strategy	Free / Assisted / Guided / Constrained
	Control Law	Admittance / Impedance / Hybrid / Other
	Frame Definition	Task / Tool / World / Sensor / Estimated
C5: Ergonomics	Ergonomic Validation	EMG / Subjective / Force reduction / None
	Metrics Reported	Muscle activity / Fatigue / Effort / Task load
	Task Duration	<1 min / 1–5 min / >5 min / Not reported
	User Study	Yes (n=?) / No
	Task Domain	Grinding / Sanding / Polishing / Assembly / Other
Quality Assessment	Components Addressed	C1–C5 coverage
	Relevance Score	1–10
	Key Contribution	Free text
	Identified Gap	Free text

Table 2.6: Paper tiering criteria and definitions.

Tier	Score	Criteria
DIAMOND	5	All five components (C1–C5) addressed with substantive methodology, at least four explicitly
GOLD	4	Four components with strong evidence, or three components with exceptional depth
SILVER	3	Three components well-addressed, suitable for supporting synthesis
RIVAL	n/a	Vision-based or pre-scanned surface approaches; methodologically sound but different assumptions than the review scope
FOUNDATIONAL	n/a	Pre-2015 seminal works establishing theoretical foundations

- Kazerooni [14]: Robotic deburring with unknown 2D geometry
- Jinno et al. [15]: Force-controlled grinding for unstructured tasks
- Houshangi [16]: Position/force control over unknown surfaces
- Tarokh and Bailey [17]: Adaptive fuzzy force control with unknown environments
- Baeten et al. [18]: Task frame formalism for hybrid vision/force control
- Iwasaki et al. [19]: Sliding mode force control for unknown environments

2.7. Search Coverage Assessment

The complete PRISMA flow diagram and quantitative outcomes of the selection process are presented in Chapter 3. This section addresses the estimated coverage of the search strategy.

Coverage Confidence

The coverage of a search strategy cannot be measured directly, since the denominator, the set of all relevant literature, is unknown. Table 2.7 therefore records a qualitative confidence assessment by component, based on how directly each component was targeted by the queries, the keyword filter, and the citation chaining.

Table 2.7: Search coverage confidence by RQ component, with measured relative recall (the fraction of included studies scoring each component that the database queries retrieved directly; the remainder entered the identification pool only through citation chaining).

Component	Confidence	Query recall	Notes
C1: Surface Estimation	High	17/25 (68%)	Directly targeted by multiple queries with vision exclusion filter
C2: Power-Assist / pHRI	High	38/49 (78%)	Directly targeted; cobot and intelligent assist device literature reached via citation chaining
C3: Shared Control	High	35/44 (80%)	Virtual-fixture terminology covered by the keyword filter and a dedicated chaining seed
C4: Anisotropic Force Control	High	27/38 (71%)	Task-frame and hybrid force/position terminology well covered
C5: Ergonomics	Moderate	24/30 (80%)	Weakest; relevant work appears in ergonomics venues outside both databases

As a measurable complement to this qualitative assessment, a relative-recall trace of the 71 included studies against the archived search exports shows that the database queries alone retrieved 52 of the 71 (73%); the remaining 19 (27%) entered the identification pool only through citation chaining, confirming the chaining arm's role in capturing relevant work that uses different terminology than the queries. The per-component split in Table 2.7 shows that the chaining arm mattered most for C1 (68% query recall) and C4 (71%), while C2, C3, and C5 were retrieved at 78–80% by the queries alone. Two caveats apply: the split measures reliance on the chaining arm within the combined strategy, not the coverage of the strategy as a whole, and relative recall is computed over the included studies only, so it cannot detect relevant work missed by both arms. The Moderate rating for C5 therefore stands despite its high query recall, since its residual risk lies in ergonomics venues that neither arm reached. The trace scripts are archived with the project materials.

The principal residual coverage risks are:

- Ergonomics and biomechanics literature published in non-robotics venues
- Very recent preprints not yet indexed in databases

- Non-English publications (excluded by E6)
- Unpublished theses and technical reports

Search Update Prior to Posting

The searches were re-executed by the first author on 12 July 2026, immediately before the public posting of this review, using the same Scopus Q1–Q7 strings (with a recency clause) and IEEE Xplore equivalents reconstructed term-for-term from them, under the unchanged screening rubric. This yielded 140 records for the window since the initial search (8 January 2026); deduplication left 125 unique records, of which 12 were already in the January reference set, giving 113 new candidates. The January keyword pre-filter was omitted, so all 113 went directly to title/abstract screening (same AI-assisted procedure as Section 2.8, all final decisions by the first author), which excluded 99. Of the remaining 14, all full texts were retrieved and assessed (the final five obtained via institutional access on 28 July 2026, after an initial retrieval attempt during the update had failed): ten met the inclusion criteria [20–29] (five with a human physically in the interaction loop, four retained as methodology references without one, under the same convention used for the January corpus, and one as a vision-based alternative-paradigm reference), while four were excluded (three fully autonomous systems with no human coupling and one simulation-only study).

Consistent with a dated search update, the ten eligible records are reported here but not merged into the analyzed corpus, which remains frozen at the 8 January 2026 search; the update therefore alters neither the corpus nor the tier assignments, and no conclusion changes: none of the 113 new records combines online surface estimation from force-motion interaction (C1) with physical human-robot coupling (C2), and none combines all five components. The three nearest candidates resolve the apparent overlap in full text: an intention-aware telerobotic ultrasonic-testing framework [22] acquires workpiece geometry through an operator-demonstrated teaching pass and closes its contact-force loop autonomously, a unified interactive force–impedance controller [20] fixes its force-controlled direction a priori from the specified desired wrench, handling unknown geometry through passivity rather than estimating it, and the closest new work on the estimation side, a friction-decoupled surface-normal estimator fusing force and velocity measurements [27], is fully autonomous and validated only in simulation on a single planar surface. With every sought report assessed in full text, the gap identified by this review stands as of July 2026.

2.8. Implementation: AI Assistance

The review protocol, comprising the research-question decomposition, the five-component rubric (C1–C5), the inclusion and exclusion criteria, the tier definitions, and the cynical assessment questions, was designed by the first author from first principles. An AI assistant supported its execution throughout: proposing candidate query formulations for each component of the decomposed research question; processing the ~12,600 exported bibliographic records (parsing, deduplication, and application of the keyword-based relevance scoring described in Section 2.3); producing first-pass assessments of the 350 candidate records against the component rubric during title/abstract screening, cross-checked by an independent validation pass and a completeness audit; executing the citation-chaining sweeps and the cynical re-assessments; and drafting the structured full-text extractions for the included papers. Screening and full-text assessment nevertheless remained the first author's throughout: every per-record assessment and every extraction draft was reviewed and accepted or corrected, with extractions spot-checked against the source PDFs, and all inclusion, exclusion, and tier decisions were made by the first author; the authors are responsible for the analysis, synthesis, and conclusions. AI assistance was additionally used to polish the manuscript's language and structure, and the July 2026 search update followed the same division of labour, with independent blind re-screening passes validating its first-pass verdicts.

Figure 2.1 makes this division of labour explicit for every stage of the review. Each hand-off between author and assistant crosses a named artifact (query strings, record sets, per-record assessments, extraction files, validation reports), and the review advances only through an author decision: no AI output enters the corpus, the tiers, or the conclusions without one. For reproducibility, the complete audit trail of the AI-executed steps (query exports, the per-record rubric assessments from both screening campaigns, the validation reports, and the update screening logs) is openly available in the review's repository.[1] The structured extractions themselves interleave verbatim excerpts from copyrighted papers and are therefore represented in the repository by a per-paper score matrix with per-cell provenance, the files being available from the authors on reasonable request. The repository additionally documents the stage protocols alongside these logs, so the procedure (decomposition, per-component query construction, evidence-first delegation, and author-gated adjudication) can be reapplied to other research questions.

2.9. Methodological Limitations

Several limitations of this review methodology should be acknowledged:

1. **Single human reviewer**: Screening and extraction were performed by a single human reviewer supported by AI tooling operating under the author-defined protocol (Section 2.8). A single-reviewer design introduces potential for inconsistent application of criteria and precludes inter-rater reliability measures; this was mitigated by explicit documentation of decisions and application of the cynical assessment questions.
2. **English language restriction**: Non-English publications were excluded (E6), potentially missing relevant work from non-English-speaking research groups.
3. **Database coverage**: While Scopus and IEEE Xplore provide strong coverage of robotics literature, some relevant work may appear in ergonomics-focused databases (e.g., Ergonomics Abstracts) not included in this search.
4. **Publication bias**: The review includes only published, peer-reviewed work, potentially excluding negative results or unsuccessful approaches documented in gray literature.
5. **Relevance scoring subjectivity**: Despite explicit criteria, the assignment of relevance scores (1–10) involves subjective judgment, particularly at tier boundaries.

These limitations are addressed in the Discussion chapter when interpreting findings and drawing conclusions about research gaps.

[1] `https://github.com/ralphdesir1/contact-force-review-companion` (private until this posting).

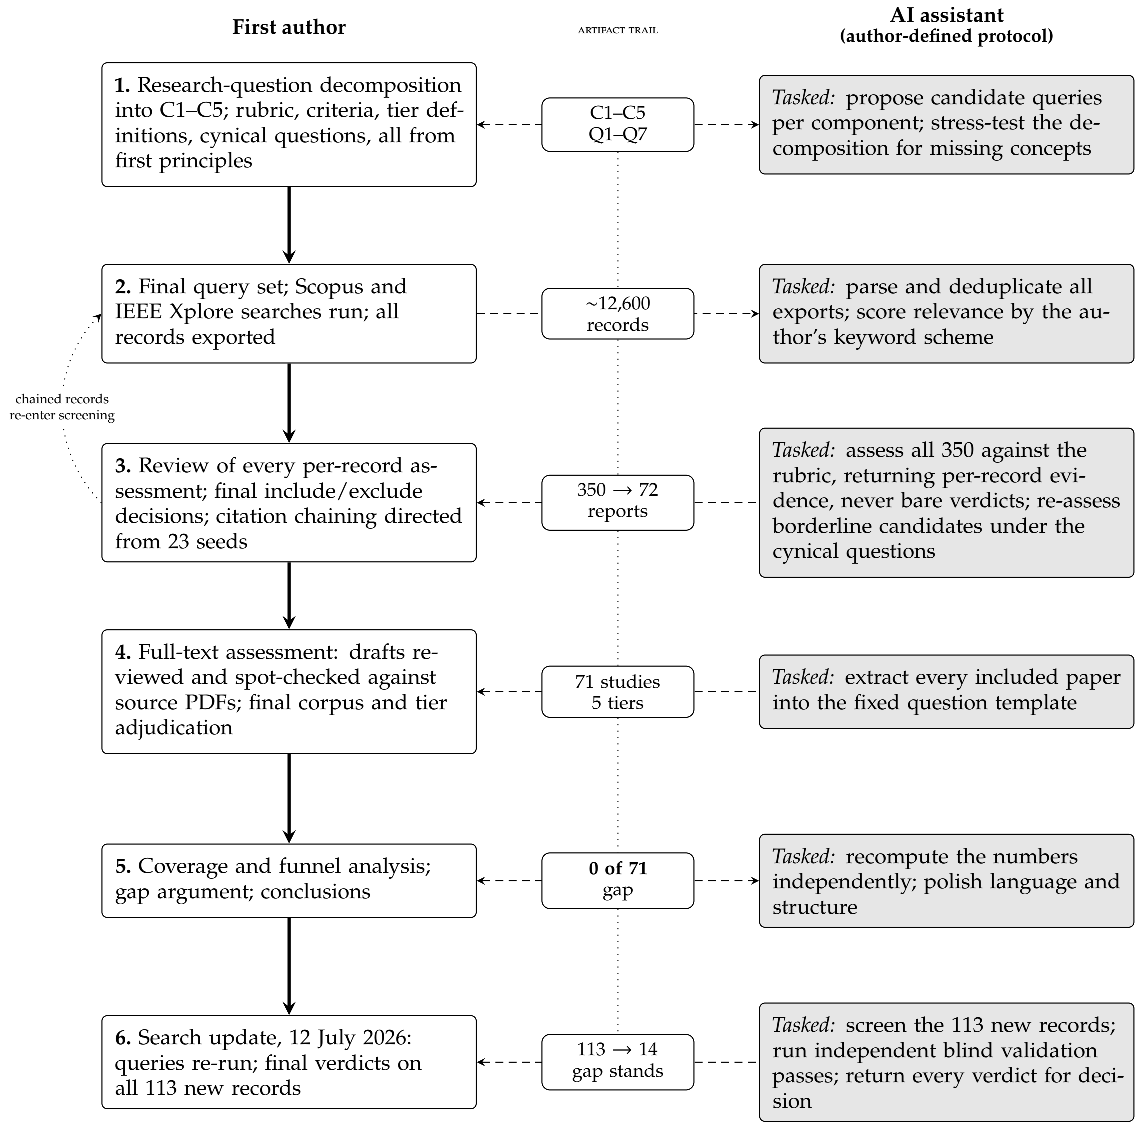

Pattern: the question is decomposed until each part is queryable; execution is delegated at scale with evidence returned at every step; every gate is adjudicated by the author; conclusions are then attacked before acceptance.

Figure 2.1: Division of labour between the first author (left lane, white) and the AI assistant (right lane, grey) at each stage of the review, including the July 2026 search update. The review's state advances only along the author lane (solid spine); dashed edges are hand-offs, each crossing a named artifact (centre column) whose per-record logs are preserved in the review's audit trail, with arrowheads giving the hand-off direction. The dotted arc marks iteration (citation chaining re-enters the record set), the double-headed dashed hand-offs at stages 1 and 5 mark exchanges that alternated between author and assistant until stable, and the dotted centre thread traces the record funnel. The grey boxes have no edges between them: the assistant holds no state across stages and makes no corpus, tier, or synthesis decision. Each is worded as the structured instruction under which the assistant was tasked.

3

Results

This chapter presents the quantitative outcomes of the systematic literature review. Section 3.1 describes the PRISMA-compliant search and selection process. Section 3.2 characterizes the included corpus by tier and relevance. Section 3.3 analyzes coverage of the five components. Section 3.4 traces the funnel from full corpus to the core intersection papers. Section 3.5 examines control law patterns among the most relevant papers. Section 3.6 summarizes key findings.

3.1. Search Results and Study Selection

The systematic search identified approximately 12,600 records across multiple sources. Table 3.1 presents the breakdown by source.

Table 3.1: Records identified by source with verification status

Source	Records	Query Files	Verified
Scopus (Q1–Q7)	1,248	7 .bib files	Yes
IEEE Xplore (primary)	3,302	7 .csv files	Yes
IEEE Xplore (supplemental)	576	5 .csv files	Yes
Forward citations (Scopus CitedBy)	7,150	24 .bib files	Yes
Backward citations (reference lists)	~330	6 .md files	Yes
Total identified	**~12,600**		

The PRISMA flow diagram (Figure 3.1) illustrates the complete selection pipeline. After removing approximately 3,066 duplicates, 9,534 unique records remained. An automated keyword-based relevance filter excluded 9,184 records that matched fewer than two RQ component keyword clusters (or one cluster without a contact-task term). This yielded 350 candidates for title/abstract screening.

Title/abstract screening applied eight exclusion criteria (E1–E8, defined in Table 2.4), removing 278 papers.

The remaining 72 papers underwent full-text review. One paper was excluded due to inaccessibility (Chinese-language publication without institutional access), yielding a final corpus of 71 studies.

3.2. Characteristics of Included Studies

The 71 included studies were classified into five tiers based on their relevance to the research question and methodological role. Table 3.2 presents the tier distribution.

A single DIAMOND paper addresses all five components; the bulk of the corpus is GOLD and SILVER work that aligns with the research question only in part, providing strong or supporting evidence on individual components. Two aspects of Table 3.2 are worth stating explicitly. The five RIVAL papers are

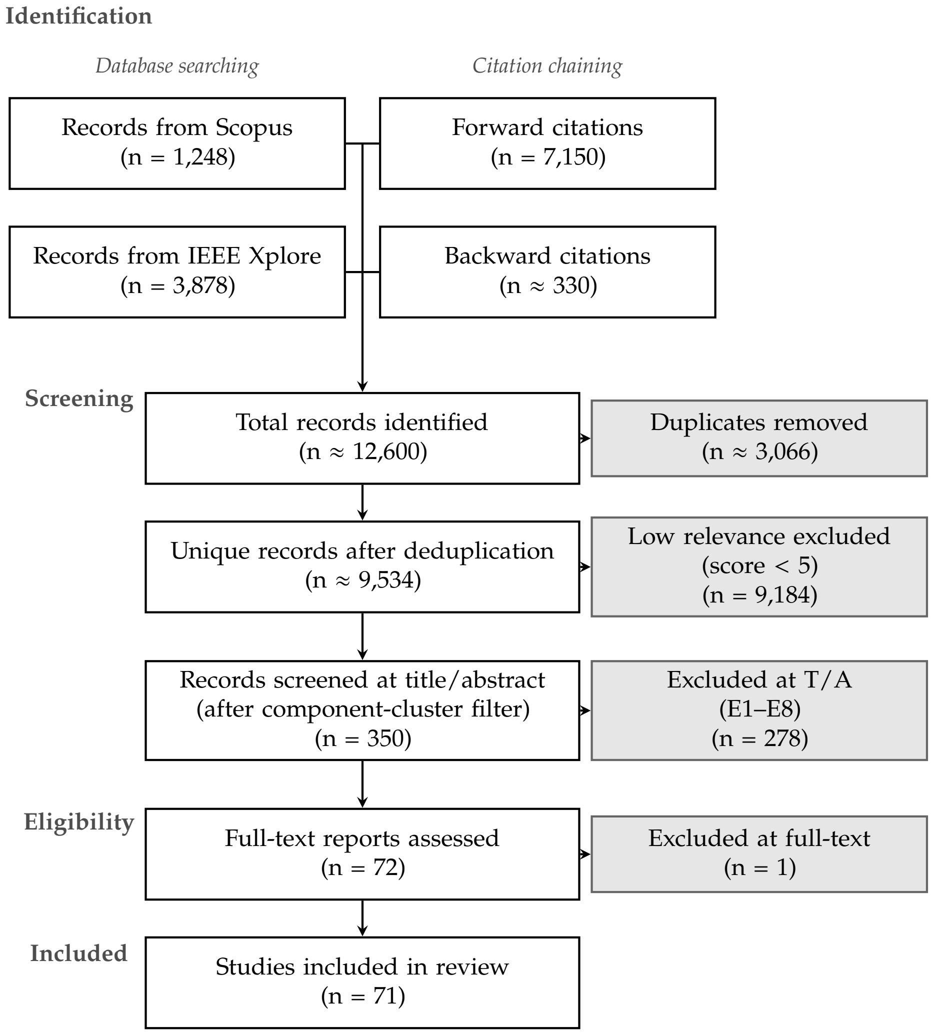

Figure 3.1: PRISMA 2020 flow diagram showing the systematic selection process from 12,600 identified records to 71 included studies. A dated search update run on 12 July 2026 (Section 2.7) screened 113 newly indexed records against the same rubric, with all 14 sought reports retrieved and assessed in full text; it did not alter the corpus or the conclusions.

Table 3.2: Paper tier distribution

Tier	Description	Count	Percentage
DIAMOND	Addresses all 5 components	1	1%
GOLD	Strong relevance to the research question	27	38%
SILVER	Moderate relevance	31	44%
RIVAL	Alternative paradigms	5	7%
FOUNDATIONAL	Pre-2015 seminal works	7	10%
Total		**71**	**100%**

not uniform: two use vision-based surface estimation rather than force-motion-based methods, and the remaining three address pHRI or impedance control without surface estimation. The FOUNDATIONAL papers are pre-2015 works retained, despite the date restriction, for the theoretical foundations they establish.

Table 3.3 shows the distribution of relevance ratings assigned during full-text assessment (a 1–10 scale, distinct from the automated keyword filter of the identification stage).

Table 3.3: Relevance score distribution

Score Range	Count	Percentage
7–10/10 (Highly relevant)	5	7%
5–6/10 (Moderately relevant)	23	32%
3–4/10 (Marginally relevant)	30	42%
1–2/10 (Low relevance)	13	18%
Total	**71**	**100%***

*Row percentages sum to 99% due to rounding.

The mean relevance score is 4.1/10 and the median is 4/10, indicating that the majority of included papers address only a subset of the components. Only 5 of 71 papers (7%) scored 7 or higher, reflecting the specificity of the research question combining five distinct technical requirements.

3.3. Component Coverage Analysis

Each paper was assessed for coverage of the five components using a three-level scale: ABSENT (0), IMPLICIT (1), or EXPLICIT (2). Table 3.4 presents the coverage distribution.

Table 3.4: Component coverage across the 71-paper corpus

Component	ABSENT	IMPLICIT	EXPLICIT	Total Addressed
C1: Surface Estimation	46 (65%)	17 (24%)	8 (11%)	25 (35%)
C2: Power-Assist/pHRI	22 (31%)	7 (10%)	42 (59%)	49 (69%)
C3: Shared Control	27 (38%)	19 (27%)	25 (35%)	44 (62%)
C4: Anisotropic Force	33 (46%)	21 (30%)	17 (24%)	38 (54%)
C5: Ergonomics	41 (58%)	21 (30%)	9 (13%)	30 (42%)

Figure 3.2 visualizes the component coverage distribution, showing both IMPLICIT and EXPLICIT contributions for each component.

Coverage is markedly uneven across the five components (Table 3.4). Power-assist and pHRI (C2) is the best served, addressed by 49 of 71 papers (69%), which reflects the mature state of physical human-robot interaction research; haptic shared control (C3) follows closely and often appears alongside it in collaborative manipulation studies. Anisotropic force control (C4) occupies the middle of the range, and ergonomic load reduction (C5) is more often mentioned than explicitly addressed, with the explicit treatments using validation instruments that range from objective EMG measurement to subjective workload assessment.

Online surface estimation (C1) is the least frequently addressed component and, more tellingly, the least explicitly addressed: only 8 papers (11% of the full corpus) treat it explicitly. This is the most significant coverage gap relative to the research question, which requires online surface geometry estimation from force-motion interaction.

3.4. Component Combination Analysis

The research question requires simultaneous consideration of multiple components: online surface estimation (C1), physical human-robot interaction (C2), and anisotropic force control (C4) form the

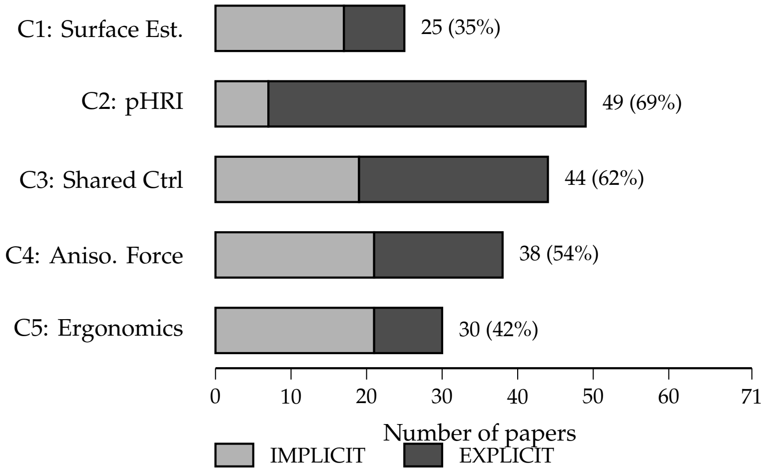

Figure 3.2: Component coverage visualization showing IMPLICIT and EXPLICIT contributions. C2 (pHRI) has highest coverage; C1 (Surface Estimation) has lowest, representing the primary gap.

core technical requirements. This section traces how the corpus narrows when requiring multiple component combinations.

The funnel from full corpus to the core intersection proceeds as follows:

1. **71 papers** in the reviewed corpus (100%)
2. **49 papers** address C2: Power-Assist/pHRI at IMPLICIT or higher (69%)
3. **16 papers** additionally address C4: Anisotropic Force Control (C2 ∩ C4, 23%)
4. **7 papers** additionally address C1: Online Surface Estimation (C1 ∩ C2 ∩ C4, 10%)
5. **3 papers** address all five components C1–C5 (4%)

Figure 3.3 visualizes this progressive narrowing as a funnel diagram, illustrating how the corpus reduces when requiring multiple core components.

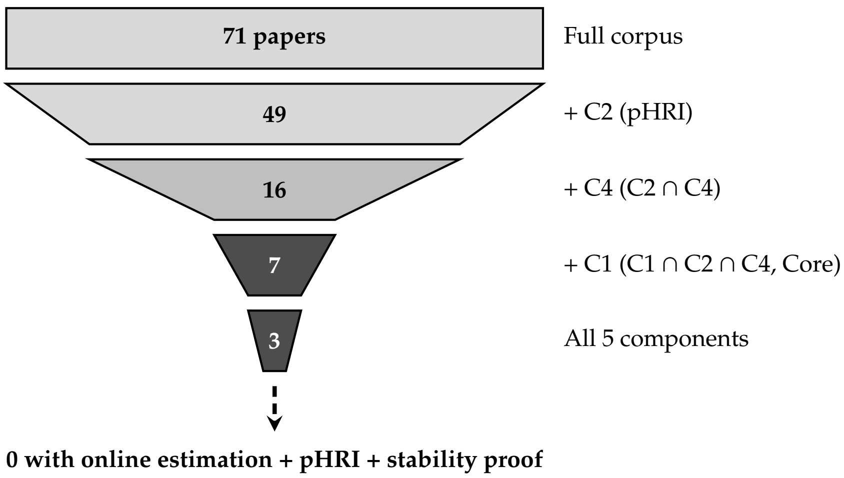

Figure 3.3: Research gap funnel: progressive narrowing from full corpus (71 papers) to the core intersection. The funnel terminates at zero papers combining online force-motion surface estimation, physical human-robot interaction, and formal stability analysis (stability analysis was coded for the 28 DIAMOND and GOLD papers).

The seven papers addressing the core intersection (C1+C2+C4) are listed in Table 3.5 with their component scores and key limitations.

Three of these seven papers also address all five components at IMPLICIT or higher: Paper 2 [30], Paper 287 [31], and Paper 164 [35]. However, each has significant limitations relative to the research

Table 3.5: Papers addressing the core component intersection (C1+C2+C4)

Paper	Citation	Year	C1	C2	C4	Key Limitation
2	[30]	2025	IMP	EXP	EXP	Surface normal assumed aligned with Z-axis
287	[31]	2015	IMP	EXP†	EXP	Instantaneous normal only ($\hat{\mathbf{n}} = -\mathbf{F}_T/\|\mathbf{F}_T\|$); joystick input
43	[32]	2015	IMP	IMP	EXP	Human not physically coupled to robot
74	[33]	2025	IMP	IMP	IMP	Fully autonomous, no pHRI in experiments
186	[34]	2021	EXP*	EXP	IMP	*Surface pre-scanned, not estimated online
10	[15]	1992	IMP	IMP†	EXP	Joystick teleoperation, no force feedback
164	[35]	2022	IMP	EXP†	IMP	Vision-based, not force-motion estimation

IMP = IMPLICIT, EXP = EXPLICIT. *Paper 186 uses pre-scanned mesh geometry.

†C2 for Papers 287, 10, and 164 reflects real-time human input via a remote interface (joystick or haptic device), not direct mechanical coupling; under the strict physical-coupling reading of C2, these papers do not satisfy it.

question:

- Paper 2 [30] achieves the highest overall relevance (9/10) and validates ergonomic benefit via EMG (43.2% assist efficiency, ASE), but assumes the surface normal aligns with the world Z-axis rather than estimating it online.
- Paper 287 [31] estimates surface normal from force direction ($\hat{\mathbf{n}} = -\mathbf{F}_T/\|\mathbf{F}_T\|$, with $\mathbf{F}_T$ the measured contact force at the tool tip), but this instantaneous estimate lacks geometric accumulation or friction compensation.
- Paper 164 [35] validates ergonomic benefit via the NASA Task Load Index (NASA-TLX), but uses vision-based (point cloud) surface estimation rather than force-motion-based estimation during contact.

Although seven papers score at IMPLICIT or higher on all three core components, none satisfies the stricter reading of C1 (online estimation from force-motion interaction) and C2 (direct mechanical coupling) simultaneously, as the limitation column shows: no paper in the reviewed corpus combines genuinely online force-motion surface estimation with direct physical human-robot coupling and anisotropic force control in a single validated system.

3.5. Control Architecture Distribution

Detailed control law analysis was performed on the 28 highest-relevance papers (DIAMOND and GOLD tiers), representing 39% of the corpus. This section reports the distribution of control approaches, reference frame usage, and stability analysis methods.

Control Law Types

Table 3.6 presents the distribution of control architectures among the 28 analyzed papers.

Admittance control dominates, accounting for half of the analyzed papers (14 of 28), which reflects its suitability for human-robot interaction where the robot responds compliantly to external forces. Impedance control is the main alternative, and the remainder divide evenly between hybrid force/position control and specialized approaches, the latter including variable stiffness and learning-based methods (Table 3.6).

Table 3.6: Control law type distribution (28 papers)

Control Type	Count	Percentage
Admittance-based	14	50%
Impedance-based	8	29%
Hybrid Force/Position	3	11%
Specialized/Other	3	11%
Total	**28**	**101%***

*Percentages sum to 101% due to rounding.

Reference Frame and Stability Analysis

Table 3.7 summarizes reference frame usage and stability analysis approaches.

Table 3.7: Reference frame usage and stability analysis (28 papers)

Category	Count	Percentage
Reference Frame:		
World/Task frame only	17	61%
Tool frame included	8	29%
Adaptive/Estimated frame	4	14%
Stability Analysis:		
Formal proof provided	16	57%
No formal proof	12	43%

Reference frame categories are not mutually exclusive.

The majority of papers (17 of 28, 61%) operate exclusively in world or task reference frames, assuming known and fixed coordinate systems. Only 4 of 28 papers (14%) employ adaptive or estimated reference frames that could accommodate unknown surface geometry:

- Karayiannidis and Doulgeri [36] use a fixed estimated normal $\hat{\mathbf{n}}$ that is not updated online during task execution.
- Luo et al. [37] derive a local frame from trajectory covariance decomposition, applicable to motion planning rather than surface contact.
- Nasiri and Wang [38] perform online surface normal estimation with friction compensation, but operate autonomously without physical human-robot interaction.
- Lu and Wen [31] estimate surface normal from force direction ($\hat{\mathbf{n}} = -\mathbf{F}_T/||\mathbf{F}_T||$) and include pHRI, but provide no formal stability analysis.

Regarding stability analysis, 16 of 28 papers (57%) provide formal proofs: 8 use Lyapunov methods, 3 use passivity-based analysis, 3 use frequency-domain methods (Routh–Hurwitz or Nyquist), and 2 use implicit approaches (linear-quadratic regulator, LQR, or positive definiteness arguments). The remaining 12 papers (43%) do not provide formal stability proofs.

No paper among the 28 analyzed combines all three of: (1) online surface normal estimation from force-motion, (2) physical human-robot interaction, and (3) formal stability proof.

3.6. Summary

The systematic search identified approximately 12,600 records, from which 71 studies were included following PRISMA guidelines. Key quantitative findings are:

1. **Component coverage is uneven**: C2 (Power-Assist/pHRI) is addressed by 49 of 71 papers (69%), while C1 (Online Surface Estimation) is addressed by only 25 of 71 papers (35%). Only 8 papers (11%) explicitly address surface estimation.

2. **The core intersection is sparsely populated**: Only 7 of 71 papers (10%) address all three core components (C1+C2+C4) at IMPLICIT or higher. Only 3 of 71 papers (4%) address all five components.
3. **Adaptive reference frames are rare**: Only 4 of 28 analyzed papers (14%) use adaptive or estimated reference frames. Of these, only Paper 287 combines surface normal estimation with pHRI.
4. **Stability analysis is incomplete**: Among the 28 DIAMOND and GOLD papers analyzed, 12 (43%) provide no formal stability proof.
5. **Admittance control dominates**: Admittance-based approaches appear in 14 of 28 analyzed papers (50%), consistent with human-robot interaction requirements.
6. **The full integration gap**: No paper in the corpus combines genuinely online force-motion surface estimation with direct physical human-robot coupling and anisotropic force control in a single validated system; among the 28 papers analyzed for control laws, none provides formal stability analysis for such a combination either.

Having characterized the corpus quantitatively, Chapter 4 provides detailed thematic analysis of how each component is addressed in the literature, followed by a discussion of implications in Chapter 5.

4
Literature Synthesis

This chapter synthesizes the thematic findings from the 71-paper corpus. While Chapter 3 presented quantitative coverage statistics, this chapter critically analyzes *how* each component is addressed in the literature, identifies methodological limitations, and evaluates the evidence supporting different approaches. Section 4.1 provides an overview. Sections 4.2–4.6 analyze each component in turn. Section 4.7 compares control law architectures. Section 4.8 examines cross-cutting integration challenges. Section 4.9 presents the research gap analysis. Section 4.10 concludes with a summary.

4.1. Introduction

The research question decomposes into five technical components: online surface estimation from force-motion interaction (C1), physical human-robot interaction with power-assist (C2), shared control with variable autonomy (C3), anisotropic force control with normal-tangent decomposition (C4), and ergonomic load reduction with quantitative validation (C5). As established in Chapter 3, no paper in the reviewed corpus addresses all five components with genuine online surface estimation.

A more nuanced picture emerges than simple coverage statistics suggest. The literature contains substantial work on each individual component, but the approaches are fundamentally incompatible with the research question's requirements in specific ways. Surface estimation methods assume autonomous operation. Human-robot interaction methods assume known geometry. Anisotropic control methods use fixed reference frames. Ergonomic validation methods test on flat surfaces. Understanding *why* these approaches fail to combine is essential for defining what closing the gap requires.

The synthesis is organized by component, with each section structured as follows: classification of existing approaches, critical evaluation of their limitations, identification of the specific gap, and relevance to the research question. The analysis draws on extraction data from all 71 papers, with detailed control law comparison for the 28 DIAMOND and GOLD papers.

4.2. Surface Estimation Methods

Of the 71 papers, 25 address surface estimation at IMPLICIT or EXPLICIT level (35% of corpus). However, the phrase "surface estimation" encompasses fundamentally different technical approaches with different applicability to the research question. This section classifies these approaches and identifies why none satisfies the requirement for online geometry estimation from force-motion interaction combined with human coupling.

Force-Based Geometry Estimation

Seven papers perform some form of surface normal estimation from force or motion measurements during contact; an eighth foundational work applies contact force control with the normal known a priori. This is the closest category to the review's requirement. Notation is unified across papers in this chapter; symbols may therefore differ from the original publications.

Doulgeri and Karayiannidis [39] present the most mathematically rigorous approach, with an adaptive law for surface normal estimation:

$$\dot{\hat{\mathbf{n}}}_{cd}(t) = -\gamma \hat{\mathbf{Q}}_{cd}(t)\,(\dot{\mathbf{p}}_t + \alpha \Delta \mathbf{p}_t) \tag{4.1}$$

where γ is a positive adaptation gain, $\hat{\mathbf{Q}}_{cd} = \mathbf{I} - \hat{\mathbf{n}}_{cd}\hat{\mathbf{n}}_{cd}^T$ is the tangent projector onto the estimated surface, $\dot{\mathbf{p}}_t$ is the tip velocity, α is a positive control gain, and $\Delta \mathbf{p}_t = \mathbf{p}_t - \mathbf{p}_{td}$ is the tip position error relative to the desired tangent-plane position. This approach estimates the normal from motion variables rather than force, which could simplify sensor requirements but requires compliant contact. The critical limitation is that this method is validated only in simulation with no physical human-robot interaction: the robot operates autonomously.

Nasiri and Wang [38] address a key practical challenge: friction corrupts force-based normal estimation. Their friction-compensated estimator is:

$$\hat{\mathbf{n}}_{surf} = \frac{\hat{\mathbf{f}}_n}{||\hat{\mathbf{f}}_n||}, \quad \hat{\mathbf{f}}_n = \mathbf{f}_s - \mathbf{f}_\tau \tag{4.2}$$

where $\mathbf{f}_s$ is the measured sensor force, $\hat{\mathbf{f}}_n$ is the estimated normal force component, and the friction force $\mathbf{f}_\tau = -\bar{\mu}||\mathbf{f}_\perp||\hat{\mathbf{v}}/||\hat{\mathbf{v}}||$ (with $\bar{\mu}$ the weighted moving-average friction coefficient, $\mathbf{f}_\perp$ the force component perpendicular to the velocity direction, and $\hat{\mathbf{v}}$ the estimated end-effector velocity) is subtracted from the measured sensor force. This explicit friction compensation improves accuracy on real surfaces, but the estimation algorithm operates at 1 kHz autonomously with no human coupling.

Lu and Wen [31] provide the only paper combining instantaneous normal estimation with human input:

$$\hat{\mathbf{n}} = -\frac{\mathbf{F}_T}{||\mathbf{F}_T||} \tag{4.3}$$

where $\hat{\mathbf{n}}$ is the estimated surface normal and $\mathbf{F}_T$ is the measured contact force vector at the tool tip. Lu and Wen's approach (estimating normal as $\hat{\mathbf{n}} = -\mathbf{F}/||\mathbf{F}||$) reveals the field's premature claim to "surface estimation." What they provide is surface *detection*: a single-sample direction that changes every control cycle. True estimation would accumulate observations into a geometric model that filters noise, rejects transients, and predicts normals at positions not yet contacted. The gap between detection and estimation is not semantic; it determines whether a system can plan contact trajectories or merely react to them. Any friction force corrupts this estimate; any noise in the force estimate, which Lu and Wen obtain from joint torque sensors and themselves describe as noisy, propagates directly to the normal direction. This review adopts a strict definition: C1 requires accumulated geometric representation, not instantaneous force projection.

Kazerooni [14] established the foundational approach for autonomous deburring, estimating contact angle from force ratios in 2D:

$$\tan(\alpha) = \frac{\mu t_y - t_x}{\mu t_x + t_y} \tag{4.4}$$

where α is the angle between the surface normal and the Y-axis, μ is the Coulomb friction coefficient, and t_x, t_y are the measured force components in the x and y directions at the contact point. This 1988 work demonstrates the principle but is limited to 2D contours with autonomous operation. Earlier foundational contributions include Stepien et al. [13], who applied tool/workpiece contact force control to robotic deburring with the surface normal known a priori, and Houshangi [16], who addressed position/force control over an unknown surface using hybrid strategies. Both operate autonomously without human coupling.

Yin et al. [40] and Jiang et al. [41] further extended force-based surface tracking to curved workpieces, though both track a planar contour: the intersection of the unknown surface with a pre-specified constraint plane. These methods demonstrate the feasibility of force-based surface following but remain fully autonomous.

Across these papers the estimation method is described in detail but its accuracy is almost never measured. Table 4.1 compares what each of the eight C1-EXPLICIT papers estimates, from which signals, and with what reported error. Only two report a quantitative error in an estimated quantity, and neither is an angular surface-normal error obtained on hardware: Doulgeri and Karayiannidis report convergence from a 24.5° initial offset in simulation, and Houshangi reports contact-point location errors from static measurements that were never closed in a control loop. The remaining papers report force or position tracking instead, which confounds estimator accuracy with controller and plant behavior and cannot substitute for it.

Parameter Estimation (Not Geometry)

Seven papers estimate environment parameters such as stiffness or contact position, but do *not* estimate surface geometry. This distinction is critical for the research question.

Tarokh and Bailey [17] use adaptive fuzzy control to handle unknown environment parameters, an early approach to contact under uncertainty. Iwasaki et al. [19] apply sliding mode control with a gain-scheduled variable hyperplane, adapting to unknown stiffness without estimating geometry.

Jung et al. [43] estimate environment stiffness algebraically:

$$\hat{k}_e = \frac{f_e}{x - x_e} \tag{4.5}$$

where $\hat{k}_e$ is the estimated environment stiffness, f_e is the measured contact force, x is the current end-effector position, and x_e is the environment surface position. This approach assumes x_e is approximately known and the contact is along a single axis. Surface orientation is not estimated. Jung2001 already includes both simulation and PUMA 560 hardware validation of this stiffness-estimation approach; Jung2004 [44] addresses the same unknown-environment force-tracking problem with a different technique, an adaptive gain that explicitly does not require environment-stiffness knowledge.

Azlan et al. [45] use function approximation techniques to estimate scalar location x_e and stiffness k_e online, but the formulation is inherently one-dimensional with no geometric content. Zhang et al. [33] employ reinforcement learning to adapt admittance parameters under unknown stiffness, but "unknown environment" refers to unknown stiffness on a *known flat surface*, not unknown geometry. Liang et al. [46] observe environment stiffness and location in real time from the measured contact force and robot position using iterative learning, again without estimating the normal direction.

These parameter estimation methods address robustness to environment uncertainty, which is valuable, but do not address the requirement of estimating surface *orientation* that varies along the contact trajectory.

Fixed or Assumed Surface Normal

The most relevant paper to the research question, Liu et al. [30] (DIAMOND tier), uses an Extended State Observer to estimate lumped disturbances including environment dynamics. However, the surface normal direction is assumed, not estimated: the paper performs no explicit online estimation of surface geometry (neither normal direction nor curvature), only lumped disturbance estimation, and the surface normal is assumed to be aligned with the world Z-axis in the experiments.

The ESO compensates for unknown environment *dynamics* (stiffness, damping) but the force control *direction* is fixed a priori. This approach works for surfaces where the normal is known or constant (horizontal polishing), but fails for curved or inclined surfaces where normal direction varies along the trajectory.

Karayiannidis and Doulgeri [36] assume a fixed estimated normal $\hat{\mathbf{n}}$ with bounded estimation error φ, but do not update the estimate online. Al Mashagbeh and Khamesee [32] assume the normal is aligned with the robot Z-axis and adapt position reactively, not the frame direction.

Pre-Scanned and Vision-Based Approaches

Three papers achieve surface awareness through means other than force-motion estimation: Kana et al. (SILVER tier) via a pre-scanned mesh, and Pruks and Ryu (RIVAL tier, journal and conference versions)

Table 4.1: Estimation accuracy reported by the eight papers scoring EXPLICIT on C1. Only two report a quantitative error in an estimated quantity, and neither is an angular surface-normal error measured on hardware: Doulgeri and Karayiannidis report the normal-vector error converging from a 24.5° initial offset in simulation, and Houshangi reports contact-point location errors from static measurements that were never closed in a control loop. The remaining six report only downstream task metrics, so across the entire C1-EXPLICIT set the accuracy of the estimate itself is unquantified.

Paper	Estimated quantity	Signals	Estimation accuracy	Convergence	Geometry tested	Platform
Kazerooni [14]	Normal angle α of a 2D contour	f_x, f_y	Not reported[a]	Not reported (algebraic)	Planar 2D contours (part edges)	Hardware, autonomous
Houshangi [16]	Contact-point location and contact type (point / line / plane)	$\mathbf{f}$, $\mathbf{m}$	Contact-point location error 0.9–4.6% (point), 0.17 cm / 5.3% (line)[b]	Not applicable (static, single-shot)	Flat table only	Hardware, autonomous, open loop[b]
Yin et al. [40]	Instantaneous normal and tangent at contact	$\mathbf{f}$, $\mathbf{m}$, μ known	Not reported[c]	∼5 s network learning delay	Curved bulb, tracked as a planar contour	Hardware, autonomous
Doulgeri & Karayiannidis [39]	Surface normal $\hat{\mathbf{n}}_{cd}$ and tangent projector	$\dot{\mathbf{p}}_t$, $\Delta\mathbf{p}_t$ (no force)	Normal error converges to zero from 24.5° initial offset	$<$ 1 s (force and position error)	Curved (circle, r = 0.239 m), soft tip	Simulation only
Jiang et al. [41]	Surface normal and local curvature	$\mathbf{f}$, $\mathbf{m}$	Not reported; results are qualitative only[d]	Not reported	Curved bulb, planar-contour trajectory	Hardware, autonomous
Kana et al. [34]	None estimated online; proxy point and normal on a pre-scanned mesh	Offline depth scan; robot / proxy position	Not applicable[e]	Not reported	Non-convex free-form blade	Hardware, human in loop (n = 1)
Nasiri & Wang [38]	Friction-compensated normal $\hat{\mathbf{n}}_{surf}$ and friction coefficient	$\mathbf{f}_s$, $\hat{\mathbf{v}}$	Not reported[f]	Not reported	Flat, dome, sinusoidal	Hardware and simulation, autonomous
Li et al. [42]	Surface contour indirectly, via a force-error-driven reference trajectory	Force error e_f only	Not reported[g]	Proved global, no time reported	Curved thin-walled shell and blade	Hardware, autonomous

Notes. A cell reads "not reported" only where the paper reports no error in the estimated quantity itself; several of these papers report task-performance metrics instead, listed below. Such metrics confound estimator accuracy with controller and plant behavior and cannot substitute for it. [a] Provides an analytical sensitivity of the angle estimate to friction perturbation ($\delta\alpha \approx 0.57°$ for $\delta\mu = 0.01$), but the experiments report contact forces, tool velocity, material removal rate and frequency response; the estimated angle is never compared against a measured ground truth. [b] The contact-point calculation is validated on static measurements only and is never closed in a control loop; the method additionally requires the manipulated object's shape to be known a priori. [c] Reports force error of 10 ± 0.35 N against a 10 N target (sensor error ±0.2 N); the estimated tangential-slope curve is shown to correspond to the bulb profile qualitatively, with no quantitative estimation error. [d] The paper states only that "a good experimental result is obtained" and provides no quantitative results, so it is not evidence of estimator performance either way. [e] Geometry is supplied by a pre-scan (mesh of 29,540 vertices and 50,118 faces), so there is no online estimate to score; the reported metric is proxy-to-robot position error at three stiffness levels. [f] Reports positional tracking error of 0.3–1 mm and up to ∼5% better positional accuracy than a non-updated normal. [g] Reports force-error MIAE of 1.5–1.8 N, position error within ±0.7–3 mm, and surface roughness of 0.649 µm against a 1.039 µm baseline. The paper's orientation-tracking error is a closed-loop tracking result, not an estimation error, and is therefore not reported as accuracy here.

via depth-camera point clouds. The remaining three RIVAL papers address pHRI control without surface estimation.

Kana et al. [34] provide an explicit statement of the pre-scanning assumption:

> "The starting point is the geometry of the surface, which we assume known and available in a discretized format, e.g. through scanning technologies."

Their proxy-based framework with geodesic trajectory generation is sophisticated, but it requires the surface mesh to be available before task execution. This fails the review's requirement for *online* estimation during contact.

Pruks and Ryu [35] use depth camera point clouds for surface estimation, achieving ergonomic benefit (validated via NASA-TLX). However, vision-based estimation is fundamentally different from force-motion-based estimation: it requires surface visibility, cannot work during occlusion (common in tool-handling), and cannot sense compliant or deformable surfaces accurately. This represents an alternative paradigm rather than a gap to be filled.

Gap Identification for C1

The surface estimation literature reveals a consistent pattern:

- Papers with force- or motion-based geometry estimation (Doulgeri, Nasiri, Kazerooni) operate autonomously without human coupling.
- Papers with human coupling (Liu, Kana) assume known or fixed surface geometry.
- The only exception (Lu and Wen) provides instantaneous estimation without geometric accumulation or friction compensation.

Closing the C1 gap requires: online surface normal estimation from force-motion interaction (following approaches like [38, 39]), integrated with physical human-robot coupling (following [30]), with filtering to build a local surface model rather than instantaneous-only estimates.

4.3. Physical Human-Robot Interaction Methods

Component C2 (Power-Assist/pHRI) has the highest coverage in the corpus, with 49 of 71 papers (69%) addressing it at IMPLICIT or EXPLICIT level. However, the majority address free-space co-manipulation rather than contact tasks with surface-following. This section analyzes the coupling methods, force rendering approaches, and their applicability to the research question.

Coupling Methods

Physical human-robot interaction requires some form of coupling between human and robot. The literature reveals three distinct coupling paradigms with different implications for contact tasks.

Direct physical coupling (hands-on-payload) appears in eight papers where human and robot share physical contact with the same object or tool. However, analysis reveals these papers address fundamentally different interaction regimes, none of which combines direct coupling with online surface geometry estimation:

- Liu et al. [30] achieve 43.2% assist efficiency (ASE) on an sEMG-derived metric, but assume the surface normal is fixed along the Z-axis, so the coupling handles known geometry only.
- Haninger et al. [47] focus on high-payload stability and collision safety, not surface geometry uncertainty.
- Mujica et al. [48] address unknown payload inertia during co-manipulation, which is orthogonal to surface geometry estimation.
- Grafakos et al. [49] use EMG for intent detection in free-space tasks without contact surface considerations.

Additional studies addressing physical coupling in pHRI include Bae et al. [50], who use virtual stiffness guidance for variable admittance; Franceschi et al. [51], who apply cooperative differential game theory for arbitration; Niaz et al. [52], who develop a learning-based adaptive admittance controller for contact-rich manufacturing; and Reyes-Uquillas and Hsiao [53], who achieve safe manual guidance through adaptive admittance. None of these combines direct coupling with online surface geometry estimation.

Direct physical coupling is essential for the target application because it ensures the human feels contact forces through the shared tool and can respond naturally to surface variations. Bidirectional force transmission creates a coupled dynamical system where human compliance and robot compliance interact, a fundamentally different control problem from teleoperation where the human receives only rendered feedback. The research question requires this coupling *combined with* online surface estimation, which none of the eight papers provides.

Teleoperation (master-slave) appears in four papers where the human operates a separate input device. Lu and Wen [31] use a joystick with no force feedback. Pruks and Ryu [35, 54] use haptic devices with virtual fixture feedback. Han et al. [55] use a Novint Falcon with bilateral force feedback. Teleoperation decouples the human's proprioceptive feedback from actual contact forces: the human senses only rendered forces from the master device, not the true contact interaction. This eliminates the natural coupled dynamics between human arm compliance and contact stability, requiring the robot to autonomously regulate contact. This restricts applicability to surfaces with well-characterized stiffness or requires high-bandwidth bilateral control that introduces stability challenges.

EMG-based coupling appears in three papers where muscle activation serves as an input channel. Bednarczyk et al. [56] achieve 67.8% FIT accuracy mapping EMG to estimated human force. Grafakos et al. [49] use co-contraction as a damping switching signal. Li et al. [57] modulate force control gain based on biceps EMG. EMG-based approaches provide intent information beyond F/T measurement but require additional sensing hardware and calibration.

Force Rendering: Admittance vs. Impedance

Force rendering defines how the robot responds to human or environment forces. The dominant approaches in pHRI are admittance control (force input → velocity output) and impedance control (position input → force output).

Admittance control appears in 14 of 28 analyzed papers (50%), with the standard form:

$$\mathbf{M}_d \dot{\mathbf{V}}_{ref} + \mathbf{D}_d \mathbf{V}_{ref} = \mathbf{F}_h \tag{4.6}$$

where $\mathbf{M}_d$ and $\mathbf{D}_d$ are virtual mass and damping matrices, $\mathbf{V}_{ref}$ is the reference velocity output of the admittance filter, and $\mathbf{F}_h$ is the human-applied force. Admittance control is natural for pHRI because the human applies force and the robot yields compliantly.

Li et al. [58] extend admittance control with MPC-based parameter optimization and energy tank passivity:

$$\mathbf{M}_d(t)\ddot{\mathbf{x}}_d + \mathbf{D}_d(t)\dot{\mathbf{x}}_d + \mathbf{K}_d(t)\mathbf{x}_d = \mathbf{F}_h \tag{4.7}$$

where $\mathbf{x}_d$ is the desired position, $\mathbf{F}_h$ is the human-applied force, and the time-varying parameters $\mathbf{M}_d(t)$, $\mathbf{D}_d(t)$, $\mathbf{K}_d(t)$ (virtual mass, damping, and stiffness) are optimized online while passivity is guaranteed via an energy tank.

Impedance control appears in 8 of 28 papers (29%), typically on torque-controlled robots whose joints can be commanded directly:

$$\mathbf{M}_d \ddot{\mathbf{e}}_p + \mathbf{D}_d \dot{\mathbf{e}}_p + \mathbf{K}_d \mathbf{e}_p = \mathbf{F}_{ext} \tag{4.8}$$

where $\mathbf{e}_p$ is the position error between desired and actual end-effector pose, $\mathbf{F}_{ext}$ is the external force, and $\mathbf{M}_d$, $\mathbf{D}_d$, $\mathbf{K}_d$ are the virtual mass, damping, and stiffness matrices defining the target impedance. Varying the impedance parameters online can inject energy and destabilize the interaction if the variation is not constrained. Ficuciello et al. [59] address this with variable impedance on redundant manipulators, exploiting the null space for intuitive pHRI. Choi et al. [60] propose a multi-task energy-aware impedance controller, while their later work [61] achieves stability via control Lyapunov function-based MPC. Dong et al. [62] use adaptive fuzzy impedance with human motion intention estimation for modular manipulators. These impedance approaches target free-space or known-geometry interaction.

Transparency and Intent Detection

A key challenge in pHRI is achieving "transparency," making the robot feel light and responsive to human intent. Duchaine and Gosselin [63] introduce velocity-based variable impedance to improve transparency:

$$\mathbf{F} = m\ddot{\mathbf{p}} + c\dot{\mathbf{p}} - \alpha\dot{\mathbf{F}}\dot{\mathbf{p}} \tag{4.9}$$

where $\mathbf{F}$ is the measured force vector, m is the virtual mass, c is the virtual damping, $\mathbf{p}$ is the end-effector position (with $\dot{\mathbf{p}}$ and $\ddot{\mathbf{p}}$ the velocity and acceleration), $\dot{\mathbf{F}}$ is a diagonal matrix whose entries are the time derivatives of the corresponding force components, and α is a weighting factor. The term $-\alpha\dot{\mathbf{F}}\dot{\mathbf{p}}$ modulates effective damping based on the product of force rate and velocity. This elegant intuition has nonetheless shaped a generation of pHRI controllers without rigorous foundation. When force is increasing in the velocity direction (human accelerating), damping decreases; when decreasing, damping increases. Duchaine and Gosselin's approach works empirically, but without formal stability proof, it cannot be safely extended to contact tasks where environment dynamics interact with parameter variation. The field's willingness to adopt this heuristic demonstrates a troubling gap between pHRI research (where "it works" suffices) and force control research (where stability guarantees are mandatory). For contact tasks with human coupling, this gap is unacceptable: any parameter adaptation law must include stability analysis.

Sensorless force estimation methods extend pHRI to robots without wrist-mounted F/T sensors. Le and Kang [64] implement contact force estimation via an adaptive third-order sliding mode observer, achieving sensorless collaborative interaction. Chien et al. [65] compare observer-based force-sensorless approaches for impedance control. Ren et al. [66] use joint torque servo with active disturbance rejection, and Wang et al. [67] validate torque-feedback-based impedance control experimentally. These observer-based methods reduce hardware requirements but introduce estimation delay that may affect contact stability.

Dimeas and Aspragathos [68] compare admittance adaptation strategies and find that, in their high-stiffness experiments, damping-only adaptation required 534% more operator energy than inertia-only adaptation, with combined inertia–damping adaptation costing 120% more; the effort measure is transferred mechanical energy, not muscular activity. However, inertia adaptation requires higher bandwidth and can introduce instability at high frequencies.

Haninger et al. [47] identify a fundamental tradeoff at high payloads: adding feedback damping for contact stability increased sustained human effort by 16%, measured via an interaction-energy metric. Damping that improves collision safety thus increases sustained human effort during normal operation.

Passivity and Stability

Variable impedance pHRI raises stability concerns because parameters change during interaction. The literature provides two main approaches.

Energy tank passivity, used by Li et al. [58], Luo et al. [37], and Nini et al. [69], maintains a virtual energy reservoir:

$$T(z_t(t)) = \frac{1}{2}z_t(t)^2, \quad \underline{T} \le T(z_t(t)) \le \bar{T} \tag{4.10}$$

where $T(z_t)$ is the tank energy as a function of the tank state variable z_t, and $\underline{T}$ and $\bar{T}$ are the lower and upper energy bounds. Parameter changes are permitted only if sufficient energy is available in the tank, preventing energy injection that could cause instability.

Passivity filter conditions, used by Bednarczyk et al. [56], constrain parameter rates:

$$0 \le K_d - \dot{D}_d \tag{4.11}$$

$$0 \le 2D_d M_d^{-1}(K_d - \dot{D}_d) - \dot{K}_d + \ddot{D}_d \tag{4.12}$$

where K_d and D_d are the time-varying impedance stiffness and damping parameters, M_d is the (constant) virtual mass, and overdots denote time derivatives. These conditions guarantee passivity for time-varying impedance but restrict the allowable adaptation range.

Neither mechanism is universal: a substantial share of the analyzed papers rely on empirical tuning or intuitive arguments rather than formal stability analysis, quantified in Section 4.7.

Gap Identification for C2

The pHRI literature is mature for free-space co-manipulation but underdeveloped for contact tasks:

- Papers with direct physical coupling either address contact but assume known surface geometry (Liu, Haninger) or co-manipulate without environment contact at all (Mujica).
- Papers with surface normal estimation (Lu, Doulgeri) either lack physical coupling or provide only instantaneous estimates.
- No paper combines pHRI with online surface geometry estimation that updates the force control frame.

Closing the C2 gap requires: physical human-robot coupling where the operator handles the tool directly (following [30]), integrated with surface estimation to enable contact on unknown geometry.

4.4. Shared Control Approaches

Haptic shared control (C3) addresses how authority is allocated between human and robot during collaborative tasks. Of the 71 papers, 44 address shared control at IMPLICIT or EXPLICIT level (62%). The literature reveals diverse authority allocation mechanisms, but none adapts authority based on surface geometry uncertainty.

Authority Allocation Mechanisms

Authority allocation can be triggered by various signals. The literature reveals five distinct categories.

Trust-based allocation uses computed or reported trust levels to modulate robot stiffness. Liao et al. [70] employ a Dynamic Bayesian Network to estimate trust and map it to stiffness:

$$K_o^d(h_h, T_k^m) = \frac{\gamma_1(T_k^m + \gamma_6)}{1 + e^{-\gamma_2 h_h}} + \frac{\gamma_3(T_k^m + \gamma_7)}{1 + e^{-\gamma_4(T_k^m - \gamma_5)}} \tag{4.13}$$

where K_o^d is the desired stiffness of the object-level impedance (the subscript o denotes the co-manipulated object, not the robot itself), h_h is the human's applied force, T_k^m is the measured trust level estimated via the Dynamic Bayesian Network, and γ_1–γ_7 are model parameters. This achieves 76% NASA-TLX workload reduction compared to autonomous operation. However, trust is estimated from the operator's verbal feedback, not from geometry uncertainty.

Intention-based allocation detects human intent from force, velocity, or EMG signals. Hamad et al. [71] define a Human Intention Index:

$$K_{HII}(t) \in \begin{cases} (0,1] & \mathbf{v} \cdot \dot{\mathbf{f}}_H > 0 \text{ and } \mathbf{v} \cdot \mathbf{f}_H > 0 \text{ (accelerating)} \\ [-1,0) & \mathbf{v} \cdot \dot{\mathbf{f}}_H > 0 \text{ and } \mathbf{v} \cdot \mathbf{f}_H < 0 \text{ (decelerating)} \\ [-1,0) & \mathbf{v} \cdot \dot{\mathbf{f}}_H < 0 \text{ (decelerating)} \\ \approx 0 & \mathbf{v} \cdot \dot{\mathbf{f}}_H \approx 0 \text{ (no change)} \end{cases} \tag{4.14}$$

where $\mathbf{v}$ is the velocity of the manipulated object, $\mathbf{f}_H$ is the human-applied force, and $\dot{\mathbf{f}}_H$ is its time derivative. K_{HII} is the continuous output of a Mamdani fuzzy inference system over these rules, not a crisp assignment. This drives adaptive force scaling but does not consider surface geometry.

Cacace et al. [72] use a neural network classifier to detect four intention classes (accompany, opposite, coinciding, deviation) with 89.4% accuracy, switching between passive and active assistance modes accordingly.

Performance-based allocation adapts parameters based on task execution metrics. Cacace et al. [73] use distance to target for sigmoid parameter scheduling:

$$\mathscr{L}(e_x) = \left(\frac{v_{max} - v_{min}}{1 + e^{(-w_1/l)(e_x - l/w_2)}} \right) + v_{min} \tag{4.15}$$

where $\mathscr{L}(e_x)$ is the scheduled admittance parameter value as a function of the distance-to-target error e_x, v_{max} and v_{min} are the upper and lower parameter bounds, l is a characteristic task length, and w_1, w_2

are shape parameters controlling the sigmoid transition. This provides smooth transitions but requires known waypoints.

EMG-based allocation uses muscle activation for authority modulation. Li et al. [57] modulate force control gain based on normalized EMG:

$$k_p = \begin{cases} (k_{p,max} - k_{p,min}) \times (1 - e_{norm}) & e_{norm} \leq \sigma \\ k_{p,min} & e_{norm} > \sigma \end{cases} \tag{4.16}$$

where $k_{p,max}$ and $k_{p,min}$ are the upper and lower bounds on the proportional gain, $\sigma \in (0, 1)$ is a switching threshold, and $e_{norm} = e_{sma}/e_{mvc}$ is the normalized EMG signal, with e_{sma} the simple moving average of raw EMG and e_{mvc} the maximum voluntary contraction (MVC) reference. Higher muscle activation yields lower assistance gain, preventing the robot from fighting the human. However, EMG-based methods require additional sensing and per-user calibration.

Additional intention- and force-based allocation approaches include Li et al. [74], who combine behavioral intention recognition with switching impedance control; Dimeas et al. [75], who incorporate manipulator performance constraints into Cartesian admittance control; and Lima et al. [76], who implement assistive force control for collaborative human-robot transportation. Devie et al. [77] optimize co-manipulation control using force sensor stiffness with unknown environments.

Bitikofer et al. [78] systematically compare admittance control dynamic models for transparent free-motion pHRI, finding that the choice of virtual dynamics significantly affects naturalness, measured as deviation from a minimum-jerk trajectory in a single-subject study. Zhou and Song [79] adapt impedance control for vehicle driving with human operation under unstructured environments, demonstrating the breadth of pHRI applications beyond manufacturing.

Safety-based allocation modulates autonomy based on collision or energy constraints. Nini et al. [69] bound the inertia adaptation rate to preserve passivity of their energy tank, as one part of a broader framework that separately enforces ISO/TS 15066 power-and-force limits via a dedicated constraint:

$$\dot{m}_P = \frac{2\omega(T - T_{min})}{\dot{x}_{max}^2 \Delta t} \tag{4.17}$$

where $\dot{m}_P$ is the passivity-derived bound on the inertia adaptation rate (per admittance degree of freedom; the paper additionally caps this against an independent, designer-chosen threshold before use), ω is a weighting coefficient that distributes the tank's available energy across the admittance degrees of freedom, T and T_{min} are the current and minimum tank energy levels, $\dot{x}_{max}$ is the maximum end-effector velocity, and Δt is the control time step. This preserves passivity but does not adapt to task geometry.

Contact Tasks with Shared Control

Only four papers combine shared control with contact tasks, and critically, none allocates authority based on surface geometry confidence. This omission reflects a deeper structural gap: existing shared control methods assume the contact surface is known, so there is no geometry uncertainty against which to allocate authority.

- Liu et al. [30]: Human guides tangent trajectory while robot regulates normal force, but the normal direction is fixed along the Z-axis. Authority allocation responds to human intent (via force input), not to estimation uncertainty.
- Li et al. [57]: EMG-modulated force control for ultrasound imaging uses known anatomical surface shapes. The EMG signal modulates force gain, not geometric adaptation.
- Kana et al. [34]: Proxy-based co-manipulation operates on pre-scanned meshes where surface normals are computed offline. Shared control allocates between human guidance and constraint enforcement, but geometry is never uncertain.
- Lu and Wen [31]: Force-based normal detection provides instantaneous estimates ($\hat{\mathbf{n}} = -\mathbf{F}/\|\mathbf{F}\|$) but does not update authority based on estimate confidence or convergence state.

A system that estimates surface geometry online must address a question none of these papers considers: how much authority should the robot assert while the normal estimate is still converging? High authority risks applying force in the wrong direction; low authority leaves the human to regulate contact unaided. This authority-uncertainty coupling is absent from the shared control literature reviewed here.

Gap Identification for C3

The shared control literature reveals:

- Authority allocation is triggered by trust, intention, performance, EMG, or safety, but never by geometry uncertainty.
- No paper provides direction-dependent authority allocation (different assistance in normal vs. tangent directions).
- Contact-task shared control assumes known surface geometry in all cases.

Closing the C3 gap calls for: authority allocation that accounts for surface estimation confidence, with potentially different assistance levels in the estimated normal direction (where contact force is critical) versus tangent directions (where human guidance is primary).

4.5. Anisotropic Force Control

Anisotropic force control (C4) refers to direction-dependent behavior: regulating force in some directions while permitting motion in others. Of the 71 papers, 38 address anisotropic control at IMPLICIT or EXPLICIT level (54%). This section distinguishes genuine anisotropy from fixed-frame single-axis control.

Explicit Normal-Tangent Decomposition

Six papers discussed here implement explicit separation of control behavior between estimated or assumed surface normal and tangent directions.

Liu et al. [30] implement the clearest version of the anisotropic approach: distinct strategies for free-motion and contact situations, with the robot tracking the human's motion intention in the tangent plane while regulating contact force in the normal direction. The normal direction force control uses ESO-based disturbance rejection:

$$\Delta f_n = f_e - f_d = M_d \ddot{z}_c + D_d \dot{z}_c \tag{4.18}$$

where $\Delta f_n = f_e - f_d$ is the net force error in the normal direction (with f_d the desired contact force and f_e the measured contact force, combining the pHRI human force and the pREI environment force into a single sensed channel), M_d and D_d are the virtual mass and damping of the admittance model, and z_c is the compliant displacement along the normal axis. However, as noted in Section 4.2, the normal direction is assumed aligned with the world Z-axis.

Doulgeri and Karayiannidis [39] implement the most rigorous normal-tangent decomposition using the tangent projector:

$$\mathbf{Q} = \mathbf{I} - \hat{\mathbf{n}}\hat{\mathbf{n}}^T \tag{4.19}$$

Force control acts along $\hat{\mathbf{n}}$ while position control acts in the subspace $\mathbf{Qx}$. The control law is:

$$\mathbf{u} = \mathbf{J}^T \left(\hat{\mathbf{n}}_d(t) f_d - k_v \dot{\mathbf{p}} - \begin{bmatrix} \alpha k_v \hat{\mathbf{Q}}_{cd}(t) \Delta \mathbf{p}_t \\ k_\phi \mathbf{e}_\phi \end{bmatrix} \right) + \mathbf{g}(\mathbf{q}) \tag{4.20}$$

where $\mathbf{u}$ is the joint torque vector, $\mathbf{J}$ is the robot Jacobian, $\hat{\mathbf{n}}_d(t)$ is the estimated generalized normal direction, f_d is the desired contact force magnitude, k_v is the velocity damping gain, $\dot{\mathbf{p}}$ is the end-effector velocity, α and k_ϕ are positive control gains for position and orientation respectively, $\hat{\mathbf{Q}}_{cd}(t)$ is the estimated tangent space projector, $\Delta \mathbf{p}_t$ is the tip position error, $\mathbf{e}_\phi$ is the orientation error, and $\mathbf{g}(\mathbf{q})$ is the gravity compensation torque. This approach adapts the force frame online, but operates autonomously without human coupling.

Karayiannidis and Doulgeri [36] use parallel force/position control with a fixed estimated normal, providing Lyapunov stability analysis but no online normal update.

Jinno et al. [15] established the task frame approach for robotic grinding:

$$\Sigma_W = \{P_w; \mathbf{n}_w, \mathbf{o}_w, \mathbf{a}_w\} \tag{4.21}$$

where Σ_W is the tool coordinate system defined by the grinder orientation, P_w is the frame origin at the tool center, $\mathbf{n}_w$ (the X_w axis) is the grinding direction, $\mathbf{a}_w$ (the Z_w axis) is the press direction, and $\mathbf{o}_w$ completes the right-handed coordinate system. Force control regulates contact along $\mathbf{a}_w$ while the tangential trajectory is commanded by the operator's joystick. The orientation adapts reactively via force error but surface geometry is not explicitly estimated. Baeten et al. [18] extended the task frame formalism to hybrid vision/force servoing with shared control, demonstrating early integration of visual feedback with force regulation in a constrained-direction framework.

Lu and Wen [31] estimate normal from force direction (Equation 4.3) and apply a generalized damper along this direction, with motion permitted orthogonally. This is the only paper combining normal estimation with human input, but as noted, the estimation is instantaneous without geometric accumulation.

Virtual Fixture-Based Anisotropy

Four papers achieve direction-dependent behavior through virtual fixtures rather than physical surface contact.

Tebaldi et al. [80] analyze stability for VF-constrained admittance:

> "The tangential component F_τ is analyzed because of its influence in the proxy dynamics... the motion in directions orthogonal to the reference is restricted."

Bazzi et al. [81] formalize direction projection for rotational VFs:

> "A smaller value of c_v implies a stronger rejection of f_v and hence a higher motion resistance."

Different scaling coefficients for different directions create anisotropic behavior, but in virtual constraint space rather than physical contact geometry. Luo et al. [37] adapt the fixture from a learned trajectory distribution, deriving direction-dependent stiffness from the covariance ellipsoid of the demonstrations, and Bae et al. [50] apply virtual stiffness normal to a straight-line path while leaving the along-path direction unconstrained.

Virtual fixture approaches are conceptually relevant (direction-dependent compliance) but address a different problem: constraining motion to paths or regions rather than regulating contact force on physical surfaces.

Zurlo et al. [82] contribute to the sensing side by developing collision detection and contact point estimation using virtual joint torque sensing, enabling cobots to localize contact without dedicated sensors. Kanekiyo et al. [83] propose a multiple virtual dynamics-based control framework for contact force realization in unknown environments, decomposing the problem into cascaded (series-connected) virtual dynamics layers. Both operate autonomously without physical human coupling.

Fixed-Frame Single-Axis Control

A substantial portion of papers scored IMPLICIT for C4 use isotropic parameters or single-axis control that does not constitute genuine anisotropy. These approaches fundamentally cannot handle unknown curved surfaces because they lack the directional selectivity required when the normal direction varies along the trajectory.

The prevalence of isotropic impedance parameters across the pHRI literature reflects an unexamined assumption: that human-robot interaction occurs in free space where direction is irrelevant. Li et al. [74] exemplify this with identity matrices: $\mathbf{M}_d = 0.8\mathbf{I}$, $\mathbf{D}_d = 10.53\mathbf{I}$, $\mathbf{K}_d = 28.96\mathbf{I}$. This isotropic formulation achieves good free-space performance precisely because contact forces, if they occur, are treated as disturbances to reject uniformly. For contact tasks, this assumption inverts the priority. The target application requires direction-dependent behavior: stiff along the surface normal (to regulate contact force) and compliant tangentially (to follow human guidance). When the normal direction is unknown and estimated online, isotropic control cannot distinguish force regulation from motion assistance, and force applied in any direction receives the same compliant response. This review argues that anisotropy is not an enhancement to pHRI control; it is a prerequisite for contact tasks. Any framework that defaults to isotropic parameters has implicitly assumed away the contact problem.

Zhang et al. [84] achieve good force accuracy (±2 N fluctuation after convergence, reached in approximately 10 iterations) using sliding-mode iterative learning control for belt grinding, and do so on curved workpieces as well as flat ones (a spline-contour steel part). The approach succeeds not because the surface is flat but because the geometry is known in advance: the grinding trajectory is generated by offline programming from a CAD model, and the contact angle θ used in the force mapping is estimated offline rather than online. The relevant contrast with the research question is therefore prior geometric knowledge versus online estimation, not flat versus curved. Where no CAD model or offline angle estimate is available, the force direction must be recovered from the interaction itself.

Balci et al. [85] demonstrate force-controlled robotic surface finishing with variable tool center point, achieving consistent contact on curved workpieces by adjusting TCP rather than the control frame. Li et al. [42] address robot compliance control for grinding thin-walled parts with unknown surfaces, adapting both deformation and orientation, representing one of the few works addressing unknown surface geometry in an industrial finishing context.

Song et al. [86] reduce force overshoot from 41.64 N to 0.12 N using adaptive feedforward compensation, demonstrating state-of-the-art force tracking on a single Cartesian axis. However, their adaptive mechanism compensates for environment stiffness uncertainty, not normal direction uncertainty. The 0.12 N overshoot is achievable precisely because the force direction is assumed constant. Extending this to rotating frames would require re-deriving the feedforward compensation for time-varying projection directions.

These papers demonstrate that high-accuracy force control is achievable with fixed frames. The open challenge is maintaining comparable accuracy when the force frame must track an estimated surface normal that evolves during task execution.

Stability Analysis for Anisotropic Control

Stability analysis for anisotropic control is challenging because different control laws apply in different directions.

Liu et al. [30] provide Lyapunov stability analysis for the combined system, but under the assumption of fixed (known) normal direction.

Doulgeri and Karayiannidis [39] provide Lyapunov stability for adaptive normal estimation with autonomous contact, showing local convergence under the condition $k_v > 2\lambda_M \max(\alpha, \beta\gamma)$, with a further combined gain condition (their Eq. 46) delimiting the stability region.

No paper provides stability analysis for the combination of: adaptive reference frame + physical human-robot interaction + contact with unknown surface. This theoretical gap remains open.

Gap Identification for C4

The anisotropic control literature shows:

- Explicit normal-tangent decomposition exists (Doulgeri, Karayiannidis, Jinno) but operates autonomously.
- Papers with pHRI (Liu, Lu) either assume fixed normal or provide instantaneous-only estimation.
- Virtual fixture approaches achieve anisotropy but for path constraints, not physical surface contact.
- Stability proofs exist for fixed frames or autonomous adaptive frames, but not for adaptive frames with pHRI.

Closing the C4 gap requires: anisotropic force control where the normal/tangent decomposition follows the estimated surface normal (not fixed to a world axis), combined with physical human coupling and formal stability analysis.

4.6. Ergonomic Validation Methods

Ergonomic load reduction (C5) is the ultimate goal of assistive devices, but the literature reveals significant variability in validation rigor. Of the 71 papers, 30 address ergonomics at IMPLICIT or EXPLICIT level (42%), but only 9 papers (13%) provide explicit quantitative validation.

EMG-Based Objective Measurement

Seven papers directly measure muscular activation via electromyography, which provides the most direct evidence of load reduction.

Liu et al. [30] achieve the strongest result: 43.2% assist efficiency (ASE) measured via a Myo surface-EMG armband with n=8 subjects. The ASE metric is:

$$ASE = \frac{EMG_{unassisted} - EMG_{assisted}}{EMG_{unassisted}} \times 100\% \tag{4.22}$$

This is the only paper combining EMG-validated load reduction with contact force control, making it the primary reference for future integration work.

Cacace et al. [72] report a fatigue reduction from Table 3 (mean muscle activation 45.3 in passive mode to 24.5 in shared mode, ≈46%) with shared control versus passive mode, but with unspecified muscles and a single trained user performing repeated trials in this case study — not the n=10 cohort, which trained the paper's separate intention-recognition classifier.

Grafakos et al. [49] measure co-contraction index on wrist flexor/extensor pairs:

$$u_c = \min(u^i, u^j) \tag{4.23}$$

where u^i and u^j are the activation levels of an agonist-antagonist muscle pair. This achieves 38% energy reduction with EMG-based damping switching in the aware group (n=10 of 20 subjects).

Arnold and Lee [87] report approximately 10% reduction in effort (CI-normalized overall and maximum muscle activation) on ankle muscles with variable impedance, using proper normalization:

$$EMG_{norm} = \frac{EMG_{task}}{EMG_{MVC}} \times 100\% \tag{4.24}$$

Bednarczyk et al. [56] use EMG-to-force mapping with 67.8% FIT accuracy, but for control input rather than outcome measurement.

Subjective Workload Assessment

Four papers use standardized questionnaires, primarily NASA-TLX.

Liao et al. [70] report 76% NASA-TLX reduction (55.1 to 13.05) with trust-based shared control versus autonomous mode, the largest subjective improvement in the corpus.

Müller et al. [88] report only 9% NASA-TLX improvement with variable impedance plus observer, demonstrating that subjective and objective metrics can diverge: objective force measurements favored the variable impedance controller *without* the observer (lowest interaction forces), while participants subjectively perceived the observer-augmented variant as requiring less effort.

This discrepancy highlights that interaction force is not a reliable proxy for perceived effort.

Force as Ergonomic Proxy

Several papers use F/T sensor measurements as proxies for human effort, which is problematic.

Cacace et al. [73] report 48% force reduction (17.5 N to 9.08 N) with variable admittance, but this measures robot-rendered force, not muscular activation.

The Haninger et al. [47] result discussed in Section 4.3, where damping added for contact stability raised sustained human effort on an interaction-energy metric, is a case in point: proxy metrics can mislead about the ergonomic cost of safety-motivated tuning.

The pervasive use of F/T sensors as ergonomic proxies reflects the field's instrumentation bias: force sensors are standard equipment in robotics labs, while EMG requires expertise in biomechanics that most robotics researchers lack. The resulting methodological convenience has distorted the literature. F/T sensors measure net interaction force, not the effort required to produce that force. Co-contraction (simultaneous activation of antagonist muscles for stability) increases muscular effort without increasing net force. Formally, $EMG \propto F_{agonist} + F_{antagonist}$ while $F_{net} = F_{agonist} - F_{antagonist}$. Claims of "effort reduction" based solely on interaction force should be viewed with skepticism; they may reflect nothing more than the robot absorbing load that the human was never bearing. This review takes the position that EMG is the only acceptable measure of muscular load reduction.

Validation Limitations

Several papers claim ergonomic benefit without measuring muscular load:

- Wu et al. [89] report effort reduction from hardware experiments (n=8), but measure only the interaction-force integral, not muscular activation.
- Nini et al. [69] focus on safety constraints with no ergonomic assessment.

Sample sizes vary widely, from n=1 [57, 90] to n=45 [88]. Statistical power for detecting meaningful differences typically requires n≥8, achieved by only a few studies.

Gap Identification for C5

The ergonomic validation literature reveals:

- EMG-based validation is rare: only 7 of 71 papers use EMG, and only Paper 2 uses EMG as an outcome measure alongside contact force control (Li et al. [57] use EMG as a control input only).
- Force-based proxies can mislead, as demonstrated by [47, 88].
- No paper validates ergonomic benefit on unknown curved surfaces requiring online normal estimation.
- No contact-task EMG validation is paired with explicit online surface-normal estimation.

Closing the C5 gap requires: EMG-based validation (following [30] methodology), on contact tasks with unknown curved surfaces, with appropriate sample size (n≥8) for statistical power.

4.7. Control Law Architectures

Detailed control law analysis was performed on the 28 DIAMOND and GOLD papers. This section synthesizes patterns and identifies architectural implications for systems addressing the gap.

Admittance Dominance

Admittance control dominates with 14 of 28 papers (50%), versus 8 for impedance (29%) and 3 for hybrid force/position (11%), reflecting the natural match with pHRI noted in Section 4.3.

In the standard admittance form (Equation 4.6), lower $\mathbf{M}_d$ and $\mathbf{D}_d$ yield faster, more transparent response. Liu et al. [30] use $\mathbf{D}_d = \text{diag}[150, 150, 300]$ Ns/m, with higher damping in the force-controlled normal direction.

A gap-closing system should adopt admittance-based control for the tangent (human-guided) directions, following the established paradigm.

Adaptation Mechanisms

Adaptation mechanisms across the 28 papers include:

- None (fixed gains): 6 papers
- Classical adaptive (FAT, ILC): 5 papers
- EMG-based: 3 papers
- Learning-based (RL, DL): 4 papers
- Gain-scheduled/switching: 5 papers

- Observer-based (ESO, ADRC): 3 papers

Observer-based approaches (Liu [30], Ren [66]) are particularly relevant because they estimate environment dynamics without requiring explicit model knowledge. A gap-closing system should consider ESO or similar observer-based disturbance rejection for the force control loop.

Learning-based approaches (Zhang [33], Wu [89]) show promise but typically lack formal stability guarantees. For industrial deployment, stability proofs are essential.

Reference Frame Usage

Reference frame definition is critical for anisotropic control. Among 28 papers:

- World/task frame only: 17 papers (61%)
- Tool frame included: 8 papers (29%)
- Adaptive/estimated frame: 4 papers (14%)

Only four papers use adaptive reference frames, and of these:

- Karayiannidis [36]: Fixed estimated $\hat{\mathbf{n}}$, not updated online.
- Luo et al. [37]: Local frame from trajectory covariance, not surface geometry.
- Nasiri and Wang [38]: Online surface normal estimation, but autonomous.
- Lu and Wen [31]: Surface normal from force direction, with pHRI.

Closing the gap requires an adaptive reference frame that tracks the estimated surface normal, combining the estimation approach of [38, 39] with the pHRI framework of [30].

Stability Analysis Distribution

Of 28 papers, 16 (57%) provide formal stability proofs:

- Lyapunov methods: 8 papers
- Passivity-based: 3 papers
- Frequency-domain (Routh–Hurwitz): 3 papers
- Implicit (LQR, positive definiteness): 2 papers

The remaining 12 papers (43%) provide no formal stability analysis [31–33, 38, 49, 50, 52, 57, 63, 79, 81, 83]. Three of them explicitly acknowledge this gap: one admits gains were determined by "trial and error" [32], another defers stability analysis to "future work" [83], and a third states that, with humans in the loop, it is "difficult to mathematically prove the stability of the system" [63]. The practical consequence is significant: variable impedance control without stability guarantees is acceptable for laboratory demonstrations with careful tuning, but unacceptable for manufacturing environments where untrained operators interact with the system. This review takes the position that any controller proposed for pHRI contact tasks must include a stability proof as a first-class deliverable, not an afterthought.

No paper provides stability analysis for the combination of adaptive reference frame + pHRI + contact. This is the central theoretical gap identified by this review.

Table 4.2 summarizes the control architectures across the 28 DIAMOND and GOLD papers, highlighting the correlation between stability proof provision and key design choices.

A notable pattern emerges: of the 4 papers using adaptive reference frames, only 2 provide formal stability proofs (Karayiannidis [36] via Lyapunov analysis and Luo et al. [37] via energy tank passivity), but neither proof covers a surface-adaptive frame. Karayiannidis uses a fixed estimated normal and Luo adapts via trajectory covariance, not surface geometry. The 2 papers with genuine surface-adaptive frames (Nasiri [38] and Lu and Wen [31]) both lack formal stability proofs. This suggests that surface-adaptive frames introduce analytical complexity that authors either cannot or choose not to address. Any attempt to close the gap must confront this challenge directly.

Table 4.2: Control law architecture summary for DIAMOND and GOLD papers

Characteristic	With Stability Proof	Without Proof	Total
Control Type			
Admittance	5	9	14
Impedance	8	0	8
Hybrid F/P	1	2	3
Other	2	1	3
Reference Frame			
World/task only	11	6	17
Tool frame included	4	4	8
Adaptive/estimated	2	2	4
Total	**16 (57%)**	**12 (43%)**	**28**

Reference-frame categories are not mutually exclusive: one paper (Luo et al. [37]) appears in both the tool-frame and adaptive-frame rows, so that block lists 29 entries across 28 papers.

Force Accuracy

Force tracking accuracy is reported quantitatively by only 8 papers, including:

- Liu et al. [30]: MAE 0.09 N, MaxE 0.36 N (best in corpus)
- Azlan et al. [45]: <0.1 N max error (simulation)
- Song et al. [86]: 0.12 N overshoot
- Zhang et al. [84]: ±2 N around 20 N target
- Zhang et al. [33]: ±1 N steady-state

A gap-closing system should target force accuracy comparable to [30] (MAE <0.5 N) to demonstrate that online surface estimation does not degrade force regulation performance.

4.8. Component Integration Challenges

The synthesis reveals that the five components exist in the literature but have not been integrated. This section examines why integration is challenging and what closing the gap requires.

The Surface-pHRI Incompatibility

The fundamental challenge in combining surface estimation with pHRI is that these two capabilities make conflicting assumptions about the control architecture.

Surface estimation methods (Section 4.2) assume autonomous operation because:

1. Estimation algorithms require controlled motion to excite surface geometry. Doulgeri's adaptation law (Equation 4.1) requires motion in the tangent plane to update the normal estimate: $\dot{\hat{\mathbf{n}}}_{cd} \propto \dot{\mathbf{p}}_t$. If the robot cannot control its own motion, the estimation may not converge.
2. Human forces introduce uncertainty that complicates separating contact force from human force. The F/T sensor measures their sum: $\mathbf{F}_{sensor} = \mathbf{F}_{contact} + \mathbf{F}_{human}$. Autonomous systems set $\mathbf{F}_{human} = 0$; pHRI systems must decompose this sum, typically using the admittance model inversion, but this assumes known contact geometry.
3. Estimation convergence typically requires multiple contact points sampled systematically, which human guidance may not provide. Nasiri's friction compensation [38] requires velocity in specific directions to distinguish normal force from friction force.

Conversely, pHRI methods (Section 4.3) assume known geometry because:

1. Variable impedance adaptation is triggered by human intent signals (force rate, EMG, co-contraction), not geometry uncertainty. No adaptation law in the reviewed corpus increases damping when the normal estimate variance is high.

2. Control frame definition requires a known reference direction for force regulation. Liu's ESO compensates for environment dynamics uncertainty, but the force direction $\mathbf{n}$ is assumed constant throughout the task.
3. Stability analysis assumes known or bounded environment parameters. Passivity conditions like $0 \leq K_d - \dot{D}_d$ (Equation 4.11) are derived for impedance parameters varying in a fixed frame, not for the frame itself rotating.

Bridging this gap requires three innovations not present in the current literature:

1. **Human-compatible estimation**: Surface estimation that converges using motion patterns typical of human-guided tasks (smooth, goal-directed trajectories) rather than systematic probing. This may require longer convergence windows or adaptive sampling that exploits natural human motion variability.
2. **Force decomposition**: A method to separate human force from contact force using the estimated surface geometry. If $\hat{\mathbf{n}}$ is the estimated normal, contact force should be $\mathbf{F}_{contact} \approx (\mathbf{F}_{sensor} \cdot \hat{\mathbf{n}})\hat{\mathbf{n}}$ plus friction, while human force acts primarily tangentially. This decomposition is circular (requires $\hat{\mathbf{n}}$ to compute, but $\hat{\mathbf{n}}$ depends on correct force decomposition), requiring iterative or coupled estimation.
3. **Uncertainty-aware authority**: Authority allocation that reduces robot force control authority when normal estimate confidence is low, preventing the robot from asserting force in potentially wrong directions.

The Frame Adaptation Problem

Anisotropic control requires a reference frame aligned with the surface. When this frame is estimated online, three coupled challenges emerge that no existing paper addresses simultaneously:

Force setpoint discontinuity: The force setpoint direction rotates as the normal estimate $\hat{\mathbf{n}}(t)$ updates. If the estimate changes rapidly (e.g., at surface curvature transitions), the desired force vector $f_d\hat{\mathbf{n}}(t)$ rotates, potentially causing force spikes or contact loss. Consider a surface where the true normal changes by 30° over 10 cm of motion: if the estimate tracks this change with delay, the force error $\mathbf{f}_d - \mathbf{f}_{actual}$ includes both magnitude error and direction error. Fixed-frame analyses like Liu's [30] bound only magnitude error.

Time-varying frame stability: Standard Lyapunov stability analysis for impedance control assumes a fixed coordinate frame. When the frame rotates, additional terms appear in the Lyapunov derivative:

$$\dot{V} = \dot{V}_{fixed} + \underbrace{\frac{\partial V}{\partial \hat{\mathbf{n}}}\dot{\hat{\mathbf{n}}}}_{\text{frame rotation term}} \tag{4.25}$$

The frame rotation term can inject or extract energy depending on the alignment between estimation error and control action. Doulgeri and Karayiannidis [39] address this for autonomous contact: their adaptation law preserves the unit norm of the estimate and they prove local asymptotic stability, which suggests the frame rotation term is manageable when the adaptation gain is bounded, although no explicit bound on $||\dot{\hat{\mathbf{n}}}||$ is derived. With pHRI, the human can induce arbitrary motions that may violate this bound.

Human adaptation to changing robot behavior: As the robot's force control frame rotates, the human experiences changing resistance directions. In tangent directions (human-guided), resistance should be low; in the normal direction (force-controlled), resistance reflects the desired contact force. If these directions shift unexpectedly, the human may perceive the robot as "fighting" their intent, increasing co-contraction (muscular effort to stabilize against uncertainty) and defeating the ergonomic purpose. No paper in the corpus studies human adaptation to robots with time-varying anisotropic compliance.

Doulgeri and Karayiannidis [39] provide the most relevant theoretical foundation, with the local stability result for autonomous adaptive-frame contact and its gain conditions stated in Section 4.5. Extending this to pHRI requires:

- Modeling the human arm as an additional impedance element with its own (uncertain, time-varying) parameters.
- Bounding the frame rotation rate $||\dot{\hat{\mathbf{n}}}||$ even when human forces can induce arbitrary motion.
- Validating that the resulting frame changes do not degrade perceived transparency or ergonomic benefit.

The Validation Gap

Ergonomic validation (Section 4.6) is performed on simple geometries that do not exercise the online estimation capability:

- Liu et al. [30]: Tests on a sinusoidal height profile demonstrate that the system handles varying surface height, but force is regulated along a fixed vertical axis while the true surface normal varies with the local slope; the resulting misalignment is neither compensated nor quantified. The reported ASE results validate their controller on this simplified geometry class.
- Arnold and Lee [87]: Tests variable impedance on a wearable ankle robot in a free-space reaching task, with no surface contact.
- Other EMG studies address free-space tasks without surface contact (Grafakos [49], Cacace [72]), where surface geometry is not a factor.

No study validates ergonomic benefit on surfaces requiring online normal estimation. This creates an inferential gap: even if future work demonstrates technical feasibility (accurate estimation, stable control), the ergonomic benefit cannot be assumed to transfer from fixed-geometry validation. Specifically:

- **Estimation overhead may increase cognitive load**: If the human must monitor whether the robot has “learned” the surface correctly, mental demand increases even if physical demand decreases.
- **Frame uncertainty may increase co-contraction**: Unpredictable robot behavior (from estimation transients) may cause the human to stiffen their arm, increasing muscular effort.
- **Curved surfaces require different motion strategies**: Human guidance on curved surfaces involves continuous wrist reorientation, potentially loading different muscle groups than flat-surface tasks.

Ergonomic benefit must therefore be validated specifically on unknown curved surfaces with online estimation active, comparing against both unassisted operation and fixed-frame assistance to isolate the contribution of adaptive estimation.

Methodological Limitations of This Review

Before presenting the gap analysis, several limitations of this systematic review must be acknowledged:

Publication bias: The reviewed corpus consists of peer-reviewed publications, which disproportionately report positive results. Methods that failed to achieve stable contact or accurate estimation may be underrepresented, potentially inflating the apparent maturity of individual components.

Language bias: Only English-language publications were included. Significant work in robotics and control originates from non-English-speaking research communities (particularly Japan, Germany, and China), and relevant methods may exist in conference proceedings or technical reports not indexed in Scopus/IEEE.

Gray literature exclusion: Doctoral theses and technical reports were not systematically searched. These sources often contain more detailed experimental methodology and negative results than journal publications.

Temporal snapshot: The search was conducted in January 2026. The rapid pace of development in learning-based control and human-robot interaction means relevant work may have appeared since the search date.

Component decomposition subjectivity: The five-component framework (C1–C5) was designed for this review. Papers were scored against these components by a single human reviewer (with AI assistance under the author-defined protocol; see Section 2.8), and borderline cases (e.g., whether instantaneous normal estimation counts as C1=IMPLICIT or EXPLICIT) required judgment calls that another reviewer might resolve differently.

These limitations do not invalidate the gap analysis (the absence of papers combining all five components is robust to reviewer interpretation), but they caution against overgeneralizing specific component scores or claiming exhaustive coverage of the field.

4.9. Research Gap Analysis

The systematic review of 71 papers reveals a specific, well-defined research gap. This section quantifies the gap with evidence from the corpus.

The Core Integration Gap

The research question requires five components simultaneously:

1. **C1**: Online surface estimation from force-motion (not vision, not pre-scanned)
2. **C2**: Physical human-robot coupling (not teleoperation, not autonomous)
3. **C3**: Shared control with appropriate authority allocation
4. **C4**: Anisotropic force control (force along estimated normal, motion tangentially)
5. **C5**: Validated ergonomic load reduction

0 of 71 papers in the reviewed corpus (derived from systematic search of Scopus and IEEE Xplore in January 2026 using queries targeting surface contact co-manipulation, force control for surface finishing, online surface estimation, anisotropic compliance, and power-assist for contact tasks) combine all five components with genuine online surface estimation.

Table 4.3 shows how component combinations progressively narrow the corpus. While individual components are well-represented, their intersection reveals the gap. One intersection is exact by construction of the corpus rather than by error: C1+C4 equals the full C1 count (25 papers), because every study that estimates surface geometry online also differentiates normal from tangential control; in this literature, estimation is always in service of anisotropic force control.

Table 4.3: Component combination coverage across the 71-paper corpus

Component Combination	Papers	Percentage
C2 alone (pHRI)	49	69%
C2 + C3 (pHRI + shared control)	43	61%
C2 + C4 (pHRI + anisotropic)	16	23%
C1 + C4 (surface est. + anisotropic)	25	35%
C1 + C2 (surface est. + pHRI)	7	10%
C1 + C2 + C4 (core intersection)	7	10%
C1 + C2 + C4 + C5 (with ergonomics)	3	4%
All five with online C1	0	0%

The funnel from 71 papers to 0 reveals where integration fails: only 7 papers address both C1 and C2, and these are exactly the 7 core-intersection papers, since every paper combining C1 and C2 also addresses C4. In two of them the surface information is pre-scanned (Kana) or vision-based (Pruks and Ryu) rather than derived from force-motion interaction, and each of the 7 has a specific limitation that prevents full integration.

Table 4.4 presents the 7 core papers with their component scores and specific limitations. This comparison reveals that the gap is not merely “nobody tried”; each paper made deliberate design choices that preclude full integration.

The closest papers are:

- Liu et al. [30]: 4 of 5 components explicit, with C1 only implicit: assumes fixed Z-axis normal.
- Lu and Wen [31]: 3 of 5 components at EXPLICIT level, but C1 is instantaneous-only estimation.

Table 4.4: The 7 core papers addressing C1+C2+C4: component scores and key limitations

Paper	C1	C2	C3	C4	C5	Key Limitation
Liu et al. [30]	IMP	EXP	EXP	EXP	EXP	Surface normal assumed fixed (Z-axis)
Lu & Wen [31]	IMP	EXP‡	EXP	EXP	IMP	Instantaneous estimation only; no filtering
Kana et al. [34]	EXP†	EXP	EXP	IMP	ABS	Requires pre-scanned surface mesh
Al Mashagbeh [32]	IMP	IMP	IMP	EXP	ABS	Human decoupled; moves marker not robot
Zhang et al. [33]	IMP	IMP	ABS	IMP	ABS	Experiments are autonomous, not pHRI
Jinno et al. [15]	IMP	IMP‡	ABS	EXP	ABS	Joystick input; no force feedback
Pruks & Ryu [35]	IMP	EXP‡	EXP	IMP	EXP	Vision-based, not force-motion

Legend: ABS = Absent, IMP = Implicit, EXP = Explicit. †Kana's EXP score reflects surface being "known" via pre-scanning, which fails the review's requirement for online estimation from force-motion interaction. ‡C2 for Lu & Wen, Jinno et al., and Pruks & Ryu reflects real-time human input through a remote interface (joystick or haptic device) rather than direct mechanical coupling; under the strict physical-coupling reading used in the gap analysis, these papers do not satisfy C2.

- Kana et al. [34]: 3 of 5 components, but C1 uses pre-scanned surface.

Specific Capability Gaps

Force-motion surface estimation + pHRI: 0 of 71 papers in the systematic corpus with *accumulated* geometric estimation. Papers with force-based (Nasiri [38], Kazerooni [14]) or motion-based (Doulgeri [39]) estimation operate autonomously. Papers with pHRI (Liu [30], Mujica [48], Haninger [47]) assume known geometry. Lu and Wen [31] represent a partial exception: they estimate normal direction from instantaneous force ($\hat{\mathbf{n}} = -\mathbf{F}/\|\mathbf{F}\|$) during human-guided operation, but this estimate is not filtered, does not account for friction, and is not accumulated into a geometric model. Closing the gap requires estimation that builds a local surface representation over time.

Adaptive force frame + pHRI: 0 of 71 papers. Lu and Wen [31] use instantaneous normal estimation but do not provide stability analysis for the time-varying frame. Doulgeri and Karayiannidis [39] provide stability analysis for adaptive frames but operate autonomously. No paper combines both.

EMG validation on unknown curved surfaces: 0 of 71 papers. Liu et al. [30] validate EMG on a sinusoidal height profile but regulate force along a fixed Z-axis, leaving the variation of the true surface normal uncompensated. This validates their controller only where the fixed-axis assumption is tolerable, not on surfaces requiring normal estimation. Closing the gap requires validation on surfaces where normal direction varies along the trajectory, exercising the online estimation capability.

Stability proof for adaptive frame + pHRI: 0 of 71 papers. Of the 16 papers with formal stability proofs (Section 4.7), all assume either fixed reference frames (Liu, Jung, Bednarczyk) or autonomous operation (Doulgeri). The theoretical combination of time-varying estimated frame + human impedance coupling + contact dynamics remains unanalyzed.

The RIVAL Alternative

Five papers (RIVAL tier) approach the problem through alternative paradigms; only the two Pruks and Ryu papers perform vision-based surface estimation, while the remaining three are pHRI or impedance approaches without surface estimation. Huang et al. [91] employ vision-based adaptive sliding mode admittance control for collaborative robots. Zhang and Xia [92] use interconnection and damping assignment passivity-based control for a compliant assistive robot, and Maithani et al. [93] implement trust-based variable impedance for cooperative pHRI. These represent an alternative paradigm with different trade-offs:

This review focuses on the force-motion paradigm specifically because it enables contact tasks where vision is occluded (grinding under guards, polishing concave surfaces) and where surface compliance matters (deformable workpieces).

Table 4.5: Force-motion vs. vision-based surface estimation

Characteristic	Force-Motion	Vision-Based
Works during contact	Yes	Limited
Works in occlusion	Yes	No
Senses compliance	Yes	No
Requires visibility	No	Yes
Pre-contact planning	No	Yes
Existing with pHRI	No (gap)	Teleoperation only [35]

Requirements for Closing the Gap

Based on the gap analysis, closing the identified gap requires four elements:

1. A force-motion-based surface normal estimator that operates online during contact without requiring prior geometric knowledge, extending approaches from [38, 39] to the pHRI context.
2. An anisotropic admittance controller where the force control frame adapts to track the estimated surface normal, combining the adaptive normal-tangent decomposition of [39] with the pHRI framework of [30].
3. Stability analysis for the combined system (adaptive frame + pHRI + contact), extending the Lyapunov approach of [39] to include human dynamics.
4. EMG-validated ergonomic benefit on unknown curved surfaces, following the methodology of [30] but removing the fixed-normal assumption.

No existing paper provides these four elements in combination; together they define the research needed to close the gap.

4.10. Summary

This chapter synthesized findings from the 71-paper corpus across the five components.

Surface Estimation (C1): Force- and motion-based geometry estimation exists for autonomous systems (Doulgeri, Nasiri) but has not been integrated with physical human coupling. Papers with pHRI assume known or fixed surface geometry. The only exception (Lu and Wen) provides instantaneous-only estimation without geometric accumulation.

Physical Human-Robot Interaction (C2): Mature for free-space co-manipulation with diverse approaches (admittance, impedance, EMG-based). Contact tasks are underrepresented, and all assume known surface geometry.

Shared Control (C3): Authority allocation mechanisms are diverse (trust, intention, performance, EMG, safety) but none considers surface geometry uncertainty. Direction-dependent authority allocation along an estimated, rather than fixed, normal does not exist in the reviewed corpus.

Anisotropic Force Control (C4): Explicit normal-tangent decomposition exists (Doulgeri, Karayiannidis) but operates autonomously. Papers with pHRI use fixed reference frames. Stability analysis for adaptive frames with human coupling does not exist in the reviewed corpus.

Ergonomic Validation (C5): EMG-based validation is rare (7 of 71 papers). Only Liu et al. use EMG as an outcome measure alongside contact force control, and only on surfaces with an assumed fixed normal. Validation on unknown curved surfaces does not exist in the reviewed corpus.

Control Architectures: Admittance control dominates (50%), appropriate for pHRI. Adaptive reference frames are rare (4 of 28 papers). Formal stability proofs are missing in 43% of papers.

Research Gap: 0 of 71 papers combine online force-motion surface estimation + physical human coupling + anisotropic control + formal stability + ergonomic validation. This is the specific gap this review identifies.

Chapter 5 will discuss implications of these findings, integration with broader technical knowledge, and limitations of this review.

5 Discussion

The preceding chapter established that the five components exist individually in the literature but have not been integrated, quantified the gap as 0 of 71 papers combining all components with genuine online surface estimation, and identified three technical incompatibilities that explain why integration has not occurred. The focus now shifts from *what the literature shows* to *what it means*: Section 5.1 delivers a direct answer to the research question posed in Chapter 1, Section 5.2 interprets the findings beyond the purely technical perspective, Section 5.3 traces evidence-based implications for the design of systems that would close the gap, and Section 5.4 interrogates the review's own assumptions and limitations.

5.1. Answering the Research Question

Chapter 1 posed the question: *how can a control strategy for a mechanically coupled assistive device maintain desired contact force along the surface normal while providing power-assist in tangential directions, when surface geometry is unknown and must be estimated online from force-motion interaction, to reduce operator muscular load during tool-handling tasks?*

The answer from the systematic review of 71 papers is that the literature provides mature solutions for each clause of this question in isolation, but no integrated framework exists. Admittance-based power-assist is well-established (Section 4.3). Normal-tangent decomposition for anisotropic force control has been demonstrated with formal stability guarantees for fixed force frames (Section 4.5). Online surface normal estimation from force-motion interaction has been demonstrated on curved surfaces in autonomous settings (Section 4.2). Ergonomic validation has documented muscular load reductions of up to 43% via EMG [30] (Section 4.6). Each component works. None works together with the others in the configuration the research question demands.

The gap is narrow but real. Three papers came closest: Liu et al. [30] addressed 4 of 5 components explicitly (C1 only implicitly) with the strongest ergonomic validation in the corpus, but assumed the surface normal is fixed along the world Z-axis; Lu and Wen [31] combined instantaneous normal estimation with human-guided operation, but without geometric accumulation, friction compensation, or stability analysis; Kana et al. [34] demonstrated sophisticated surface-following with human co-manipulation, but required the surface mesh to be pre-scanned. The specific limitations of all 7 core papers are documented in Table 3.5; what matters for the present argument is that each limitation reflects a deliberate design choice, not an oversight. Liu chose a fixed normal because it enabled stability analysis. Lu chose instantaneous estimation because it avoided the convergence problem. Kana chose pre-scanned geometry because it enabled geodesic trajectory planning. The gap persists not because researchers have neglected the problem, but because solving one component's requirements actively complicates another's.

The gap is structural. It sits at the intersection of three research communities that publish in different venues and use different formalisms. Surface estimation researchers [38, 39] publish in autonomous manipulation venues and assume the robot has full control authority. Physical human-robot interaction researchers [63, 68] publish in human-robot collaboration venues and assume the contact environment is known or absent. Ergonomics researchers validate on simplified geometries where the control problem is already solved. The synthesis revealed this pattern implicitly: autonomous estimation papers and known-geometry pHRI papers coexist in the corpus, developing largely independently and publishing in different venues with incompatible assumptions. Applying the cynical assessment questions (Section 2.4) to the 28 DIAMOND and GOLD papers sharpens this observation: only 5 (18%) genuinely handle unknown surface geometry, while 16 (57%) have a human in the real-time control loop, but only Lu and Wen [31] appear in both categories, and their operator commands the robot through a joystick rather than through mechanical coupling, with surface estimation limited to instantaneous force-direction projection without geometric accumulation.

The research question therefore remains open. Closing it requires bridging these communities, not merely extending a single line of work.

5.2. Interpreting the Findings

Chapter 4 analyzed each component through a technical lens: what methods exist, how they compare, and where they fall short. The following subsections interpret those findings through broader perspectives.

Deployment Readiness

The finding that 12 of 28 DIAMOND and GOLD papers (43%) lack formal stability analysis (Section 4.7) is not merely an academic concern. Variable impedance controllers without stability guarantees cannot support a transferable safety argument under ISO 10218 and ISO/TS 15066 for collaborative robot deployment. Conformity with their power- and force-limiting requirements is demonstrated by risk assessment and measurement; an empirically tuned controller can pass, but its bounds hold only for the tested platform and configuration, so every retuning reopens the safety case. A controller validated through empirical tuning on a specific hardware platform provides no transferable safety argument. The gap between laboratory demonstration and deployable system is therefore wider than the literature suggests: not only are the five components not integrated, but nearly half the existing single-component controllers lack the formal analysis needed for industrial deployment.

The correlation identified in Table 4.2 reinforces this concern: of the 4 papers using adaptive reference frames, only 2 provide formal stability proofs, neither for a surface-adaptive frame. The 2 papers with genuine surface-adaptive frames (Nasiri [38] and Lu and Wen [31]) both lack proofs, while the 2 with proofs (Karayiannidis [36] and Luo [37]) use non-surface-adaptive formulations. The analytical complexity introduced by surface-adaptive frames discourages the stability analysis that deployment requires. Any system combining a surface-adaptive frame with pHRI coupling faces this challenge directly.

Ergonomic Evidence on Real Workpieces

The finding that 0 of 71 papers validate ergonomic benefit on unknown curved surfaces (Section 4.6) has implications beyond the technical gap. It means the occupational health case for assistive devices in surface finishing remains unsubstantiated for the majority of real manufacturing scenarios. Industrial surface finishing overwhelmingly involves curved or freeform workpieces: turbine blades, automotive body panels, cast components. Liu et al.'s [30] 43% reduction in muscular effort is the strongest result in the corpus, but it was validated on a surface with a fixed vertical normal. Whether this benefit persists when the robot must simultaneously estimate the surface and regulate anisotropic force is an empirical question that the literature cannot currently answer.

The question matters practically because the justification for deploying assistive robots in surface finishing rests on ergonomic benefit. If the benefit degrades or disappears when the system operates on geometrically uncertain surfaces, the deployment case collapses regardless of the controller's technical sophistication.

The Instrumentation Barrier

The prevalence of F/T sensor measurements as ergonomic proxies, criticized in Section 4.6, reflects a pattern deeper than methodological convenience. Robotics laboratories are equipped with force/torque sensors; electromyography requires expertise in biomechanics, electrode placement protocols, and signal processing that most robotics researchers lack. The resulting instrumentation asymmetry has created a structural barrier to ergonomic validation. Papers report "effort reduction" based on interaction force measurements that, as Haninger et al. [47] and Müller et al. [88] independently demonstrated, can diverge from subjectively perceived effort and task-level outcomes. The field has optimized for what its laboratories can measure, not for what the occupational health application requires.

Breaking this pattern requires either interdisciplinary collaboration (robotics teams partnering with biomechanics laboratories) or investment in EMG instrumentation and expertise within robotics groups. The latter approach is exemplified by Liu et al.'s [30] methodology of surface EMG measurement on upper-body muscle groups.

Why the Gap Persists

The synthesis in Chapter 4 documented the technical incompatibilities between surface estimation and pHRI (Section 4.8). Viewed across the full corpus, a circular dependency emerges that explains the persistence of the gap. Surface estimation researchers require autonomous operation for controlled excitation of the surface geometry. pHRI researchers require known geometry to design stable controllers. Ergonomics researchers require working integrated systems to validate benefit. No group can start without the others finishing first. The deadlock is not a coordination failure; it reflects genuine technical constraints. Doulgeri and Karayiannidis's [39] estimation law requires tangential motion to update the normal estimate, and human-guided motion may not provide the excitation patterns the estimator needs. Liu et al.'s [30] stability proof holds for a fixed force-control direction, and extending it to a rotating frame introduces additional terms in the Lyapunov derivative (Equation 4.25) that may violate the boundedness conditions. Each community's simplifying assumptions are not arbitrary; they are necessary for the results they achieve.

Breaking this circular dependency requires addressing all three constraints simultaneously, accepting that the resulting system will be more constrained than any single-component solution.

5.3. Implications for System Design

Chapter 4 identified four requirements for closing the gap (Section 4.9). Each design implication below is traced to specific evidence from the review, justifying why these particular choices follow from the gap analysis rather than from researcher preference.

Admittance Control as Base Architecture

The review found that 14 of 28 DIAMOND and GOLD papers (50%) use admittance control for pHRI tasks (Section 4.7). The dominance of admittance control is not arbitrary. Admittance control takes the measured human force as input and commands robot motion, the causality available on the non-backdrivable, position- or velocity-controlled industrial manipulators typical of contact applications; impedance control commands joint torque in response to measured motion, and renders compliance faithfully only on backdrivable, torque-controlled hardware. The architectural choice is therefore driven largely by the available hardware. Admittance control also carries a known caveat directly relevant to this application: in stiff contact, and with a stiffly gripping human, low virtual inertia and damping can destabilize the loop, which is why damping design recurs throughout the corpus. Liu et al.'s [30] directional damping ($D_d = \text{diag}[150, 150, 300]$ Ns/m, with higher damping along the force-controlled normal) demonstrates how admittance parameters can encode anisotropy directly. Admittance control for the tangential directions remains an evidence-based choice, grounded in the literature's convergence on this architecture and in the hardware realities of pHRI contact tasks.

Force-Motion Estimation over Vision

The RIVAL comparison (Table 4.5) established that vision-based and force-motion-based surface estimation serve fundamentally different operating conditions. Vision requires surface visibility and cannot operate during contact when the tool occludes the workpiece. Force-motion estimation works during contact and in occlusion, but requires physical interaction with the surface. For surface finishing tasks where the tool is in continuous contact and typically occludes the contact region, force-motion estimation is not merely an alternative to vision; it is the only viable sensing modality. A focus on force-motion estimation is therefore a response to the operational constraints of the target application, not a limitation of scope.

EMG as Validation Metric

The critique of force-based ergonomic proxies (Section 4.6) established that interaction force reduction does not reliably indicate muscular load reduction. Co-contraction, where antagonist muscles activate simultaneously for joint stability, increases effort without increasing net force. Müller et al.'s [88] finding that objective force measurements and subjective perceived effort diverged, and Haninger et al.'s [47] finding that safety-enhancing damping increased sustained human effort by 16% on an interaction-energy metric, demonstrate that force-derived proxies can mislead in both directions. Any claim of ergonomic benefit must therefore be validated with EMG rather than repeating the field's methodological shortcut. The standard is not imposed externally; it follows directly from the review's own evidence that the alternative is unreliable.

Stability Proof as Deliverable

The finding that both papers with genuine surface-adaptive reference frames lack stability proofs (Section 4.7) establishes both the need and the difficulty. A new controller combining an adaptive frame, pHRI coupling, and contact dynamics introduces more potential instability mechanisms than any existing controller in the corpus. Doulgeri and Karayiannidis's [39] Lyapunov analysis provides a starting point, demonstrating local stability for autonomous contact with an adaptive frame under the condition $k_v > 2\lambda_M \max(\alpha, \beta\gamma)$, with the full stability region delimited by an additional combined gain condition. Extending this to include human dynamics as an additional impedance element with uncertain, time-varying parameters is the theoretical gap identified in Section 4.8. Omitting stability analysis would not merely repeat the field's existing gap; it would do so for a controller with strictly more complexity than those already lacking proofs.

What Would Constitute a Meaningful Demonstration

A future system need not exceed Liu et al.'s [30] 43% ASE to demonstrate a contribution. Liu's result was achieved with a fixed surface normal, the simplest possible geometric assumption. The meaningful demonstration is a system that works on unknown curved surfaces where the normal must be estimated online. Even a lower ergonomic benefit with online estimation active would demonstrate something no existing paper shows: that pHRI assistance transfers to geometrically uncertain contact tasks. The appropriate comparison is not against Liu's best-case result, but against unassisted operation on the same curved surface and against fixed-frame assistance on curved surfaces where the fixed-frame assumption is violated.

5.4. Limitations of This Review

Section 4.8 acknowledged five methodological limitations: publication bias, language bias, gray literature exclusion, temporal snapshot, and component scoring subjectivity. This section interrogates whether these limitations, or others not yet acknowledged, could undermine the central gap claim.

Methodological Limitations

The review searched Scopus and IEEE Xplore but not ergonomics-specific databases (e.g., PubMed, Ergonomics Abstracts). The coverage confidence analysis (Table 2.7) rated C5 (ergonomic load reduction) as Moderate, the lowest confidence of the five components. Papers combining surface estimation with ergonomic validation may exist in biomechanics or occupational health literature that the search strategy did not reach. This limitation is partially mitigated by the inclusion of forward and backward citation tracing from the 23 seed papers, which would capture highly cited cross-disciplinary work, but cannot guarantee coverage of recent publications in non-engineering venues.

The single-reviewer design introduces subjectivity in component scoring. No inter-rater reliability was computed because no second reviewer was available. Borderline scoring decisions, such as whether a paper's reactive force adaptation constitutes IMPLICIT surface estimation (C1) or merely robust force control, required judgment that another reviewer might resolve differently.

The English-language restriction excludes work from research communities where surface finishing robotics is active, particularly in Japan and Germany. Jinno et al.'s [15] foundational work on force-controlled grinding demonstrates that relevant methods originate from non-Anglophone research traditions, and current work in these communities may appear in domestic conferences or technical reports not indexed in the searched databases.

Could the Gap Be an Artifact?

The central claim of this review is that 0 of 71 papers combine all five components with genuine online surface estimation. This claim must be interrogated for two potential artifacts.

Framework circularity. The five-component decomposition (C1–C5) was designed for this review. A differently framed research question might reveal different gaps or none at all. If the question asked only about "adaptive force control for contact tasks" without requiring online surface estimation, the gap would narrow substantially: Liu et al. [30] already demonstrate this for a fixed, pre-assumed surface normal, and the remaining contribution would be incremental. The 5-component framework therefore shapes the gap it claims to find. This circularity is mitigated, but not eliminated, by two observations. First, the individual components are well-represented in the literature (C2 appears in 69% of papers, C4 in 54%), confirming that the framework captures real research activity rather than imposing alien categories. Second, the empty intersection is robust: even relaxing the C1 requirement from accumulated geometric estimation to any form of online normal awareness derived from force-motion interaction, and relaxing C5 from EMG to any ergonomic measure, the count remains 0. The gap is not an artifact of overly strict scoring.

Scoring sensitivity. Would a different reviewer score Lu and Wen's [31] surface estimation (C1) as EXPLICIT rather than IMPLICIT? If so, the gap narrative changes: one paper would address 4 of 5 components at EXPLICIT level. However, even under this generous scoring, Lu and Wen's estimation is limited to instantaneous force-direction projection ($\hat{\mathbf{n}} = -\mathbf{F}/||\mathbf{F}||$) without filtering, friction compensation, or geometric accumulation. The qualitative gap between this and a surface estimation capability that could support trajectory planning on curved surfaces remains regardless of how the score is recorded. Similarly, re-scoring Kana et al.'s [34] C1 as failing the online requirement (which the extraction record confirms) does not change the count because their pre-scanned approach fundamentally differs from force-motion estimation. The gap claim is robust to reasonable scoring variations at the boundaries.

The RIVAL Challenge

The classification of vision-based approaches as RIVAL alternatives rather than as solutions to the research question deserves scrutiny. If Pruks and Ryu [35] achieve validated ergonomic benefit with vision-based surface estimation, is force-motion estimation truly necessary? This review's scoping assumes that surface finishing involves tool-surface occlusion where the contact region is not visible. But what fraction of real manufacturing surface finishing involves truly occluded contact? Convex workpieces (automotive panels, turbine blade exteriors) may permit camera placement where the contact zone remains visible. The assumption that vision is generally occluded during surface finishing is plausible for concave surfaces and enclosed geometries, but is an empirical claim that this review cannot verify. If future work demonstrates that vision-based estimation suffices for the majority of industrial surface finishing scenarios, the force-motion approach would serve a narrower niche than this review implies.

Conversely, force-motion estimation offers capabilities that vision cannot: sensing surface compliance, operating on transparent or reflective surfaces, and functioning in environments where dust, coolant, or sparks degrade optical sensing. These advantages may prove more significant than occlusion in practice, but this too is an empirical question beyond the scope of a literature review.

What the Review Can Claim

Despite these limitations, the following claims are robust.

The absence of papers combining online force-motion surface estimation with pHRI and anisotropic force control is confirmed by systematic search of two major engineering databases using queries designed to capture each component (Section 2.2), supplemented by citation tracing. Even accounting for scoring subjectivity and database coverage limitations, no plausible re-scoring of the reviewed corpus would yield a paper that performs accumulated geometric estimation from force-motion interaction during physically coupled human-robot contact with direction-dependent force control and EMG-validated ergonomic benefit; the July 2026 search update (Section 2.7) likewise found none among 113 newly indexed records.

The technical incompatibilities identified in Section 4.8, particularly the surface-pHRI incompatibility and the frame adaptation problem, are structural findings derived from the mathematical formulations of the reviewed papers. These are not scoring artifacts; they follow from the equations. Surface estimation requires controlled excitation that human guidance may not provide. Stability analysis for fixed frames does not extend to rotating frames without additional analytical work. These constraints exist regardless of how papers are scored or how many additional papers might be found.

The identified gap is genuine. Whether it is the most important gap in the field, or whether an alternative approach would serve industry better, are questions this review raises but cannot answer.

Chapter 6 draws these threads together and turns them into recommendations addressed to the separate research communities whose boundaries this chapter has traced.

6 Conclusion

The following sections summarize the findings of the systematic literature review, formulate recommendations for the field, and reflect on the scope and limitations of the review's contribution.

6.1. Summary of Findings

The systematic review of 71 papers, drawn from Scopus and IEEE Xplore using queries targeting surface contact co-manipulation, force control for surface finishing, online surface estimation, anisotropic compliance, and power-assist for contact tasks, reveals that the literature provides individually mature solutions for every clause of the research question. Physical human-robot interaction is the most-covered component (49 of 71 papers, 69%), with admittance control the dominant architecture among the 28 papers analyzed for control laws (14 of 28). Shared control with diverse authority allocation mechanisms is well-represented (44 of 71, 62%). Anisotropic force control appears in 38 of 71 papers (54%), with formal stability guarantees demonstrated in individual fixed-frame cases. Ergonomic considerations appear in 30 of 71 papers (42%), though only 7 use EMG; the strongest documented result is a 43% muscular load reduction [30]. Surface estimation is the least-represented component (25 of 71, 35%), and its force-motion-based subset has produced mathematically rigorous approaches with proven convergence properties. Each component is technically mature.

The gap lies in integration. The component combination analysis (Table 4.3) shows a progressive narrowing: 49 papers address pHRI, 16 combine pHRI with anisotropic control, 7 address the core intersection of surface estimation plus pHRI plus anisotropic control, and 0 of 71 combine all five components with genuine online surface estimation from force-motion interaction. The 7 core papers that address the core intersection (C1+C2+C4) each fall short in a specific, documented way (Table 3.5). The three closest, Liu et al. [30], Lu and Wen [31], and Kana et al. [34], respectively assume a fixed surface normal, provide only instantaneous estimation without geometric accumulation, and require a pre-scanned surface mesh. Each limitation reflects a deliberate design choice that enabled the results the paper achieved, not an oversight that is trivially resolved.

The gap is structural. It sits at the intersection of three research communities with incompatible assumptions: surface estimation researchers assume autonomous operation for controlled excitation; pHRI researchers assume known geometry for stable controller design; ergonomics researchers validate on simplified geometries where the control problem is already solved. Re-examining the 28 DIAMOND and GOLD papers under the cynical assessment questions (Section 2.4) is consistent with this divide: the set of papers with genuine unknown-surface handling (5, or 18%) and the set with a human in the real-time control loop (16, or 57%) overlap in exactly 1 paper, and that paper (Lu and Wen [31]) couples its operator through a joystick rather than mechanically, with surface estimation limited to instantaneous force-direction projection; under a strict physical-coupling reading of C2, the overlap is empty. A circular dependency prevents incremental resolution: estimation requires controlled motion the human may not provide, stable control requires the geometry the estimator has not yet converged on, and ergonomic validation requires the integrated system that does not yet exist.

6.2. Recommendations

The following recommendations derive from specific findings of this review. Each is linked to the evidence that motivates it.

Research Opportunities

For surface estimation research: validate with human-in-loop disturbances. The review found that none of the 5 genuine surface estimation papers among the 28 DIAMOND and GOLD papers tests estimation under direct physical human-robot coupling (Section 4.2). The closest (Lu and Wen [31]) uses joystick-based guidance with only instantaneous force-direction estimation. This means the estimation algorithms have never been tested under the conditions where they would actually be used in assistive applications: with a human physically coupled to the robot, introducing unmodeled forces that corrupt the measurement signal. The field treats surface estimation as a perception problem; it is also an interaction problem. Future work should validate estimation convergence using human-guided motion patterns rather than systematic autonomous probing.

For pHRI research: move beyond flat and known surfaces. With one partial exception (Lu and Wen [31], who estimate normal direction from instantaneous force on a flat surface), contact-task pHRI papers among the 49 addressing C2 assume known or fixed surface geometry when surface contact is involved (Section 4.3). Controllers validated only on horizontal flat surfaces provide no evidence of performance on the curved and freeform workpieces that dominate real surface finishing. The admittance parameters, stability conditions, and transparency metrics reported for flat-surface tasks may not transfer to surfaces where the contact normal rotates along the trajectory. Experimental validation on curved surfaces with varying normal direction should become standard practice.

For ergonomic validation: adopt EMG as standard for muscular load claims. Only 7 of 71 papers use electromyography (Section 4.6). Force/torque sensor proxies can mislead: Haninger et al. [47] found that safety-enhancing damping increased sustained human effort by 16% on an interaction-energy metric, and Müller et al. [88] found that objective force measurements and subjective perceived effort diverged. The field's reliance on force proxies has left the occupational health case for assistive devices unsubstantiated in terms of actual muscular load. EMG measurement, following established protocols for electrode placement and MVC normalization, should be the minimum standard for claims of ergonomic benefit.

For stability analysis: require formal proofs for variable impedance pHRI controllers. Of the 28 DIAMOND and GOLD papers, 12 (43%) lack formal stability analysis (Section 4.7). Among the 4 papers using adaptive reference frames, only 2 provide formal stability proofs, neither for a surface-adaptive frame: the 2 papers with genuine surface-adaptive frames both lack proofs. ISO 10218 [10] and ISO/TS 15066 [11] specify power and force limits for collaborative operation; without formal analysis, compliance must be re-established empirically for every platform, payload, and configuration, yielding no transferable safety argument. Formal stability proofs, whether Lyapunov-based, passivity-based, or via energy tanks, should be mandatory for controllers intended for physical human-robot interaction in contact tasks.

Application

For system designers: use the 5-component framework as a design checklist. The review's component decomposition (C1–C5) with coverage statistics (Table 4.3), documented trade-offs (Sections 4.2–4.6), and specific reference implementations provides a structured resource for teams building assistive surface finishing systems. Each component has identified failure modes: force-based estimation without friction compensation corrupts the normal estimate; isotropic admittance parameters cannot distinguish force regulation from motion assistance; force-based ergonomic proxies diverge from actual muscular load. These documented failure modes can inform design decisions before implementation.

Consequence for the Field

The circular dependency will persist without interdisciplinary work. The structural explanation identified in this review (Section 5.1) implies that incremental improvements within single research communities, better estimation algorithms tested autonomously, better pHRI controllers tested on flat surfaces, better ergonomic metrics applied to simplified tasks, will continue to produce solutions that are individually strong but collectively unusable for the target application: assisting human workers on geometrically uncertain surfaces. Breaking this cycle requires research that simultaneously addresses estimation under human coupling, control with uncertain geometry, and ergonomic validation on realistic workpieces. The requirements identified in Section 4.9 outline one path; others are possible, but all must span the community boundaries this review has identified.

6.3. Closing Remarks

The systematic review mapped 71 papers across five components with quantified coverage, applied the cynical assessment questions to the 28 highest-relevance papers, identified structural technical incompatibilities between surface estimation and physical human-robot interaction, and characterized a specific integration gap. The resulting analytical mapping, including the component coverage statistics, the cross-community structural explanation, and the documentation of why existing partial solutions do not compose, did not previously exist in the literature: a systematic search of approximately 12,600 records, updated to July 2026, surfaced no comparable synthesis. It provides a foundation for researchers working at the intersection of surface estimation, physical human-robot interaction, and ergonomic validation.

The gap analysis establishes that no existing paper solves the integration problem: not one of the 71 reviewed studies, and none of the 113 further records screened in the July 2026 update. It does not establish that the integration is feasible. Whether a surface normal estimator can converge under human-guided motion, whether a Lyapunov stability proof can be obtained for the combined adaptive-frame and pHRI system, and whether ergonomic benefit survives on geometrically uncertain surfaces are empirical and analytical questions that future work must answer. The review provides evidence-based reasons to attempt the integration; it cannot guarantee success.

The motivation for this work remains the one stated in Chapter 1: workers performing surface finishing tasks develop musculoskeletal disorders from sustained force application. The literature shows that assistive devices can reduce this burden on geometrically simple surfaces. Extending this benefit to the curved and freeform workpieces that dominate real manufacturing requires the integration this review has mapped.

Author Contributions

Peter Ralph Desir conceived the review, developed the search and screening protocol, conducted the search, screening, and synthesis, and drafted the manuscript, with AI assistance as described in Section 2.8. Arno Stienen supervised the review and reviewed the manuscript.

References

[1] World Health Organization, *Musculoskeletal health*, WHO Fact Sheet, Accessed: 2026-07-28, 2022. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/musculoskeletal-conditions

[2] S. Liu, B. Wang, S. Fan, et al., "Global burden of musculoskeletal disorders and attributable factors in 204 countries and territories: A secondary analysis of the Global Burden of Disease 2019 study," *BMJ Open*, vol. 12, no. 6, e062183, 2022. DOI: 10.1136/bmjopen-2022-062183

[3] J. de Kok et al., "Work-related musculoskeletal disorders: Prevalence, costs and demographics in the EU," European Agency for Safety and Health at Work, Tech. Rep., 2019. DOI: 10.2802/66947

[4] S. Belli and F. De Keulenaer, "OSH pulse 2025 – occupational safety and health in the era of climate and digital change," European Agency for Safety and Health at Work, Tech. Rep., 2025. DOI: 10.2802/0978422

[5] B. P. Bernard and V. Putz-Anderson, "Musculoskeletal disorders and workplace factors: A critical review of epidemiologic evidence for work-related musculoskeletal disorders of the neck, upper extremity, and low back," National Institute for Occupational Safety and Health, Tech. Rep. 97-141, 1997.

[6] S. Bevan, "Economic impact of musculoskeletal disorders (MSDs) on work in Europe," *Best Practice & Research Clinical Rheumatology*, vol. 29, no. 3, pp. 356–373, 2015. DOI: 10.1016/j.berh.2015.08.002

[7] National Research Council and Institute of Medicine, *Musculoskeletal Disorders and the Workplace: Low Back and Upper Extremities*. Washington, DC: National Academies Press, 2001. DOI: 10.17226/10032

[8] A. W. Stedmon, H. Howells, J. R. Wilson, and I. Dianat, "Ergonomics/human factors needs of an ageing workforce in the manufacturing sector," *Health Promotion Perspectives*, vol. 2, no. 2, pp. 112–125, 2012, PMC3963628. DOI: 10.5681/hpp.2012.014

[9] International Federation of Robotics, "World robotics 2025 – industrial robots," IFR Statistical Department, Tech. Rep., 2025, Published 25 September 2025. Accessed: 2026-07-28. [Online]. Available: https://ifr.org/worldrobotics/report-2025

[10] International Organization for Standardization, *ISO 10218-1:2011 – robots and robotic devices – safety requirements for industrial robots – part 1: Robots*, International Standard, ISO, Geneva, Switzerland, 2011.

[11] International Organization for Standardization, *ISO/TS 15066:2016 – robots and robotic devices – collaborative robots*, Technical Specification, ISO, Geneva, Switzerland, 2016.

[12] M. J. Page et al., "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, n71, 2021. DOI: 10.1136/bmj.n71

[13] T. M. Stepien, L. M. Sweet, M. C. Good, and M. Tomizuka, "Control of tool/workpiece contact force with application to robotic deburring," *IEEE Journal on Robotics and Automation*, vol. RA-3, no. 1, pp. 7–18, 1987. DOI: 10.1109/JRA.1987.1087077

[14] H. Kazerooni, "Robotic deburring of two-dimensional parts with unknown geometry," *Journal of Manufacturing Systems*, vol. 7, no. 4, pp. 329–338, 1988. DOI: 10.1016/0278-6125(88)90043-X

[15] M. Jinno, T. Yoshimi, and A. Abe, "Force controlled grinding robot system for unstructured tasks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 2, Raleigh, NC, 1992, pp. 1117–1124. DOI: 10.1109/IROS.1992.594529

[16] N. Houshangi, "Position/force control of a robot manipulator over an unknown surface," in *IEEE International Conference on Systems, Man and Cybernetics (SMC)*, 1992, pp. 1128–1132. DOI: 10.1109/ICSMC.1992.271638

[17] M. Tarokh and S. Bailey, "Adaptive fuzzy force control of manipulators with unknown environment parameters," *Journal of Robotic Systems*, vol. 14, no. 5, pp. 341–353, 1997. DOI: 10.1002/(SICI)1097-4563(199705)14:5<341::AID-ROB1>3.0.CO;2-P

[18] J. Baeten, H. Bruyninckx, and J. De Schutter, "Shared control in hybrid vision/force robotic servoing using the task frame," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002, pp. 2128–2133. DOI: 10.1109/IRDS.2002.1041581

[19] M. Iwasaki, N. Tsujiuchi, and T. Koizumi, "Force control for unknown environment using sliding mode controller with gain-scheduled variable hyperplane," in *IEEE 28th Annual Conference of the Industrial Electronics Society (IECON)*, vol. 3, 2002, pp. 1812–1817. DOI: 10.1109/IECON.2002.1185246

[20] F. Shao, S. Endo, S. Hirche, and F. Ficuciello, "Interactive force-impedance control," *IEEE Robotics and Automation Letters*, vol. 11, no. 5, pp. 6488–6495, 2026. DOI: 10.1109/LRA.2026.3681128

[21] T. Wang, S. Xiao, Z. Lin, X. Chen, and L. Yang, "MAGIC: Mixed reality-enhanced active constraint generation and impedance control for human–robot co-manipulation," *IEEE Transactions on Automation Science and Engineering*, vol. 23, pp. 8216–8233, 2026. DOI: 10.1109/TASE.2026.3683134

[22] Y. Zhang et al., "Intention-aware hybrid motion-force control with weighted normal estimation for telerobotic ultrasonic testing," *Robotics and Computer-Integrated Manufacturing*, vol. 103, p. 103 359, 2027. DOI: 10.1016/j.rcim.2026.103359

[23] X. Ma, T. Wang, K. Li, J. Liu, X. Li, and F. Guo, "Adaptive variable impedance constant force control of robot grinding for casting surface defects," *The International Journal of Advanced Manufacturing Technology*, 2026. DOI: 10.1007/s00170-026-18248-6

[24] S. M. M. Rahman, "HRC metrology: Assessment criteria, metrics and methods for human–robot co-manipulation tasks," *Machines*, vol. 14, no. 3, p. 336, 2026. DOI: 10.3390/machines14030336

[25] F. Zeng, X. Wang, Z. Qing, H. Li, X. Zhang, and X. Hu, "The cross-domain transferable RACN–WGAN-based external force estimation for sensor-less force control and human–robot collaboration tasks," *Advanced Engineering Informatics*, vol. 74, p. 104 738, 2026. DOI: 10.1016/j.aei.2026.104738

[26] C. Zhou, H. Liao, H. H.-T. Chan, G. Liu, and W.-H. Liao, "Fuzzy neural network-based admittance control for robot under unknown environments," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2025, pp. 1858–1863. DOI: 10.1109/ROBIO66223.2025.11377432

[27] A.-J. Li, J. Li, and S. Yu, "Passivity-guaranteed variable impedance control for unknown surface treatment via friction-decoupled geometry estimation," in *6th International Symposium on Intelligent Robotics and Systems (ISoIRS)*, 2026. DOI: 10.1109/ISOIRS70157.2026.11545297

[28] W. Lu, L. Ma, D. Huang, J. Lin, and D. Wang, "Variable admittance control in feature space based on human intention," *IEEE/ASME Transactions on Mechatronics*, 2026. DOI: 10.1109/TMECH.2025.3647518

[29] S. Li, Q. Duan, Y. Lei, and G. Jiang, "Research on adaptive impedance control for serial manipulators in surface tracking applications," in *7th International Symposium on Robotics and Intelligent Manufacturing Technology (ISRIMT)*, 2025, pp. 218–223. DOI: 10.1109/ISRIMT67769.2025.11413154

[30] C. Liu, K. Zhao, Z. Lu, and C. Li, "Observer-based fuzzy adaptive admittance control for unknown environment-coupled physical human-robot interaction with prescribed tracking," *IEEE Transactions on Fuzzy Systems*, 2025. DOI: 10.1109/TFUZZ.2025.3646735

[31] L. Lu and J. T. Wen, "Human-directed robot motion/force control for contact tasks in unstructured environments," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2015, pp. 1165–1170. DOI: 10.1109/CoASE.2015.7294255

[32] M. Al Mashagbeh and M. B. Khamesee, "Force control for position interface industrial manipulator working in unknown environment," *Microsystem Technologies*, vol. 21, no. 12, pp. 2557–2563, 2015. DOI: 10.1007/s00542-015-2492-9

[33] Y. Zhang, J. Yao, and C. Qian, "Learning-based variable admittance control combined with NMPC for contact force tracking in unknown environments," *Actuators*, vol. 14, no. 7, p. 323, 2025. DOI: 10.3390/act14070323

[34] S. Kana, K.-P. Tee, and D. Campolo, "Human-robot co-manipulation during surface tooling: A general framework based on impedance control, haptic rendering and discrete geometry," *Robotics and Computer-Integrated Manufacturing*, vol. 67, p. 102 033, 2021. DOI: 10.1016/j.rcim.2020.102033

[35] V. Pruks and J.-H. Ryu, "Method for generating real-time interactive virtual fixture for shared teleoperation in unknown environments," *The International Journal of Robotics Research*, vol. 41, no. 9–10, pp. 925–951, 2022. DOI: 10.1177/02783649221102980

[36] Y. Karayiannidis and Z. Doulgeri, "Force/position regulation for a robotic finger in compliant contact with an unknown surface," in *13th Mediterranean Conference on Control and Automation (MED)*, Limassol, Cyprus, 2005, pp. 77–82. DOI: 10.1109/.2005.1466995

[37] S. Luo, M. Cheng, and R. Ding, "Adaptive virtual fixture based on learning trajectory distribution for comanipulation tasks," *IEEE Transactions on Human-Machine Systems*, vol. 55, no. 2, pp. 165–175, 2025. DOI: 10.1109/THMS.2025.3540123

[38] E. Nasiri and L. Wang, "Hybrid force motion control with estimated surface normal for manufacturing applications," in *21st International Conference on Ubiquitous Robots (UR)*, 2024, pp. 125–132. DOI: 10.1109/UR61395.2024.10597498

[39] Z. Doulgeri and Y. Karayiannidis, "Force/position control self-tuned to unknown surface slopes using motion variables," *Robotica*, vol. 26, no. 6, pp. 703–710, 2008. DOI: 10.1017/S026357470800430X

[40] Y. Yin, H. Hu, and Y. Xia, "Active tracking of unknown surface using force sensing and control technique for robot," *Sensors and Actuators A: Physical*, vol. 112, no. 2–3, pp. 313–319, 2004. DOI: 10.1016/j.sna.2004.01.044

[41] Z. Jiang, Y. Xu, and J. Xie, "Active tracking unknown surface based on force control for robot," in *Second International Conference on Digital Manufacturing and Automation (ICDMA)*, 2011, pp. 143–145. DOI: 10.1109/ICDMA.2011.42

[42] Y. Li, Z. Xu, S. Li, Z. Liao, S. Li, and X. Zhou, "Robot compliance control framework for grinding thin-walled parts with unknown surface: Deformation and orientation adaptation," *Robotics and Computer-Integrated Manufacturing*, vol. 98, p. 103 147, 2026. DOI: 10.1016/j.rcim.2025.103147

[43] S. Jung, T. C. Hsia, and R. G. Bonitz, "Force tracking impedance control for robot manipulators with an unknown environment: Theory, simulation, and experiment," *The International Journal of Robotics Research*, vol. 20, no. 9, pp. 765–774, 2001. DOI: 10.1177/02783640122067651

[44] S. Jung, T. C. Hsia, and R. G. Bonitz, "Force tracking impedance control of robot manipulators under unknown environment," *IEEE Transactions on Control Systems Technology*, vol. 12, no. 3, pp. 474–483, 2004. DOI: 10.1109/TCST.2004.824320

[45] N. Z. Azlan, H. Yamaura, and I. Suwarno, "Function approximation technique-based adaptive force-tracking impedance control for unknown environment," *International Journal of Robotics and Control Systems*, vol. 5, no. 1, pp. 381–406, 2025. DOI: 10.31763/ijrcs.v5i1.1029

[46] X. Liang, H. Zhao, X. Li, and H. Ding, "Force tracking impedance control with unknown environment via an iterative learning algorithm," in *3rd International Conference on Advanced Robotics and Mechatronics (ICARM)*, 2018, pp. 158–164. DOI: 10.1109/ICARM.2018.8610769

[47] K. Haninger, M. Radke, A. Vick, and J. Krüger, "Towards high-payload admittance control for manual guidance with environmental contact," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4275–4282, 2022. DOI: 10.1109/LRA.2022.3150051

[48] M. Mujica, M. Crespo, M. Benoussaad, S. Junco, and J.-Y. Fourquet, "Robust variable admittance control for human-robot co-manipulation of objects with unknown load," *Robotics and Computer-Integrated Manufacturing*, vol. 79, p. 102 408, 2023. DOI: 10.1016/j.rcim.2022.102408

[49] S. Grafakos, F. Dimeas, and N. Aspragathos, "Variable admittance control in pHRI using EMG-based arm muscles co-activation," in *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2016, pp. 1900–1905. DOI: 10.1109/SMC.2016.7844516

[50] J. Bae, K. Kim, J. Huh, and D. Hong, "Variable admittance control with virtual stiffness guidance for human-robot collaboration," *IEEE Access*, vol. 8, pp. 117 335–117 346, 2020. DOI: 10.1109/ACCESS.2020.3004872

[51] P. Franceschi, N. Pedrocchi, and M. Beschi, "Adaptive impedance controller for human-robot arbitration based on cooperative differential game theory," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 7881–7887. DOI: 10.1109/ICRA46639.2022.9811853

[52] P. P. Niaz, E. Erzin, and C. Basdogan, "Learning-based adaptive admittance controller for efficient and safe pHRI in contact-rich manufacturing tasks," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 9299–9306. DOI: 10.1109/IROS58592.2024.10802530

[53] D. Reyes-Uquillas and T. Hsiao, "Safe and intuitive manual guidance of a robot manipulator using adaptive admittance control towards robot agility," *Robotics and Computer-Integrated Manufacturing*, vol. 70, p. 102 127, 2021. DOI: 10.1016/j.rcim.2021.102127

[54] V. Pruks and J.-H. Ryu, "A framework for interactive virtual fixture generation for shared teleoperation in unstructured environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 10 234–10 241. DOI: 10.1109/ICRA40945.2020.9196579

[55] L. Han, J. Zhang, and H. Wang, "Shared control in pHRI: Integrating local trajectory replanning and cooperative game theory," *IEEE Transactions on Robotics*, vol. 41, pp. 1263–1277, 2025. DOI: 10.1109/TRO.2025.3532510

[56] M. Bednarczyk, H. Omran, and B. Bayle, "EMG-based variable impedance control with passivity guarantees for collaborative robotics," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4307–4312, 2022. DOI: 10.1109/LRA.2022.3149575

[57] T. Li, H. Xing, H. D. Taghirad, and M. Tavakoli, "EMG-based hybrid impedance-force control for human-robot collaboration on ultrasound imaging," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 670–675. DOI: 10.1109/IROS47612.2022.9981615

[58] Z. Li, H. Wei, H. Zhang, and C. Liu, "A variable admittance control strategy for stable and compliant human-robot physical interaction," *IEEE Robotics and Automation Letters*, vol. 10, no. 2, pp. 1138–1145, 2025. DOI: 10.1109/LRA.2024.3519885

[59] F. Ficuciello, L. Villani, and B. Siciliano, "Variable impedance control of redundant manipulators for intuitive human-robot physical interaction," *IEEE Transactions on Robotics*, vol. 31, no. 4, pp. 850–863, 2015. DOI: 10.1109/TRO.2015.2430053

[60] S. Choi, S. Ha, and W. Kim, "A multi-task energy-aware impedance controller for enhanced safety in physical human-robot interaction," *IEEE Robotics and Automation Letters*, vol. 10, no. 2, pp. 1345–1352, 2025. DOI: 10.1109/LRA.2024.3519871

[61] S. Choi, S. Hwang, and W. Kim, "Stable variable impedance control via CLF-MPC for physical human-robot interaction," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 14 750–14 757. DOI: 10.1109/IROS60139.2025.11247511

[62] B. Dong, T. An, H. Dong, R. Sun, C. Li, and B. Ma, "Adaptive fuzzy impedance control of human-robot interaction modular robot manipulators based on human motion intention estimation," in *14th International Conference on Information Science and Technology (ICIST)*, 2024, pp. 514–519. DOI: 10.1109/ICIST63249.2024.10805395

[63] V. Duchaine and C. M. Gosselin, "General model of human-robot cooperation using a novel velocity based variable impedance control," in *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces (WHC)*, 2007, pp. 445–451. DOI: 10.1109/WHC.2007.59

[64] Q. D. Le and H.-J. Kang, "Implementation of sensorless contact force estimation in collaborative robot based on adaptive third-order sliding mode observer," *Systems Science & Control Engineering*, vol. 10, no. 1, pp. 507–516, 2022. DOI: 10.1080/21642583.2022.2063201

[65] S.-H. Chien, J.-H. Wang, and M.-Y. Cheng, "Performance comparisons of different observer-based force-sensorless approaches for impedance control of collaborative robot manipulators," in *IEEE Conference on Industrial Cyberphysical Systems (ICPS)*, 2020, pp. 326–331. DOI: 10.1109/ICPS48405.2020.9274790

[66] T. Ren, Y. Dong, D. Wu, and K. Chen, "Impedance control of collaborative robots based on joint torque servo with active disturbance rejection," *Industrial Robot*, vol. 46, no. 4, pp. 518–528, 2019. DOI: 10.1108/IR-06-2018-0130

[67] H. Wang, Z. Wang, and H. Wang, "Impedance control strategy and experimental analysis of collaborative robots based on torque feedback," in *IEEE International Conference on Robotics and Biomimetics (ROBIO)*, Dali, China, 2019, pp. 2951–2957. DOI: 10.1109/ROBIO49542.2019.8961470

[68] F. Dimeas and N. Aspragathos, "Online stability in human-robot cooperation with admittance control," *IEEE Transactions on Haptics*, vol. 9, no. 2, pp. 267–278, 2016. DOI: 10.1109/TOH.2016.2518670

[69] M. Nini, A. Pupa, C. Secchi, C. Fantuzzi, and F. Ferraguti, "The art of not getting smacked: ISO/TS 15066-compliant variable admittance control for safe human-robot interaction," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 6336–6343. DOI: 10.1109/IROS60139.2025.11247254

[70] Z. Liao and Y. Wang, "Trust-based variable impedance control of human-robot cooperative manipulation," *Robotics and Computer-Integrated Manufacturing*, vol. 88, p. 102 730, 2024. DOI: 10.1016/j.rcim.2024.102730

[71] Y. M. Hamad, Y. Aydin, and C. Basdogan, "Adaptive human force scaling via admittance control for physical human-robot interaction," *IEEE Transactions on Haptics*, vol. 14, no. 4, pp. 750–761, 2021. DOI: 10.1109/TOH.2021.3071626

[72] J. Cacace, A. Finzi, and V. Lippiello, "Shared admittance control for human-robot co-manipulation based on operator intention estimation," in *15th International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, 2018, pp. 61–70. DOI: 10.5220/0006831800610070

[73] J. Cacace, R. Caccavale, A. Finzi, and V. Lippiello, "Variable admittance control based on virtual fixtures for human-robot co-manipulation," in *IEEE International Conference on Systems, Man and Cybernetics (SMC)*, 2019, pp. 1569–1574. DOI: 10.1109/SMC.2019.8914084

[74] J. Li, W. Zheng, and L. Ding, "Human-robot collaborative strategy based on behavioral intention recognition and switching impedance control for manipulator," in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, 2025, pp. 1–6. DOI: 10.1109/AIM64088.2025.11175742

[75] F. Dimeas, V. C. Moulianitis, C. Papakonstantinou, and N. Aspragathos, "Manipulator performance constraints in Cartesian admittance control for human-robot cooperation," in *IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, 2016, pp. 3049–3054. DOI: 10.1109/ICRA.2016.7487469

[76] B. G. C. Lima, E. Ferrentino, P. Chiacchio, and M. Vento, "Assistive force control in collaborative human-robot transportation," in *32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2023, pp. 2646–2651. DOI: 10.1109/RO-MAN57019.2023.10309518

[77] S. Devie, P.-P. Robet, Y. Aoustin, M. Gautier, and A. Jubien, "Optimized force and co-manipulation control using stiffness of force sensor with unknown environment," in *11th International Workshop on Robot Motion and Control (RoMoCo)*, Wasowo Palace, Poland, 2017, pp. 99–104. DOI: 10.1109/RoMoCo.2017.8003898

[78] C. K. Bitikofer, E. T. Wolbrecht, R. M. Maura, and J. C. Perry, "Comparison of admittance control dynamic models for transparent free-motion human-robot interaction," in *International Conference on Rehabilitation Robotics (ICORR)*, 2023, pp. 1–6. DOI: 10.1109/ICORR58425.2023.10304709

[79] S. Zhou and J. Song, "Impedance control for vehicle driving with human operation under unstructured environment," in *International Conference on Internet Computing and Information Services (ICICIS)*, 2011, pp. 159–162. DOI: 10.1109/ICICIS.2011.46

[80] D. Tebaldi, D. Onfiani, and L. Biagiotti, "On the analysis of stability, sensitivity and transparency in variable admittance control for pHRI enhanced by virtual fixtures," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 16 665–16 672. DOI: 10.1109/IROS60139.2025.11246940

[81] D. Bazzi, F. Roveda, A. M. Zanchettin, and P. Rocco, "A unified approach for virtual fixtures and goal-driven variable admittance control in manual guidance applications," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6378–6385, 2021. DOI: 10.1109/LRA.2021.3093283

[82] D. Zurlo, T. Heitmann, M. Morlock, and A. De Luca, "Collision detection and contact point estimation using virtual joint torque sensing applied to a cobot," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7533–7539. DOI: 10.1109/ICRA48891.2023.10160661

[83] M. Kanekiyo, H. Arita, K. Nakashima, and K. Tahara, "Desired contact force realization in unknown environments via multiple virtual dynamics-based control framework," in *IEEE 21st International Conference on Automation Science and Engineering (CASE)*, 2025, pp. 3130–3137. DOI: 10.1109/CASE58245.2025.11164077

[84] T. Zhang, Y. Yu, and Y. Zou, "An adaptive sliding-mode iterative constant-force control method for robotic belt grinding based on a one-dimensional force sensor," *Sensors*, vol. 19, no. 7, p. 1635, 2019. DOI: 10.3390/s19071635

[85] B. Balci, J. Roberts, J. Donovan, and P. Corke, "Force controlled robotic surface finishing with variable tool centre point (TCP)," in *IEEE 19th International Conference on Automation Science and Engineering (CASE)*, 2023, pp. 1–6. DOI: 10.1109/CASE56687.2023.10260672

[86] T. Song, L. Fang, and Y. Qian, "A novel robust adaptive impedance control scheme for contact force tracking with unknown environment," *IEEE Robotics and Automation Letters*, vol. 10, no. 7, pp. 6744–6751, 2025. DOI: 10.1109/LRA.2025.3568570

[87] J. Arnold and H. Lee, "Variable impedance control for pHRI: Impact on stability, agility, and human effort in controlling a wearable ankle robot," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2429–2436, 2021. DOI: 10.1109/LRA.2021.3062015

[88] F. Müller, J. Janetzky, U. Behrnd, J. Jäkel, and U. Thomas, "User force-dependent variable impedance control in human-robot interaction," in *IEEE 14th International Conference on Automation Science and Engineering (CASE)*, 2018, pp. 1328–1335. DOI: 10.1109/COASE.2018.8560340

[89] M. Wu, Y. He, and S. Liu, "Adaptive impedance control based on reinforcement learning in a human-robot collaboration task with human reference estimation," *International Journal of Mechanics and Control*, vol. 21, no. 1, pp. 21–31, 2020.

[90] M. Hanafusa and J. Ishikawa, "Operability evaluation of human-adaptive impedance control for human-collaborative robots," in *22nd IEEE International Conference on Industrial Technology (ICIT)*, 2021, pp. 80–86. DOI: 10.1109/ICIT46573.2021.9453547

[91] C. Huang, M. Zhu, S. Song, Y. Zhao, and J. Jiang, "Vision-based adaptive LT sliding mode admittance control for collaborative robots with actuator saturation," *Robotica*, vol. 42, no. 6, pp. 1986–2003, 2024. DOI: 10.1017/S0263574724000729

[92] T. Zhang and J. Xia, "Interconnection and damping assignment passivity-based impedance control of a compliant assistive robot for physical human-robot interactions," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 538–545, 2019. DOI: 10.1109/LRA.2019.2891434

[93] H. Maithani, J. A. Corrales Ramón, and Y. Mezouar, "Trust-based variable impedance control for cooperative physical human-robot interaction," in *IEEE International Conference on Mechatronics (ICM)*, 2019, pp. 706–711. DOI: 10.1109/ICMECH.2019.8722839