Journal of Intelligent & Robotic Systems (2026) 112:69
https://doi.org/10.1007/s10846-026-02401-0

REVIEW

# Toward Seamless Physical Human-Humanoid Interaction: Insights from Control, Intent, and Modeling with a Vision for What Comes Next

Gustavo A. Cardona[1] · Shubham S. Kumbhar[1] · Panagiotis Artemiadis[1]

Received: 25 September 2025 / Accepted: 14 April 2026


**Abstract**

Physical Human-Humanoid Interaction (pHHI) is a rapidly advancing field with significant implications for deploying robots in unstructured, human-centric environments. In this review, we examine the current state of the art in pHHI through three core pillars: (i) humanoid modeling and control, (ii) human intent estimation, and (iii) computational human models. For each pillar, we survey representative approaches, identify open challenges, and analyze current limitations that hinder robust, scalable, and adaptive interaction. These include the need for whole-body control strategies capable of handling uncertain human dynamics, real-time intent inference under limited sensing, and modeling techniques that account for variability in human physical states. Although significant progress has been made within each domain, integration across pillars remains limited. We propose pathways for unifying methods across these areas to enable cohesive interaction frameworks. This structure enables us not only to map the current landscape but also to propose concrete directions for future research that aim to bridge these domains. Additionally, we introduce a unified taxonomy of interaction types based on modality, distinguishing between direct interactions (e.g., physical contact) and indirect interactions (e.g., object-mediated), and on the level of robot engagement, ranging from assistance to cooperation and collaboration. For each category in this taxonomy, we provide the three core pillars that highlight opportunities for cross-pillar unification. Our goal is to suggest avenues to advance robust, safe, and intuitive physical interaction, providing a roadmap for future research that will allow humanoid systems to effectively understand, anticipate, and collaborate with human partners in diverse real-world settings.

**Keywords** Physical human-humanoid interaction · Humanoid models and control · Human state prediction and estimation

## 1 Introduction

In recent years, advances in artificial intelligence and the decreasing cost of technological components have led to an unprecedented surge in the development and deployment of humanoid robots. Market forecasts predict remarkable growth, with some estimates suggesting that billions of humanoid robots could be deployed around the world by 2040 [1]. Their anthropomorphic form and dexterous movements make them especially well-suited to navigate, manipulate, and interact within environments designed for people, which motivates an urgent push to expand their capabilities and versatility. Cutting-edge platforms, such as Honda's ASIMO [2], Boston Dynamics' Atlas [3], Unitree's G1 [4], and Agility Robotics' Digit [5], with the help of multiple researchers from industry and academia, have shown the promise of these robots for performing locomotion tasks. Experts foresee them moving beyond the limited scope of past industrial applications and becoming a fundamental part of daily life. In areas such as healthcare [6], humanoid robots can assist with tasks including lifting and transporting patients, guiding rehabilitation exercises, and serving as agile assistants in clinics and hospitals, thereby helping to reduce the workload of medical personnel. In domestic and service roles [7], humanoids could perform tasks of cleaning, fetching, and education, while also providing essential support to the elderly and people with disabilities. Furthermore, in dangerous and inaccessible environments, such as disaster zones [8] or outer space [9], humanoid robots could

✉ Panagiotis Artemiadis
partem@udel.edu

Gustavo A. Cardona
cardona@udel.edu

Shubham S. Kumbhar
shubhamk@udel.edu

1 Department of Mechanical Engineering, University of Delaware, 210 S College Ave, Newark 19716, DE, USA

carry out critical missions, keeping humans safe from harm, among many other applications.

Humanoid robots are poised to become indispensable partners across diverse settings, including factory floors, disaster relief operations, hospital wards, and domestic environments, fundamentally transforming human activities and interactions [10]. Applications encompass multiple domains in which physical coupling is essential for task execution: $i$) healthcare and assisted living, including sit-to-stand assistance, patient transfers, gait support, and rehabilitation, $ii$) logistics, manufacturing, and construction, such as co-carrying bulky payloads, collaborative positioning, and tool handling, $iii$) maintenance within constrained or hazardous environments, including nuclear decommissioning and ship or aircraft maintenance, $iv$) disaster response, such as rubble handling and victim extraction support, and $v$) immersive teleoperation and training, where physical coupling enhances intent communication and safety. In these contexts, physical Human–Humanoid Interaction (pHHI) uniquely demands simultaneous balance-safe whole-body control, intent anticipation under partial observability, and human-aware feasibility and cost reasoning. These diverse applications demonstrate that the primary technical challenge extends beyond achieving isolated competencies in walking [11], balancing [12], or grasping [13]. The central objective is to enable safe, intuitive, and effective pHHI, which constitutes the main focus of this review.

pHHI refers to the bidirectional exchange of forces, motion, and information between a human and a humanoid robot through intentional physical contact to achieve a shared task. In contrast to robot-alone tasks, pHHI embeds the human as a dynamic, uncertain, and co-equal element of the control loop. This introduces complexities beyond those already present in humanoid control for locomotion, manipulation, or balance, as human responses to robot actions, as well as subtle variations in contact, timing, or posture, can substantially influence task outcomes.

Unlike purely virtual or verbal interaction, pHHI requires continuous physical coupling in which both partners adapt their behaviors in response to real-time variations in force, posture, and intent. These couplings, also referred to as *interaction modalities* [14], can be broadly classified into two types. In *direct interaction*, such as body-to-body contact (e.g., assisting a person to stand [15]), the humanoid must exhibit whole-body compliance, reliable contact detection, and human-safe impedance control. In *indirect, object-mediated interaction*, such as co-manipulation [16] or teleoperation via a haptic interface [17], effective coordination depends on shared object models, force/torque coupling, and accurate task state estimation. In both cases, the physical coupling serves not only as a means to accomplish the task but also as a communication channel: the human and the humanoid both convey their intents through applied forces, while adapting to reach a common consensus. Any mismatch in intent reduces efficiency and can compromise the stability of the coupled system. Consequently, humanoids must explicitly integrate safety mechanisms such as compliance modulation and impedance control, as their actions directly affect human safety during interaction. Therefore, in pHHI these couplings inherently combine dynamic stability, compliance, intent inference, and safety assurance into a unified control problem.

The existence of multiple autonomous entities often puts these control problems under the constraints of *shared autonomy* and human-centered ergonomics. The concept of shared autonomy [18] introduces an additional layer of autonomy that spans a spectrum from pure teleoperation to fully proactive assistance, with cooperative modes in which roles, timing, and force-sharing change dynamically as the task phases and human state evolve. Some researchers also refer to this as arbitration. This perspective emphasizes that effective pHHI requires not only robust physical coupling and control but also adaptive role allocation, ensuring both efficiency and safety in human–robot partnerships.

These defining aspects of pHHI distinguish it from conventional humanoid control, making it far more than a straightforward application of existing methods. pHHI requires dedicated consideration of the coupled human-robot system. In contrast to fixed-base manipulators, humanoid robots possess a floating base and are underactuated. They typically function within multi-contact regimes, involving interactions between the feet and the ground, as well as the hands, objects, or the environment. These characteristics fundamentally alter pHHI in several key aspects: 1) *Whole-body coupling:* Interaction wrenches applied at the hands influence centroidal momentum and overall balance. An arm-generated interaction force may result in a fall if whole-body stability is not maintained. 2) *Hybrid multi-contact feasibility:* Contact constraints are unilateral and limited by friction, while transitions between contacts introduce hybrid dynamics. Planners and controllers must account for friction cones, support polygons, and the feasibility of transitions. 3) *Integrated optimization of locomotion and manipulation:* Effective cooperation frequently necessitates concurrent planning of footstep timing, body posture, and hand or object motion, all within real-time computational constraints. 4) *Enhanced safety requirements:* In addition to collision avoidance and impact limits, systems must guarantee balance recoverability, prevent slipping, maintain predictable compliance, and implement safe fallback behaviors in response to human disturbances. 5) *Uncertainty-aware and safety-constrained learning:* Distributional shifts may lead to loss of balance or unsafe contacts. Therefore, learned components should quantify uncertainty and be constrained by certified safety mechanisms with explicit fallback mode switching. Additionally, classical kinematic, dynamic, and trajectory-centric paradigms must be extended to account for human biomechanics, latent intent, comfort, social expectations, and safety attributes, which are noisy, time-varying, and person-specific.

Given all the challenges of pHHI, we believe there are four key components that must be considered: hardware, humanoid modeling and control, human intent estimation, and computational human models. The foundational layer is the hardware [19]. This establishes the upper limit of system performance. It determines the robot's sensor and control bandwidth and its resilience to impact. Hardware also determines the extent to which authority can be safely transferred to a human partner and the system's ability to handle unexpected events. Although hardware is recognized as a prerequisite, we do not provide a comprehensive review of mechanisms and actuators here. Instead, this review concentrates on the three algorithmic pillars that build upon this foundation to enable robust pHHI, as shown in Fig. 1.

The first of these pillars centers on humanoid modeling and control. Here, we consider the principal features that the model and controller should have, such as stability, safety, and adaptability. Controllers must maintain balance and feasibility when facing unpredictable human-imposed forces (stability). They must protect both partners during contact events and transitions (safety). They should also intelligently reallocate momentum and authority as the task, environment, or partner's behavior changes (adaptability). These demands shape the design of control architectures, ranging from single locomotion activities to loco-manipulation. These requirements establish the foundation on which the other pillars can reliably build.

The second pillar centers on human intent estimation. Raw sensory observations are transformed into actionable beliefs and short-horizon forecasts under three critical themes. First, multi-sensor fusion is used to create low-latency, fault-tolerant cues by combining data from pose, wrench, physiological, and gaze sensors. Next, the choice of modality involves matching appropriate probabilistic, decision-theoretic, or learning-based models to each signal stream, ensuring timely and meaningful outputs. Finally, generalization ensures performance across diverse partners, tasks, and scenarios, allowing the system to resist distribution shifts and explicitly represent uncertainty that downstream controllers can utilize for safer decisions.

The third and final algorithmic pillar centers on the development of computational human models. Here, the robot moves beyond simple collision avoidance; it must consider what is comfortable, feasible, and fair for the person it is assisting. This pillar is defined by ergonomics, fidelity, and personalization. These models must capture metrics of comfort, effort, and physical limits (ergonomics). They must be reliable enough to support predictions and constraint-setting. This can range from biomechanical to neuromechanical levels (fidelity). They must also adapt online to user-specific parameters and preferences (personalization). Exported as constraints, costs, or targets, these models allow the robot's

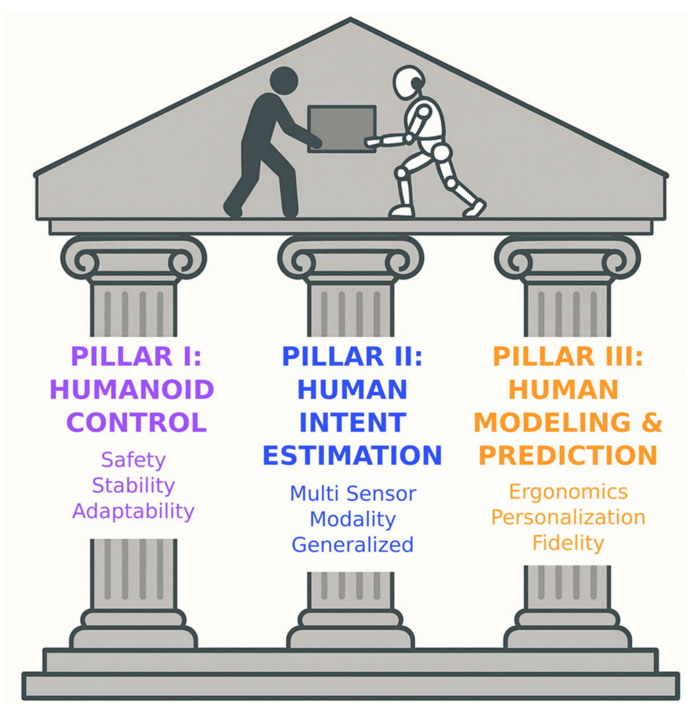


**Fig. 1** Conceptual structure of pHHI with three algorithmic pillars: Humanoid Control (Stability, Safety, Adaptability), Human Intent estimation (Multi-Sensor, Modality, Generalization), and Human Models (Ergonomics, Fidelity, Personalization)

controller to share effort and respect human limits intelligently. In doing so, these models enable the robot to move from simply preventing failure to actively promoting safe, efficient, and human-centered collaboration.

A primary challenge in the field is not the absence of sophisticated methods within each domain. Rather, there is a fundamental lack of integration among them. The existing literature often presents solutions that optimize individual aspects such as control, prediction, or human modeling in isolation. However, effective pHHI in real-world scenarios relies on the composition of these elements. For example, intent estimation must communicate calibrated uncertainty to the controller. Human models must define ergonomic limits as explicit constraints for the planner. The controller must provide stability assurances for higher-level processes. This review is motivated by the need for a unified framework with well-defined interfaces. Such a framework should connect sensing and prediction to whole-body action, while maintaining human comfort and safety.

The structure of this review paper begins with Section 2 that defines pHHI, outlines representative scenarios, and introduces a two-axis taxonomy (i.e., interaction modality and shared autonomy). Sections 3–5 develop the three pillars: Section 3 covers humanoid dynamics and control. Section 4 addresses human intent estimation and prediction. Section 5 reviews computational human models. In each pillar, we synthesize foundational and recent work, assess relevance to pHHI, and identify remaining challenges. In Section 6, we then transition to a broader discussion of the field's primary challenges and outline an architecture for a unified framework that integrates these proposed pillars. Additionally, we identify key future research directions. Finally, Section 7 presents the conclusions of this comprehensive analysis.

## 2 Review of Physical Human-Humanoid Interaction

Advancements in physical human-humanoid interaction are deeply rooted in the evolution of the robot's ability to control its complex dynamics while safely managing external forces from a human partner. The literature reveals a progression from fundamental safety principles to highly integrated, task-specific applications. A comprehensive review of human-humanoid interaction and cooperation is presented in [20], which provides a broad overview of this landscape. Our synthesis identifies two core themes that consider only physical interaction: i) task-specific implementations for key pHHI scenarios such as collaborative transport and physical assistance, and ii) control strategies that seek to ensure safe and compliant behavior.

### 2.1 Task-Specific Implementations of pHHI

The discussion is anchored by classifying pHHI along two practical axes: (i) interaction modality, which distinguishes between direct body-to-body contact and indirect object-mediated coupling, and (ii) robot engagement, which ranges from assistance to cooperation to collaboration. This taxonomy in Table 1 demonstrates that higher levels of engagement typically require stronger intent inference, more comprehensive human-feasibility models, and more stringent safety-certified whole-body control.

**Table 1** Two-axis taxonomy of pHHI (interaction modality × robot engagement) with representative tasks and dominant technical requirements. Increasing engagement generally requires stronger intent belief, richer human feasibility models, and tighter safety-certified whole-body control

| Interaction modality | Assistance (robot supports) | Cooperation (shared effort) | Collaboration (shared decisions) |
|---|---|---|---|
| Direct body-to-body contact | **Examples:** sit-to-stand support; balance assistance; guided posture correction.<br>**Dominant needs:** compliant contact regulation; stability margins; comfort constraints; predictable haptic behavior. | **Examples:** coupled locomotion guidance; assisted stepping; physical guidance with phase alignment.<br>**Dominant needs:** phase/intent estimation; impedance scheduling; hybrid contact safety; disturbance-aware control. | **Examples:** shared recovery behaviors; coordinated maneuvering with mutual adaptation.<br>**Dominant needs:** belief-aware planning; uncertainty-to-safety gating; mode switching; explicit failure/fallback logic. |
| Indirect object-mediated coupling | **Examples:** follower co-carrying; stabilizing an object while the human leads.<br>**Dominant needs:** wrench regulation at the object; object-state estimation; safe compliance; slip detection. | **Examples:** load-sharing carry; collaborative positioning (e.g., aligning a panel).<br>**Dominant needs:** role inference; ergonomics-aware cost models; footstep + hand co-planning; constraint-aware control. | **Examples:** tight-space transport; collaborative manipulation/assembly with discrete choices.<br>**Dominant needs:** multi-step planning under uncertainty; contact switching feasibility; calibrated intent + robust constraints; safe fallback under ambiguity. |

#### 2.1.1 Collaborative Transportation and Co-manipulation

A substantial amount of research has concentrated on the fundamental task of two partners collaboratively carrying an object, as shown in Fig. 2a. This task requires precise coordination of both locomotion and manipulation [21]. Early studies introduced homotopy-based controllers [22] designed for this purpose and examined the problem from a haptic feedback perspective. Notably, the work in [23, 24] addressed the challenge using specialized walking pattern generators that adapt to each partner's influence. Then, in [25], the approach incorporates both vision and haptic sensing to assess each partner's contribution, ultimately developing a comprehensive framework for human-humanoid collaborative carrying [16]. Other researchers have investigated model-based control methods for managing load distribution [26] or utilized vision systems to allow a robot to follow a human's lead while carrying a large object [27]. In [28], an object was jointly transported by finding the robot's CoM targets based on an interaction admittance model. A critical challenge in this field is making the interaction proactive rather than merely reactive. In [29, 30], the authors examined how a humanoid robot can anticipate a human's intentions, facilitating a smoother and more efficient joint transport task.

#### 2.1.2 Assistive Care and Ergonomics

Another important application area is providing direct physical assistance in healthcare and daily living activities as shown in Fig. 2b. The sit-to-stand and stand-to-sit motions are key benchmark tasks that require the robot to provide enough support without being excessive. Recent advances have demonstrated the effectiveness of robust controllers designed for this purpose, utilizing force sensing to synchronize with human movements [15, 31]. In addition to basic assistance, there is a new focus on ergonomic physical human-robot interaction (pHRI), where the robot's actions are optimized to minimize physical strain and reduce injury risk for the human partner. For instance, in [32], a control approach is explicitly developed for ergonomic payload lifting, where the robot's controller takes into account a model of human ergonomics. This represents a significant shift from simply completing a task to doing so in a manner that prioritizes the well-being of the human user. Other innovative forms of assistance include centaur-like systems that help individuals carry heavy loads over long distances [33].

### 2.2 Control Strategies

At the core of pHHI is the requirement that the robot must be physically compliant and safe to humans. Early and foundational work focused on establishing control laws that enable the robot to yield to external forces predictably, a principle essential for managing the uncertainties associated with physical contact. We distinguish methods by decision level: $i$) interaction-port behavior shaping, such as impedance, admittance, or hybrid force–motion control, which regulates local compliance and contact wrenches, $ii$) whole body feasibility and motion generation, including whole-body control, model predictive control, or trajectory optimization,

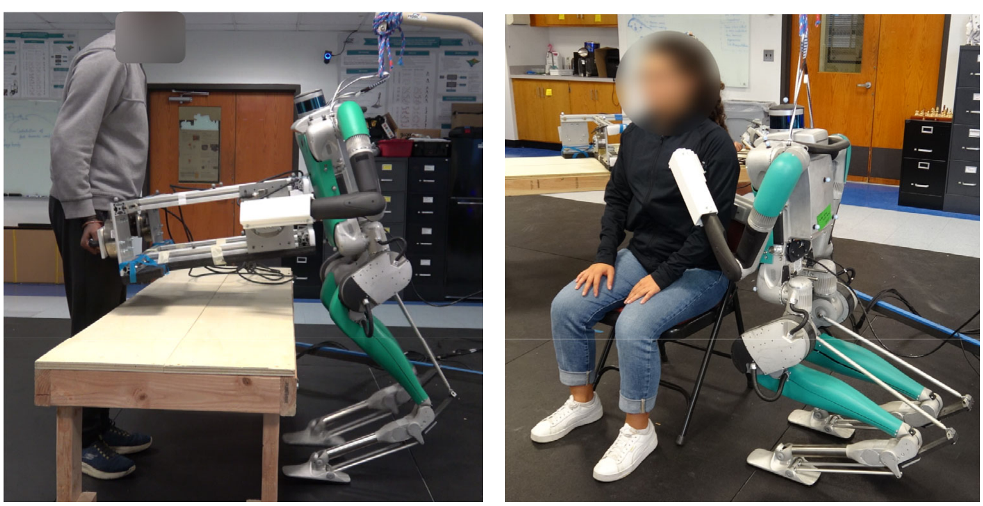

(a) Indirect comanipulation (b) Direct interaction

**Fig. 2** Comparison of interaction modalities in physical human–humanoid interaction. (**a**) Indirect, object-mediated co-manipulation (e.g., collaborative carrying), where coordination is mediated through a shared object and relies on shared object models and interaction-wrench regulation. (**b**) Direct, body-to-body assistance (e.g., emulated rehabilitation), which requires whole-body balance, reliable contact detection, and safe impedance/admittance control which enforce multi-contact constraints and balance while allocating forces across the body, and $iii$) learning-based augmentation, which is applied either as residual or adaptive components on top of ($i$) and ($ii$), or as learned predictors and reference generators that provide priors, costs, or targets for planning and control.

#### 2.2.1 Compliance and Force Control

A fundamental aspect of safe interaction in robotics is the ability to effectively modulate interaction forces. This regulation is often achieved through a technique called *impedance control* [34], in which the controller establishes a desired dynamic relationship, similar to a virtual spring-damper system, between the robot's movements and the external forces it experiences [35]. This approach allows the robot to be “soft" when pushed and “stiff" when it needs to apply force. This principle is essential and has been successfully applied in various applications, such as helping individuals stand up [31], where compliant control enables the robot to assist the person without creating a rigid or uncomfortable motion. However, a significant challenge is maintaining whole-body balance while being compliant. In [36], this issue was addressed by developing a full-body compliant controller that can maintain balance even in the presence of large and unpredictable external forces exerted by a human partner. In [28], an *admittance controller* is used, which, instead of regulating the dynamic relationship between position and force, finds a position for the Center of Mass (CoM) of the robot based on external forces generated by the expected interaction with an object jointly transported with a human.

#### 2.2.2 Whole-Body Control via Optimization

Optimization-based whole-body control treats motion and force generation as a constrained optimization problem that enforces dynamics and contact feasibility while balancing task objectives in real time. In practice, two frameworks dominate: hierarchical formulations (cascaded Quadratic Programming (QP) techniques with strict priorities) and weighted-sum formulations (a single QP with tunable objective trade-offs). In *hierarchical optimization*, also known as *Stack-of-Tasks* [24, 25, 29, 37], Whole-Body Control (WBC) is formulated as a sequence of strictly prioritized QPs. At the highest level, balance, contact feasibility, and actuation limits are enforced. Lower-priority tasks, such as interaction, end-effector positioning, and posture, are projected into the nullspace of higher-priority tasks. This approach is particularly practical for pHHI because it ensures that human–robot contact remains within stability and friction limits, even during disturbances caused by the human partner. However, this structure limits the ability to balance soft objectives, such as comfort and speed, across priorities, and it may introduce task-switching artifacts. In contrast, weighted-sum optimization addresses WBC by solving a single QP that minimizes a weighted sum of task costs, such as tracking accuracy, effort, and interaction-wrench regulation, subject to dynamics and contact constraints [16, 28, 38, 39]. This approach enables continuous, real-time trade-offs that are advantageous in pHHI. However, the effectiveness of this method depends on the selection of appropriate weights and the design of the cost function. Critical safety requirements should be enforced as hard constraints or protected through explicit feasibility checks.

Although WBC formulations frequently incorporate interaction objectives, such as wrench regulation or compliant behavior, within their cost functions, they differ from local impedance or admittance control by explicitly enforcing whole-body feasibility across contact points and balance constraints. Additionally, WBC allocates forces and motion across the entire humanoid system.

#### 2.2.3 Learning-Based Augmentation (Residuals, Predictors, and References)

Learning-based approaches complement optimization by acquiring whole-body control policies or residual adaptations directly from data. These methods enhance robustness and flexibility in environments where frequent physical contact is present. Recent studies have demonstrated the scalability of skill acquisition on humanoid robots [40]. In addition to policy learning, two further research directions are considered. The first is Learning from Human Demonstration (LfD), which enables data-efficient imitation of expert behaviors. The second is Learning and Adapting Physical Interaction, which focuses on online adjustment of assistance, impedance, and force sharing during pHHI.

**Learning from Human Demonstration (LfD)** Instead of being explicitly programmed, robots can learn collaborative tasks by observing humans or being physically guided by them. This method is effective in teaching complex physical interactions. In [41], it was demonstrated that a humanoid robot could learn to lift objects through physical guidance. This approach, commonly referred to as imitation learning, enables the robot to capture the subtleties of the interaction, such as timing and force profiles, directly from an expert human [42].

**Learning and Adapting Physical Interaction** Beyond one-shot learning, some approaches enable robots to adapt their behavior continuously during interactions. The authors in [43] developed a framework for “physical interaction learning," enabling robots to adjust their behavior in real-time to better coordinate with a human partner during cooperative tasks. Another strategy employs finite state machines to track the progress of a collaborative task effectively. This method involves passing through distinct phases, such as "approaching," "grasping," and "lifting," while providing appropriate assistance at each stage [44].

### 2.3 Synthesis: Reactive vs Predictive vs Proactive

Table 2 summarizes the control strategies discussed in this section into three groups and outlines their key features, advantages, limitations, typical applications, and critical safety implications.

As summarized in the table, compliance/force control delivers intuitive, "forgiving" contact by shaping task-space dynamics. Its softness trades away authority and often relies on a lower layer to maintain balance and limits. Optimization-based WBC enforces multi-contact feasibility. Hierarchical QPs ensure feasibility but can show discontinuities at task or priority switches. Weighted-sum QPs allow smooth trade-offs but require explicit safety encoding. Learning-based methods add personalization and fill modeling gaps, however, without a constraint-aware wrapper, they are non-robust to shifts and rare events. In practice, the learning-based methods work best as residuals or reference generators under a certified WBC/MPC layer.

In most of the approaches reviewed in this section, robotic initiative is limited to reactive behaviors [15, 21, 22, 24, 26–28, 31, 33]. Robots respond to measured interaction forces, follow the motion of their human partners, and redistribute loads to maintain compliance and adhere to safety standards. Recent studies have examined predictive and proactive strategies. In predictive approaches [16, 25, 32], controllers are adapted based on forecasts of ergonomic risks or anticipated human contributions. However, the robot does not initiate task goals. In contrast, proactive strategies [29, 30] involve anticipating human partner movements or inferring intent through subtle actions, followed by redistributing momentum or adjusting footstep timing. Sensing modalities for reactive strategies primarily use haptic feedback, while predictive and proactive approaches incorporate kinematic and visual data.

Despite impressive progress in modeling and control strategies for pHHI tasks, research still tends to focus on isolated yet critical subproblems. This fragmentation is understandable, as safe, compliant controllers, balancing under contact, intent inference, and human variability are each challenging problems on their own. However, this conservative approach leaves a gap between lab demonstrations and deployable interaction. A common and cohesive framework is needed: one that first *models and controls the humanoid* at the whole-body, multi-contact level, while guaranteeing safety in an efficient and tractable way. Second, it *estimates and forecasts human intent* from noisy multimodal signals. Third, it embeds *human biomechanical/behavioral models* so that the robot optimizes not only the task success but also comfort and effort. In this review paper, we use this lens to survey the landscape and highlight bridging scenarios where currently separate lines of work can be stitched together into an actionable framework. We approach this by framing the field around three key pillars: humanoid modeling and control, human intent estimation, and computational human models. This approach helps reveal the requirements, challenges, and opportunities for building safe, adaptive, and human-centered systems.

## 3 Pillar I: Dynamics Models and Control Strategies for Humanoid Robots

Humanoid control operates at the intersection of free-floating, underactuated, hybrid, and multi-contact dynamics, making even seemingly simple motions complex. This complexity transforms physical interactions with people into a challenge of balancing feasibility, responsiveness, and safety. In this section, we explore this challenge by surveying the main modeling abstractions, ranging from full-order whole-body and centroidal dynamics to reduced-order templates and learned residual models, as shown in Fig. 3. Additionally, we mention their corresponding control paradigms, including classical control techniques, optimal control formulations with safety certificates, and learning-based controllers. For each method, we analyze its benefits (e.g., fast balance reflexes, force-aware coordination, comfort, and predictability), its limitations (e.g., model mismatch and latency), and how well it scales across different interaction modalities and levels of shared autonomy. Our goal is not to create an exhaustive catalog but to provide a coherent overview of fit-for-purpose tools and their complementary roles, setting the foundation for developing a safe, intuitive, and robust control stack for pHHI.

### 3.1 Dynamic Models for Humanoids

#### 3.1.1 Whole-Body Dynamics

At the core of this discussion are whole-body models of humanoid robots, which consist of numerous rigid links, closed contact constraints, and a free-floating base, making the entire system underactuated. These models, whether expressed in the *Lagrangian or Newton-Euler* form, serve as the foundation for modern control strategies, including Resolved Momentum Control (RMC) and optimization-based WBC systems [45–49]. In the context of pHHI, the accuracy of these models is crucial. When a person interacts with the robot by stabilizing it at the forearm or sharing a load through their hands, the forces involved result in motion effects at the base and ground contacts. Even minor inaccuracies in mass, inertia, or contact modeling can lead to

**Table 2** Comparison of primary control strategies for physical human-humanoid interaction (pHHI). This table outlines their key features, advantages, limitations, typical applications, and critical safety implications

| Control Strategy | Key Features | Advantages for pHHI | Limitations | Typical Use Cases | Safety Implications | Rep. Ref. |
|---|---|---|---|---|---|---|
| Compliance / Force Control | Regulates end-effector/task-space dynamics; explicit force–motion blending. | Safe and forgiving contact; intuitive feel; predictably yields to human wrenches. | Trade-off between softness and tracking authority; heuristic contact handling. | Hand-guiding, sit-to-stand assistance, and object co-transport. | *Inherently passive and stable*; provides a fundamental layer of safety. | [28, 31, 36] |
| Whole-Body Control (WBC) via Optimization | *(Hierarchical QP)* Cascaded QPs with strict priorities; hard constraints. | Feasible-by-construction; coordinates the entire body under contact. | Rigid priorities; discontinuities at task/priority switches; needs reliable estimates. | Balance during interaction (e.g., stabilizing the forearm while walking). | Safety via hard constraints (joint limits, friction), but *sensitive to unexpected forces*. | [24, 25, 29, 37] |
| | *(Weighted-Sum QP)* Single QP trading off multiple costs. | Smooth real-time trade-offs (e.g., comfort vs. speed). | Sensitive to weight tuning; safety must be explicitly encoded as constraints. | Co-manipulation, sharing tasks based on vision and haptics. | Relies on cost function design; no inherent safety guarantees without explicit constraints. | [16, 28, 38, 39] |
| Learning-based Methods | Policies / residuals learned from data; imitation from demonstrations. | Personalization to partner; captures nuances of timing/grip; complements physics models. | Distribution shift, weak formal guarantees, and data/annotation burden. | Teaching collaborative skills; adjusting assistance or impedance online. | *Provides no formal guarantees*; requires a safety-critical supervisor (e.g., CBF). | [40–44] |

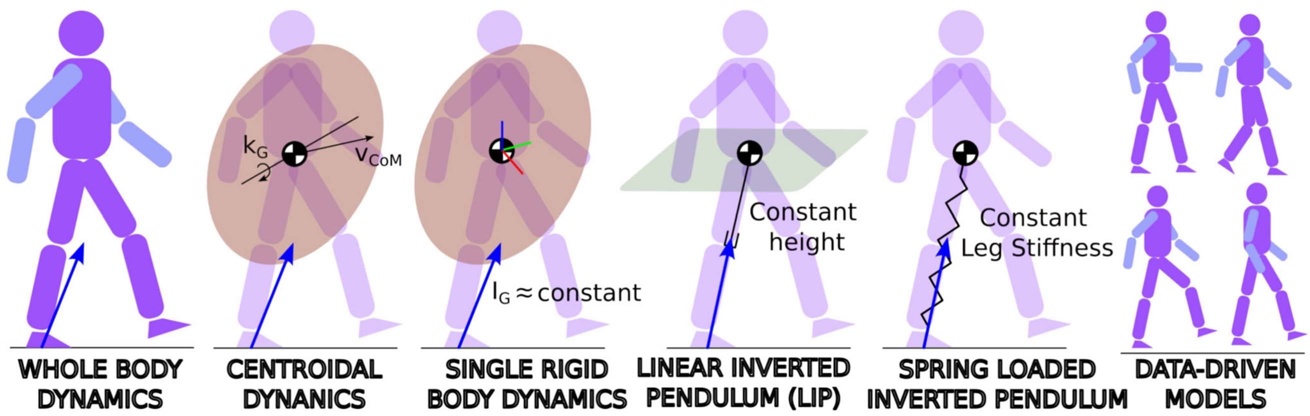


**Fig. 3** Modeling methodologies surveyed in this paper, arranged from high-fidelity to reduced-order and learned abstractions: whole-body multibody dynamics (WB), centroidal dynamics, single-rigid-body (SRB) approximations, Linear Inverted Pendulum (LIP), Spring-Loaded Inverted Pendulum (SLIP), and data-driven models

noticeable issues, such as stiffness spikes or drift during interactions. Using full models enables us to analyze balance, friction, and actuation limits simultaneously, an essential factor for practical direct assistance and co-manipulation. However, a challenge in pHHI arises from the combination of parametric uncertainty, the time-varying nature of the human impedance, and the need for real-time fidelity models. Without additional layers of compliance and predictive control, achieving perfect model matching is often unrealistic.

#### 3.1.2 Simplified Abstracted Models

To respond quickly to human actions, controllers utilize reduced-order models that consider different levels of abstraction. For instance, the Single Rigid Body (SRB) [48] and *Linear Inverted Pendulum Model (LIPM)* and its cart-table interpretation simplify balance to Center of Mass (CoM) motion and a feasibility region for Zero Moment Point (ZMP) or Center of Pressure (CoP). This simplification provides efficient guidelines for foot placement and CoM adjustment [50–54]. Another well-used model is the *Spring-Loaded Inverted Pendulum (SLIP)* [55, 56], which incorporates a compliant leg and better captures the bounce dynamics compared to the rigid-leg assumption of the LIPM model. This feature allows it to buffer impact forces and exhibit self-stabilizing properties. The *capture point* and *Divergent Component of Motion (DCM)* extend these concepts into smooth stepping and push-recovery strategies that are effective even when faced with sudden human movements [57–59]. In pHHI, these are the foundation of concepts for the robot's balance framework: when a partner misjudges the timing of a handover or shifts a shared object, the robot can execute reactive actions, such as stepping, adjusting the CoM, or adjusting its gait. Actions can be computed within milliseconds, rather than engaging in the full complexity of the whole-body dynamics. However, these abstractions prioritize computational speed over model fidelity. They typically assume coplanar and rigid contacts, and negligible angular momentum. Additionally, in the case of LIPM, they often neglect swing and impact dynamics as well as hand or object contacts and consider nearly constant CoM height. As a result, ZMP or capture margins may be inaccurately estimated during large disturbances or when operating on uneven terrain. In pHHI, where a human applies time-varying forces and introduces individualized constraints, this modeling gap becomes more apparent. Anticipatory control and personalization remain limited unless the model incorporates external wrench estimation, non-coplanar contacts, and regulation of variable height and momentum.

*Centroidal dynamics* serve as a compact model of CoM motion and global angular momentum, sitting between full fidelity and point-mass models. This approach captures how all joints contribute to maintaining whole-body balance [46, 48, 49]. Centroidal planners schedule feasible contact phases and momentum profiles. In practice, a whole body controller with inverse dynamics maps the centroidal momentum schedules to torque commands that respect the friction, contact, and actuation limits on the hardware, as demonstrated on Atlas-class platforms [47–49]. When viewed through a pHHI lens, this middle layer becomes crucial. Direct assistance tasks, such as sitting-to-standing movements and standing during gait, require precise regulation of the interaction forces in the hands while ensuring global balance in the feet.

However, centroidal models often neglect limb kinematics and generally assume predefined contact schedules, rigid contacts, and precise CoM or momentum estimates. Translating momentum plans into whole-body torques may encounter joint or velocity constraints, self-collisions, or reachability limitations, particularly when significant hand forces or non-coplanar grasp conditions are present. In pHHI, unmodeled partner impedance and sensing delays can compromise external-wrench regulation. Finally, inaccurate centroidal estimates or friction cone assumptions may result in infeasible or excessively aggressive momentum commands.

#### 3.1.3 Data-Driven Modeling

Due to the challenges of obtaining precise models for complex humanoids and the unpredictability of multi contact scenarios, *data-driven modeling* has become an essential tool. Data-driven dynamics, including learned residuals and models, complement physics in situations where uncertainty is a significant factor, such as partner-specific stiffness, unknown payloads, and unmodeled compliance. Neural policies and residuals are increasingly integrated into whole-body control stacks and loco-manipulation planners [60–69]. In the context of pHHI, these data-driven methods prove most useful as enhancements to the overall control architecture. For instance, a residual could compensate for an individual's grip stiffness, or a learned predictor might adjust expectations based on a partner's characteristic timing. However, learned models are sensitive to distribution shifts (e.g., new partners, surfaces, payloads) and provide weak guarantees of safety and stability.

### 3.2 Control Strategies for Humanoids

#### 3.2.1 Classical Control

**Whole-Body Control (WBC) via Inverse Dynamics** A common approach to WBC on humanoids is operational-space inverse dynamics, which prioritizes tasks and regulates internal forces. This method builds on torque control at key coordinates [70], organizing constraint tasks, motion tasks, and posture in a hierarchy. Lower-priority tasks are projected into the null space of higher-priority ones, enabling compliant interaction at the torque level [71–73]. Internal-force control has also been introduced, allowing unified regulation of the center of mass, operational tasks, and internal wrenches [74]. Momentum-based methods, such as Resolved Momentum Control, handle both linear and angular momentum as a whole-body task [45], thereby helping to maintain balance and manage interaction dynamics within inverse-dynamics controllers. When interacting with unknown objects, WBC can utilize online estimation of operational forces to adapt motion and force allocation as payloads change, thereby maintaining balance and contact [75]. Whole-body impedance control at the torque level has been demonstrated on robots like DLR Justin, enabling self-collision avoidance and compliant multi-task manipulation [76, 77].

In pHHI, torque-level whole-body control (WBC) is advantageous because it provides direct access to interaction ports, meaning the points of contact, such as the hands, forearms, and torso. This approach incorporates impedance control (which resists motion in response to force) and internal force or momentum regulation, allowing for transparent and adjustable contact forces. As a result, the system can adapt to unknown forces from partners or objects.

The primary limitations of this approach in pHHI would be the management of inequality constraints, such as friction cones (the allowable directions of force to avoid slipping) and torque or CoP bounds, as well as managing model uncertainty, such as payload shifts and soft contacts. Additionally, the method depends on precise torque sensing or estimation. In practice, many control architectures combine inverse-dynamics whole-body control with supervisory safety monitors, while comprehensive inequality constraint handling is typically addressed in the optimization section discussed later.

Promising directions for improvement include embedding contact and wrench observers directly within the WBC loop, and leveraging explicit centroidal momentum tasks to absorb partner-induced twists with the whole body rather than relying solely on ankle torques [46, 48, 49]. Additional enhancements involve adding safety layers that bound the CoM, ZMP, and friction regardless of the nominal task stack [78], incorporating learned residuals to address model discrepancies without neglecting physical principles [62, 63], and implementing automatic gain tuning to reduce the need for manual adjustments when switching between humans or tasks [79].

**Impedance, Admittance, and Hybrid Force/Position Control** Impedance, admittance, and hybrid force/position control are crucial concepts in humanoid loco-manipulation. If WBC determines the forces to apply, impedance and admittance govern how these forces are exchanged. In practice, humanoid manipulation controllers surround their interaction ports, such as hands and forearms, with whole-body impedance. This design allows for compliant motion when faced with unexpected human wrenches, ensuring that stance contacts remain feasible [28, 80–82]. This is the layer that individuals experience, it influences whether interaction feels soft, whether handovers are forgiving, and whether teleoperation remains safely energetic despite operator variations and uncertainties.

For indirect interactions, such as carrying a tray, hybrid force/position control separates the axes, allowing vertical force to be regulated while maintaining precise horizontal motion. This helps prevent sudden changes in load sharing. The challenges here involve classic trade-offs: if the system is too soft, the robot loses control, especially with heavy loads. If too stiff, the human may feel uncomfortable surges in force, and high admittance can endanger balance without proper constraints.

Improvements for pHHI could include variable impedance schedules that adapt to different phases of a task, for example, a softer approach, a firm lift, and a gentle handoff. Comfort can also be enhanced by limiting wrench rate and jerk, and by modulating admittance gains based on ZMP and CoP margins together with Control Barrier Function (CBF) constraints [81, 83]. Finally, co-designing with the whole-body controller enables impedance to shape the interaction ports while ensuring that the WBC preserves overall feasibility [82].

#### 3.2.2 Optimal Control and MPC/MPPI

Optimal control frames humanoid behavior as a constrained optimization problem, with Model Predictive Control (MPC) offering anticipatory control by optimizing footsteps, contacts, and force sharing over a receding horizon, allowing for replanning as conditions change.

**Modular Approaches** The high complexity and dimensionality of humanoids motivate a modular framework consisting of a high-level planner and a low-level controller. The high-level planner performs predictive planning of COM motion, footstep locations, forces, etc., while the low-level controller executes these plans on the physical robot. Both the high-level and the low-level are often formulated using optimization schemes. Model Predictive Control (MPC) has been widely used for footstep planing in humanoid locomotion [84]. Linear Inverted Pendulum (LIP) based MPC coupled with a low-level quadratic program has been widely used for legged locomotion on flat surfaces [50, 85, 86]. This method regulates the CoM and ZMP, placing steps to counteract pushes without solving full-order dynamics [50]. However, a limitation of this approach, especially for pHHI, is contact myopia: arm and interaction forces, as well as human-induced wrenches, are not perfectly modeled, which can lead to abrupt interactions. The authors explored extending LIP to account for interaction and combined impedance control in [28], which allows minor human-robot mismatches to be absorbed compliantly. Another reduced-order model that has been widely used for humanoid walking is the Single Rigid Body (SRB) [48] and the Spring-Loaded Inverted Pendulum (SLIP) [87, 88].

Most reduced-order models are obtained by making assumptions about the centroidal dynamics [46]. Approaches that use centroidal dynamics with MPC optimize CoM and global momentum while explicitly considering contact schedules, friction cones, and wrench bounds, eventually realizing the plan through whole-body control. This methodology has been successfully applied for rough terrain navigation and contact-rich maneuvers [47]. Additionally, it has been extended with multi-contact previews [81] to coordinate hand and foot movements, effectively leveraging centroidal dynamics with full kinematic consistency [48]. For pHHI, centroidal MPC could co-optimize foot placement. Nonetheless, there are challenges related to latency (due to large quadratic programs/nonlinear programs) and model fidelity under human variability.

**Hierarchical Approaches** Stack-of-Tasks (SoT) methods organize whole-body control into a hierarchy of tasks, where each task is assigned a strict priority level. High-priority constraints (e.g., contacts, balance, joint/torque limits) are enforced first, while lower-priority tasks are projected into the corresponding null spaces, resulting in predictable behavior in the event of task conflicts. Foundational ideas trace back to operational-space control and task-priority Inverse Kinematics (IK) for redundant robots [89–91]. In modern humanoids, SoT is often realized as hierarchical QPs that solve a cascade of convex problems at each cycle, enabling real-time multi-contact behaviors while preserving hard constraints [92, 93]. These formulations have been demonstrated for dynamic whole-body motion with unilateral contact/friction constraints, as well as on torque-controlled humanoids with momentum regulation [94, 95].

**Sampling-based MPC (MPPI)** The Model Predictive Path Integral (MPPI) method bypasses large, nonlinear programs by sampling multiple rollouts each cycle and averaging the best results. This produces a gradient-free controller capable of tolerating nonconvex costs and contact discontinuities, making MPPI particularly suitable for contact-rich situations and rapid "what-if" re-planning during cooperative tasks [96, 97]. For pHHI, key considerations include the trade-off between sample budget and latency, cost shaping to account for human comfort and safety, and addressing stochastic variability. The sampling-based nature of MPPI could also be leveraged with learned coupled dynamics of the humanoid and the human. This could lead to robust motion planning and efficient collaboration.

**Differential Dynamic Programming (DDP)** Differential Dynamic Programming (DDP) is a method for solving optimization that alternates forward rollouts with a backward quadratic expansion of the value function, yielding both feedforward control updates and time-varying Linear Quadratic Regulator (LQR) feedback gains. For legged and humanoid systems, DDP (and first-order variants such as iterative Linear Quadratic Gaussian (iLQG) / iterative Linear Quadratic Regulator (iLQR)) has been used to generate and stabilize whole-body motions online (e.g., get-up, disturbance recovery), and to respect input bounds via control-limited updates [98, 99]. Recent formulations address contact-rich locomotion by optimizing multi-phase rigid-contact dynamics and even hybrid switching/impact events while handling friction and torque limits through augmented-Lagrangian and barrier techniques [100, 101].

#### 3.2.3 Learning-based control

With the rise of machine learning, many researchers are using these techniques for humanoid control, aiming to achieve more general or adaptive behaviors that are difficult to design manually. These approaches are particularly appealing for pHHI because human behavior is complex to model. In theory, a learning robot could adapt to individual human partners or acquire complex interactive skills through examples or trial and error.

**Reinforcement Learning (RL)** Deep Reinforcement Learning (RL) has enabled the development of whole-body locomotion and emerging loco-manipulation skills through extensive simulation training. Recent advancements include transformer-based locomotion policies for humanoid robots [60], sim-to-real curricula that enhance the robustness of hardware [61], and contact-rich loco-manipulation controllers for tasks like pushing and carrying boxes [64, 65]. The appeal of RL for pHHI lies in its ability to personalize actions and discover strategies; a policy can learn when to step, brace with the arms, or adjust grip and effort based on cues from a partner, which behaviors can enhance model-based designs [102].

We would like to direct readers to [103] for a detailed review of learning-based methods for humanoids. The authors categorize these approaches into two primary types: end-to-end frameworks [104–108], which directly map sensory inputs to control actions, and hierarchical frameworks [69, 109, 110], which decompose locomotion into layers such as high-level planning and low-level control. Importantly, they highlight how future research can leverage the complementary strengths of optimization- and learning-based methods within unified control frameworks. For example, learning can be made more sample-efficient by incorporating guided policy learning [104, 111, 112], where expert demonstrations or planner-generated trajectories are used to guide policy training in both end-to-end and hierarchical settings. Likewise, optimization methods can be made more robust to uncertainty and unmodeled dynamics by incorporating residual learning, where learned corrections augment model-based baselines. These strategies for fusing optimization and learning not only advance locomotion control but also provide inspiration for pHHI, particularly in developing coupled models, human-aware control, and improved robustness in collaborative tasks.

In practice, on-policy training methods (such as Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO)) offer stability but require a substantial amount of data, while off-policy methods (like Soft Actor Critic (SAC) and Deep Deterministic Policy Gradient (DDPG) can be more efficient. However, they are often challenging to fine-tune. This has led to a reliance on domain randomization and structured curricula to facilitate transfer from simulation to real-world applications. The main challenges for pHHI include safety, as unsafe exploration is unacceptable when humans are involved, and the sim-to-real gap. Possible solutions involve exposing policies only as high-level references, while simultaneously implementing impedance and walking behavior control (WBC) safety layers, using conservative action spaces, and shaping rewards to prioritize comfort, for example, by minimizing sudden movements and maintaining bounded forces.

**Imitation Learning (IL) & Inverse RL** Instead of starting from scratch, Imitation Learning (IL) utilizes demonstrations, such as teleoperation or motion capture, to acquire natural, human-like whole-body behaviors. This can then be optionally refined using RL for greater robustness. For pHHI, IL can capture socially appropriate timing, effective handovers, or partner-following strategies [66, 67, 69, 113]. Meanwhile, Inverse Reinforcement Learning (IRL) learns reward functions that align with human preferences, such as ensuring stability while minimizing abrupt movements, which enhances comfort during assistance [114]. However, there are significant challenges, including covariate shift (when the robot's experiences deviate from those observed during demonstrations) and limited exposure to edge cases. Possible solutions involve aggregating datasets and employing IL-to-RL pipelines to enhance robustness. Additionally, using uncertainty-aware execution can help; if the robot's confidence is low, it can revert to compliant tracking to maintain safety and reliability.

**Behavior/Policy Foundations, Language & Teleoperation** Recent research has investigated language-guided and "foundation-style" policy models that can compile and adapt whole-body skills from extensive datasets. This approach allows for rapid adaptation to new tasks with minimal prompts or few examples [115–118]. When combined with universal whole-body teleoperation, these systems can gather high-quality demonstrations and facilitate clear shared autonomy during contact-rich collaborations [69, 119]. In pHHI, such models help align robot behavior with human intent, utilizing either language or operator cues, while also preserving the subtleties of contact learned from diverse prior data. The trade-offs involved include scale (data and computational resources), assurance, and interpretability. In practice, these models are deployed alongside certified balance and safety layers, as well as hierarchical whole-body control, ensuring that high-level instructions are executed smoothly and predictably. Although this approach shows great potential for achieving outstanding performance in complex tasks, significant work remains to be done, particularly in avoiding hallucination, ensuring safety guarantees, and preventing Large Language Model (LLM)-Jailbreaking [120, 121].

#### 3.2.4 Stability/Safety Certificates: ZMP / CoP, Capture-Point / DCM, HZD, CBF-QP

When considering pHHI, a humanoid must ensure balance and safety in real-time, even if higher-level MPC or learning may hesitate or err. At the low-levels, techniques such as ZMP/CoP feasibility and DCM/Capture Point stepping provide rapid corrections in milliseconds, keeping the CoM in a recoverable state and enabling smooth push recovery. For continuous locomotion, Hybrid Zero Dynamics (HZD) guarantees stable gaits for underactuated bipeds, offering a solid foundation upon which regulation and disturbance rejection can be built [122–124]. As a runtime safety measure, Control Barrier Function Quadratic Programs (CBF-QPs) enforce forward-invariant sets, such as no-slip cones, CoM envelopes, and joint/torque limits, regardless of the nominal controller, filtering out unsafe regions during execution [78]. The main challenges include certifying the composition across multi-phase loco-manipulation, designing uncertainty-aware sets considering unknown human forces and sensor delays, and maintaining comfort by avoiding abrupt or overly cautious reactions. Promising avenues for improvement include: i) personalizing barrier sets with partner-specific stiffness and bounded force models, ii) implementing rate-limited barriers to control jerk and force rate, iii) co-designing with impedance and centroidal/WBC to minimize interventions from barriers, and iv) integrating barriers with learned uncertainty, tightening sets when models are weak and relaxing them when confidence is high, all while maintaining certified dynamics at the core.

### 3.3 Pillar I Insights and Discussion

A humanoid controller designed for pHHI should exhibit four key features: *safety*, to ensure human well-being during close physical contact; *stability and robustness*, to guarantee reliable performance under external disturbances and varying loads; *adaptability*, to accommodate diverse human partners, tasks, and environments; and *ease of implementation*, to enable practical deployment on complex humanoid systems. Guided by the humanoid control architectures discussed in Pillar I, we examine how three major control paradigms, classical, optimization-based, and learning-based methods, perform with respect to these key requirements for pHHI, as illustrated in Fig. 4.

The first feature considered is safety, defined as the protection of the human during controller operation. Safety encompasses the robot's compliance during human contact, robustness under disturbances or model mismatches, and the level of trust that the human perceives. Classical and optimization-based controllers generally provide comparable safety, as safety can be enforced by design and, in optimization-based methods, modulated through hard and soft constraints. Both approaches, however, remain vulnerable to model inaccuracies and external disturbances. Incorporating impedance or admittance models can mitigate these vulnerabilities. Classical controllers may produce more aggressive motions, while optimization-based controllers can address this issue. Learning-based controllers offer smoother motion but currently lack full explainability and interpretability, which reduces their perceived safety due to uncertainty regarding potential failures that could endanger humans. Additionally, it is challenging to add explicit safety constraints during policy learning. Constrained Proximal Policy Optimization (CPPO) [112] enables adding a threshold on the cumulative violation of all constraints, but lacks individual thresholds and hard constraints. Safety can also be ensured to some extent by guiding policy learning to avoid unsafe actions.

The second feature considered is stability and robustness. This accounts for the robot's balance and its robustness in compensating for or anticipating disturbances that could lead to a misstep or fall. In this category, we consider that learning-based methods outperform classical and optimization-based methods. The primary reason is that learning-based methods are more effective at adapting and compensating for unmodeled effects, switching states, and disturbances. An optimization-based approach could explicitly trade off balance, tracking, and effort, yielding whole-body responses or employing predictive methods that decide to step early (specifying where/when/how far) in order to maintain stability. Classical methods can be myopic, reacting only after errors have grown, often too late, and lacking predictive capability.

The third feature considered is adaptability, defined as the controller's ability to personalize to user preferences, manage changing arbitration among objectives, and generalize across tasks. Learning-based methods generally demonstrate superior adaptability, as learned policies can be conditioned on context and preference signals such as embeddings, demonstrations, or reward shaping. These policies can be updated online through residuals, enabling rapid personalization without redesign. Learning-based controllers also adapt effectively to changes in partners, payloads, or goals. Optimization-based controllers can adapt by re-solving with updated costs, priorities, and constraints or by incorporating online parameter estimates. However, each modification requires problem re-specification and incurs solver latency, while generalization is constrained by the selected model and task encoding. Classical controllers encode preferences implicitly through gains and heuristics, and significant changes in arbitration or task necessitate manual retuning and often re-architecting of the control loops.

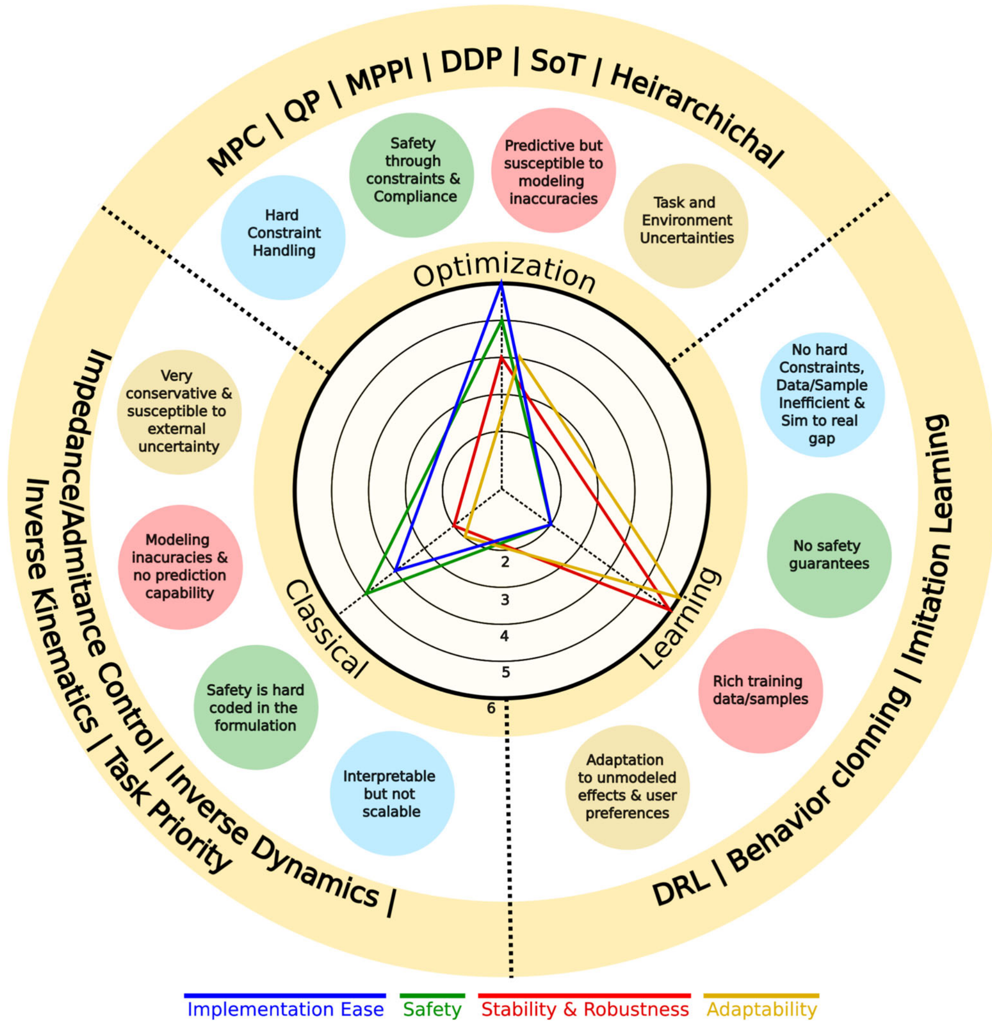


**Fig. 4** Conceptual Framework of Pillar I: Humanoid Modeling & Control. Classical, optimization-based, and learning-based control paradigms are compared using a radar chart that highlights safety, stability and robustness, adaptability, and implementation effort. The scores presented are qualitative and comparative rather than absolute measurements

The final feature considered is ease of implementation. Optimization-based approaches are generally preferred because they enable the formulation of tasks and safety requirements as costs and hard constraints, and benefit from the use of established commercial solvers. While tuning of weights and models remains necessary, these methods scale effectively as interaction complexity increases. Classical controllers provide transparency for individual behaviors, but their implementation becomes more complex as the number of contacts and arbitration scenarios increases, requiring additional gains, heuristics, or operational modes. Learning-based methods present greater implementation challenges, as they demand extensive safety-validated datasets or high-fidelity simulations, must address simulation-to-reality transfer and exploration risks, and often require supplementary safety filters because constraints are not inherently integrated.

## 4 Pillar II: Human Intent Estimation

A robust humanoid controller, as presented in the previous section, allows the robot to manage its dynamics and remain resilient to external disturbances. However, proactive physical collaboration with a humanoid robot depends on the robot's ability to anticipate its partner's next actions and the intensity with which they will act. In this section, we first define human intent, which encompasses everything from low-level desired trajectories and force profiles to higher-level goals, roles, and cognitive states. This clarification helps us understand what needs to be estimated and the timeframes involved. Next, we examine various sensing modalities for detecting intent, including kinematic and pose data, force-torque and tactile feedback, physiological signals, and gaze cues. We emphasize what each modality reveals and conceals in contact-rich settings. Based on these signals, we review models and algorithms for predicting intent. These include probabilistic filters, movement primitives, inverse optimal control/inverse reinforcement learning for inferring objectives, and Partially Observable Markov Decision Process (POMDP) formulations for actions taken under uncertainty. We also explore modern learning methods, including Recurrent Neural Networks (RNNs), Transformers, and generative models, with a focus on real-time applications, personalization, and calibrated uncertainty. Throughout our discussion, we highlight the strengths and limitations of these approaches in the context of pHHI. We conclude by indicating how intent estimates should be integrated with whole-body impedance, preview planning, and safety layers, setting the stage for the integration pathways developed in later sections.

### 4.1 Definition of Human Intent: From Desired Trajectories to Cognitive States

In the context of this review, "intent" refers to the underlying variables that help to explain and predict a partner's future movements and force profile. This includes understanding where they are heading, what they intend to do with shared objects, and the strength of their actions. At the most basic operational level, intent can be represented as a desired trajectory, essentially, the desired positions, velocity profiles, or force profiles (wrench envelopes). If a robot knows these trajectories, it can adjust its interactions and anticipate movements, such as steps or grasps. At a higher level of task planning, intent may represent specific goals (e.g., turning left, stopping, or handing over an object) or role assignments (such as leader or follower) that influence how individuals collaborate. At the cognitive level, intent is associated with plans, preferences, and social cues, such as gaze and timing, which are critical to building trust and enhancing legibility in interactions. Research on shared control and intention-based systems shows that making intent explicit can improve interaction fluency and reduce effort. However, it is crucial to maintain transparency and calibrated autonomy for user acceptance [18, 125, 126].

It is important to note that in robotics, intent must often be inferred in situations of partial observability, under tight deadlines, and in the presence of disturbances caused by contact. This intent then needs to be translated into whole-body behaviors that ensure comfort and safety, a requirement that presents challenges for both perception and control [14].

### 4.2 Sensing Modalities for Human Intent Estimation

A humanoid's ability to understand intent is limited by its sensory capabilities and the latency of its sensory input. Different sensory channels reveal various aspects of a person's underlying plan. For example, body posture can provide information about geometry and timing, but it may become unclear if the view is obstructed. Additionally, the forces exerted during contact can indicate immediate goals but may be confused by environmental factors that affect human intent. Physiological signals and gaze direction can offer cognitive-level clues that appear earlier than the corresponding motion, but they require careful calibration and pose challenges in terms of reliability.

#### 4.2.1 Kinematic and Motion-Based Sensing: Leveraging Body Pose and Skeletal Data

Kinematics is fundamental for intent inference in pHHI because an individual's pose and its changes over time provide strong indicators of near-future actions. Recent advances in markerless, multi-person tracking tailored for Human-Robot Interaction (HRI) have minimized jitter and occlusion artifacts, resulting in stable pose streams that downstream predictors can reliably use [127]. When direct contact sensing is not available, dependable 3D pose baselines and monocular human dynamics estimation provide practical inputs for early intent decoding [128, 129]. Using these pose streams, lightweight predictors can quickly identify imminent actions (e.g., a pivot during co-carrying) to facilitate timely adjustments in impedance or planning. In walking scenarios, intent recognition has been framed as a multi-class problem (continue, turn, stop), incorporating phase-aware features to trigger balance-sensitive assistance [130]. Building on this foundation, vision-based skeleton tracking with RGB-D (i.e., images with both color (RGB) and depth (D) information) or monocular cameras offers a valuable, non-contact means of understanding posture, velocity, and configuration, supporting proactive behaviors. By observing the early stages of motion, the robot can forecast short-horizon trajectories [131, 132], allowing it to anticipate user needs before physical contact occurs. This capability enables proactive handovers and preemptive collision avoidance [133–135]. In pHHI, these predictions can be directly integrated into whole-body or impedance controllers, helping to shape interaction forces while maintaining balance.

However, limitations exist; kinematic data can be ambiguous without context related to objects or forces (the same motion may have different underlying goals), and close-quarter interactions can heighten issues like occlusions and sensitivity to lighting [131, 132, 136]. To achieve cohesion and robustness in practice, several strategies could be implemented, such as tracking techniques (including temporal priors, kinematic constraints, and multi-view/depth systems) to stabilize pose and contact-aware features that align predictions with the timing of tasks.

#### 4.2.2 Force and Torque-Based Sensing: Direct Physical Cues

Haptic measurements, specifically the interaction forces and torques at the contact point, provide the clearest insight into a partner's short-term intentions, particularly in co-manipulation scenarios. End-effector and joint force/torque sensors, or estimators derived from motor currents, effectively capture human push/pull patterns, desired direction changes, and load-sharing behaviors with high fidelity and low latency. This generates an immediately actionable signal for impedance/admittance and whole-body control systems [14, 137]. Learning from these signals enhances shared object manipulation by mapping force histories to necessary motion corrections and role allocations. In co-transport situations, combining recent kinematics with wrench data enables rapid estimation of desired velocity and force direction. Additionally, constraint-aware formulations ensure that inferred goals remain physically feasible for the human–object–robot system [138]. When dedicated sensors are not available, several studies demonstrate the recovery of virtual interaction wrenches from the robot's dynamics (specifically, torque residuals), retaining much of the functionality for intent decoding and compliance shaping [139, 140].

While the strengths of this modality, such as immediacy and actionability, are significant, they also highlight its limitations. Pure force sensing is inherently reactive; the robot can only learn from the effort applied, which can result in sluggish responses or require unnecessary human exertion. Practical challenges include sensor placement, bias and noise, contact switches that complicate estimation, and the difficulty of separating human intent from environmental disturbances. Cohesive design approaches help mitigate these issues by implementing filtering and wrench observers linked to the contact state or object state data to clarify intent.

#### 4.2.3 Physiological and Gaze-Based Sensing

Physiological signals and gaze patterns provide anticipatory evidence of human intent, often preceding overt motion. Eye tracking reveals where a person's visual attention is focused and which targets they are selecting, strong predictors of the next reach, route, or grasp. This information enables actions such as pre-positioning, safer handovers, and earlier collision avoidance [141]. Research on handover timing and gaze indicates that well-timed cues can enhance compliance and perceived fluidity during interactions. These cues can be integrated into intent estimators to better coordinate role exchanges and initiatives during close contact [133, 134]. On the physiological side, ElectroMyoGraphy (EMG) captures muscle activation before any measurable force is generated. Additionally, wearable Inertial Measurement Units (IMUs) can segment activities, detect gait phases, and trigger short-term motion primitives that are generally applicable across users [142, 143]. In pHHI, these early signals can guide whole-body and impedance controllers to anticipate the most likely action.

The advantages of these channels lie in their ability to reduce lead time and clarify tasks. However, there are practical and user-dependent challenges: eye trackers and EMG can be intrusive and require calibration. At the same time, signals can vary significantly between users, depending on factors such as placement and context. Additionally, performance can degrade due to fatigue, perspiration, or changes in lighting and occlusions [144–146]. Consequently, cohesive designs treat these modalities as supporting evidence rather than definitive measurements. They fuse data from physiological signals, gaze, kinematic, and force information using probabilistic filters.

#### 4.2.4 Multi-Modal Fusion

No single sensor is reliable across the wide range of pHHI scenarios, so effective intent understanding and safe assistance increasingly depend on sensor fusion. Combining visual data with force and tactile cues consistently enhances continuous intent recognition and attention estimation compared to using a single modality, thereby reducing false triggers in tasks such as hand-guiding and shared carrying [147]. Markerless tracking, tailored for HRI, which incorporates temporal priors, multi-view perspectives, and kinematic constraints, stabilizes the visual component in cluttered environments [127]. Meanwhile, gaze information helps clarify which target a reach is directed at, while wearables like EMG sensors and IMUs provide phase and early activation cues. The outcome is the gathering of complementary evidence at various latencies and levels of abstraction: for instance, gaze can clarify the intended target of a reach, vision and kinematics may predict motion, and force measurements can confirm physical intent upon contact.

In practical applications, early cues (such as gaze, EMG, and IMU signals) serve as priors that inform predictions and control set points, while haptic feedback and kinematics provide quick posteriors to effectively close the control loop. The potential benefits of this approach are improved disambiguation and fault tolerance. However, challenges include managing latency and computational demands, as well as the necessity for reliability-aware attention and online calibration/personalization. Key open issues encompass adapting fusion weights based on different users and contexts, addressing sensor bias and drift during contact transitions, and ensuring comfort and legibility (for example, through rate-limiting updates to stiffness or wrenches) as beliefs evolve. A practical approach involves stabilizing vision, fusing gaze, EMG, and IMU data as intent priors, confirming with force and torque evidence, enforcing feasibility, and communicating confidence levels to the controller. This allows the system to balance between reactive compliance and proactive assistance smoothly [144–146].

## 4.3 Models and Algorithms for Human Intent Prediction

Once sensing collects information to expose parts of a human's plan, the main challenge is to transform these insights into actionable forecasts. Specifically, predicting what the human will do next, the timing and intensity of their actions, and our level of confidence in these predictions. Here, we discuss some of the models and algorithms most frequently used to capture human intent.

### 4.3.1 Probabilistic Models: Hidden Markov Models (HMMs) and Bayesian Filtering

Probabilistic intent modeling views human goals as latent stochastic processes that generate observed motion and force. Classical HMMs and recursive Bayesian filters maintain a belief over candidate intentions or phases and update this belief online as new sensory evidence arrives (such as from vision, wrenches, gaze, and EMG). This process yields an interpretable posterior that controllers can use for risk-aware decisions [136, 148]. On the trajectory side, probabilistic movement primitives (ProMPs) and Gaussian mixture or regression models encode distributions over motions. These models enable early action recognition and short-horizon forecasts from partial observations. Phase and speed estimators condition a motion library online for collaborative tasks, even with limited data [148–150].

The advantages of these models are threefold: they are data-efficient, provide explicit uncertainty (calibrated confidence for safe decision-making), and facilitate reflex-rate inference that interacts with impedance, WBC, or MPC. However, they also have some weaknesses: they depend on task- or library-specific motion sets, have limited prediction horizons, and can be sensitive to personalization, contact reconfiguration, and high-dimensional continuous dynamics.

### 4.3.2 Partially Observable Markov Decision Processes (POMDPs)

When intent is hidden and observations are noisy, the model naturally fits a Partially Observable Markov Decision Process (POMDP). In this framework, the human's intent is considered a latent state, while the robot maintains a belief over hypotheses derived from multimodal evidence. It chooses actions that both advance the task and reduce ambiguity. In scenarios like collaborative transport or handovers, this approach results in policies that gently probe the human's intent, such as slight object reorientation or compliant micro-motions, to encourage clarifying responses from the human while adhering to strict safety limits. This reflects findings on implicit communication through motion [151]. POMDPs also formalize the concept of belief-dependent controllers, which can adjust their impedance, being softer when uncertainty is high and firmer when confidence is high. However, exact POMDP solutions are complex for humanoid robots, as beliefs are often high-dimensional and continuous. Therefore, systems often rely on approximations that retain many advantages, such as fast Bayesian filters to generate intent or phase beliefs [150]. These filters are paired with belief-aware receding-horizon controllers that encode intent as constraints, costs, or scenarios [18].

### 4.3.3 Inverse Optimal Control (IOC) and Inverse Reinforcement Learning (IRL): Inferring Human Objectives

Humans tend to act approximately optimally with respect to a hidden objective, and researchers in IRL seek to recover that objective from demonstrations. This approach focuses on translating intent into a cost landscape rather than simply making kinematic predictions [138]. In the area of collaborative loco-manipulation, IRL has been effectively used to derive rewards that align with human preferences, such as stability, smoothness, and effort [114]. This results in robot behaviors that users perceive as more natural and comfortable [152]. Another perspective views physical interaction as a form of communication, where the robot updates its objective in real-time based on human feedback. This co-active learning process helps the robot adapt its strategies to reflect the user's implicit goals during the task [153]. The semantics-first approach to intent in pHHI facilitates generalization to new contexts, whether involving unfamiliar objects, paths, or role distributions.

However, this approach presents challenges. Many different objectives can explain the same motion (a problem known as a lack of identifiability), and the features involved are important. Additionally, data collection can be limited, and full IOC and IRL can be computationally intensive, particularly under conditions of partial observability and complex, contact-rich, whole-body dynamics. To address these challenges, practical strategies include: structuring rewards using physically meaningful features, such as ergonomics, effort, and contact safety, accompanied by regularization constraints. Combining offline IOC with online, co-adaptive updates based on corrections made during tasks to enhance personalization.

### 4.3.4 Machine Learning and Deep Learning Approaches

More recently, deep learning has become the dominant approach, offering powerful tools for learning complex patterns from large datasets without the need for extensive feature engineering.

**Recurrent Neural Networks (RNNs, LSTMs, GRUs) for Temporal Dynamics** Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) networks and Gated Recurrent Unit (GRUs), effectively utilize the sequential structure in pose, wrench, and gaze data streams [154]. This enables them to recognize intent and forecast short-term movements more accurately than traditional methods [155]. In the context of pHHI, these models support early intent detection (such as turning, stopping, or continuing while walking) by employing gait-phase-aware classifiers. These classifiers can trigger balance-sensitive assistance and control preview based on detected intent [130]. RNNs are trained on sliding windows of data and deployed in real-time, providing low-latency predictions. These predictions can be integrated into impedance, momentum, or WBC and QP controllers, allowing for pre-alignment of frames, biasing of stiffness in the predicted direction, and scheduling of stepping or handover timings. This results in proactive yet easily interpretable assistance.

Their advantages include strong temporal expressiveness, moderate computational requirements, and ease of integration into real-time systems. However, they also exhibit some weaknesses, such as sensitivity to dataset biases and domain shifts, performance drift over longer prediction horizons, and instability during phase or contact transitions when there is no explicit structural guidance. To create cohesive designs, the following strategies are recommended: First, incorporate phase and contact information (e.g., phase channels or event flags) to stabilize transitions. Second, pair RNNs with physics-based priors (e.g., Probabilistic or Dynamic Movement Primitives), which allows the network to learn residuals. Third, implement multi-task heads (for classification and trajectory prediction) with curriculum or scheduled sampling to reduce compounding errors.

**Transformer Models: Conditioning on Robot Actions for Enhanced Prediction** Transformers use attention mechanisms to capture long-range dependencies and make robust early predictions based on partial histories. They have been employed to anticipate human actions in collaborative settings [156], predict forces and velocities from video during co-transport, allowing robots to proactively adapt to load changes [157, 158], and tokenize wearable IMU streams into motion primitives that generalize across different users, yielding latent "intent tokens" that interface smoothly with control systems [159]. Generative models provide a way to achieve more human-like and diverse motion prediction. Rather than predicting a single future trajectory, models such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) can learn a distribution of plausible future motions [160, 161]. This capability enables robots to consider multiple potential human actions. The strength of these models lies in their expressiveness and their ability to capture the multi-modal nature of human behavior. However, they face challenges related to controlling the generation process to align with specific task constraints and ensuring the stability of their training. Compared to recurrent models, the attention mechanism connects distant cues (for example, an early gaze or pose shift to a later force peak). Thus, enhancing short-horizon forecasts that can be utilized by whole-body or impedance controllers for timing, stiffness adjustments, and safe anticipatory actions. A crucial area of development is conditioning on the robot's intended actions and compliance state. By inputting the robot's future motions (or candidate plans) alongside human signals, the predictor can address the question, "*What will the human do in response to this robot motion?*". This approach reduces ambiguity and produces consistent forecasts, which are better suited for shared autonomy, such as anticipating the timing of handovers, adjusting foot placements during co-carrying, or facilitating proactive load sharing.

The practical challenges of this approach are well recognized: transformers are data-hungry and compute-intensive, and conditioning on high-dimensional robot states while ensuring real-time performance and calibrated uncertainty remains complex. Emerging strategies include pairing lightweight transformer heads (such as streaming or sparse attention and sliding windows) with distilled or frozen backbones. Additionally, pretraining on simulations or logs with domain adaptation, exposing uncertainty for assistance gating, and implementing a certified execution layer underneath can help manage residual model errors.

**Generative Models (VAEs, GANs, Diffusion Models) for Expressive Motion and Scene Understanding** Generative models provide a way to achieve more human-like and diverse motion prediction. Rather than predicting a single future trajectory, models such as Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs) can learn a distribution of plausible future motions [162]. This capability enables robots to consider multiple potential human actions [131]. The strength of these models lies in their expressiveness and their ability to capture the multi-modal nature of human behavior. However, they face challenges related to controlling the generation process to align with specific task constraints and ensuring the stability of their training.

### 4.4 Pillar II: Insights and Discussion

Practical human intent estimation requires a careful matching of sensing modalities and computational models to the specific demands of the interaction. Figure 5 considers the

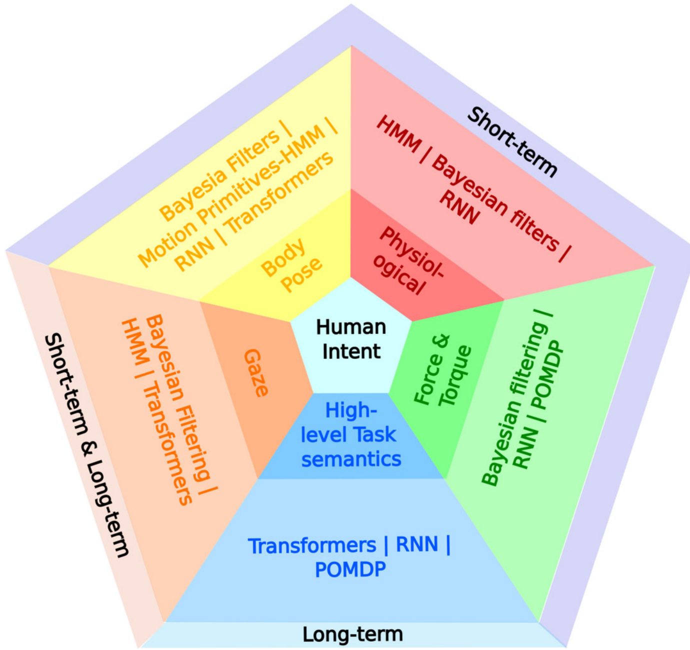


**Fig. 5** A map of human intent modalities, where the outer layer lists the most commonly used methods for each signal type, organized by their effective time horizon. Short-term physical cues (e.g., body pose, force/torque) are matched with reactive models, while long-term semantic cues (e.g., gaze, task goals) are paired with predictive, goal-aware algorithms

discussed intent modalities (Body pose, physiological, Force and torque, high-level task semantics, and gaze) and classifies them for both short-term and long-term intentions. Additionally, it highlights the more common methods and algorithms that can be utilized more effectively for these modalities. The core challenge lies in bridging different time horizons: from immediate, reflex-rate predictions that guide reactive control to long-term, goal-oriented forecasts that enable proactive assistance. Then, Table 3 classifies the sensing methods or scopes discussed in this section, describes what is estimated, and outlines their advantages and limitations.

The selection of sensing modality is closely related to the temporal characteristics of intent. Short-horizon intent, including next-step prediction or minor force adjustments, is effectively captured by high-bandwidth modalities such as body pose and force or torque measurements. These data streams deliver phase-specific information suitable for lightweight, uncertainty-aware models, including HMMs, Bayesian filters, and Probabilistic Movement Primitives (ProMPs). These approaches provide rapid controller updates, supporting impedance modulation and whole-body controller previews at reflex rates. For low-latency sequence prediction, RNNs are mainly preferred since they efficiently generate short-term forecasts from these dense data streams.

Long-horizon intent estimation focuses on inferring task-level goals such as target selection or role arbitration and relies on modalities that provide rich semantic information. Gaze serves as an early indicator of human focus and anticipated objectives, while high-level task semantics establish the context for interaction. Sophisticated computational models are required to interpret these cues. Methods, including IOC and IRL, are designed to infer underlying objectives, while POMDPs explicitly represent beliefs and uncertainty. Recently, Transformer architectures have demonstrated strong performance by using attention mechanisms to learn long-term, goal-oriented dependencies from sequential data. These semantic-based models are effective for generalizing to novel tasks because they capture the rationale behind actions rather than only surface-level features. As a result, they generalize more reliably across different objects and environmental configurations than approaches based solely on trajectory matching.

Robust performance across a diverse user population requires system customization. Physiological signals, such as EMG and IMU data, exhibit significant user specificity in intent predictions, emphasizing the need for personalization. For these signals and for gaze data, per-user calibration is typically necessary. Techniques such as few-shot normalization or online adaptation can further improve predictive accuracy. Personalization in modeling may involve adding lightweight residual adapter layers to RNN or Transformer models to reflect individual timing and stiffness preferences. An effective and generalizable system should integrate multiple sensing modalities, utilizing early indicators from gaze or EMG as priors and confirming predictions with late-arriving posteriors from pose and force measurements.

**Table 3** A comparative analysis of intent estimation algorithms, highlighting their primary sensing requirements, output types, and key trade-offs

| Sensing / Scope | Model / Algorithm | What is predicted/estimated | Advantages | Limitations / Risks | Rep. Refs. |
|---|---|---|---|---|---|
| Sensing and Fusion | | | | | |
| Multi-modal (vision, wrench/FT, tactile, gaze, EMG/IMU) | Fusion (stabilized vision, gaze/EMG/IMU priors; force/kinematics posteriors) | Intent priors/posteriors; contact events; complementary cues at different latencies | Improves disambiguation and fault tolerance; clarifies targets and timing | Latency/compute management; sensor bias/drift at contact transitions; need personalization | [126, 141, 144–146] |
| Probabilistic and Decision-Theoretic Models | | | | | |
| Pose / force / gaze / EMG streams | HMMs & Bayesian filtering | Belief over goals/phases; early action recognition; short-horizon traj./force forecasts | Data-efficient; explicit uncertainty for safe decisions; reflex-rate inference | Library/task-specific motion sets; limited horizons; sensitive w.r.t. personalization/contact reconfig/high-dim dynamics | [136, 147–150] |
| Belief over hidden intent under uncertainty | POMDPs | Belief state + probing actions to reduce ambiguity; impedance adjusted by confidence | Formalizes belief-aware control; gentle probing (micro-motions) consistent with implicit communication | Exact solutions intractable (continuous, high-dim beliefs); use approximations (fast filters + belief-aware MPC) | [18, 151] |
| Objective Inference and Learning | | | | | |
| Demonstrations / human feedback | IOC / IRL | Reward/cost aligned with preferences (stability, smoothness, effort); objective updated online | Behaviors perceived as natural/comfortable; semantics-first aids generalization | Identifiability; limited data; computational cost under partial observability & contact-rich dynamics | [114, 138, 152, 153] |
| Sequential multi-modal data | RNNs / LSTMs / GRUs | Early intent detection; short-term movement forecasts; gait-phase-aware triggers | Better accuracy than classical baselines; low-latency sliding-window inference | Data hungry; distribution shift; needs calibrated uncertainty | [130, 154, 155] |
| Sequences, video, IMUs (robot-conditioned) | Transformers (action-conditioned); intent tokens | Anticipate actions; predict forces/velocities from video; tokenize IMUs into motion/“intent” tokens | Capture long-range dependencies; improved short-horizon forecasts when conditioned on robot plans | Compute/data intensive; real-time + calibrated uncertainty is challenging | [156–158] |
| Trajectory/scene futures | Generative models (VAEs, GANs, Diffusion) | Multi-modal future samples (plausible human actions) | Expressive; captures multi-modal human behavior | Training stability; aligning generation with task constraints | [130, 159] |
| Demonstrations; large datasets; teleop | Imitation Learning / RL; foundation policies; language + teleoperation | Policies that compile/adapt whole-body skills; few-shot adaptation | Rapid adaptation; clear shared autonomy with teleop; leverage diverse prior data | Scale / interpretability; avoid hallucination; LLM safety (jailbreaking) | [66, 67, 69, 113, 115–121] |

# 5 Pillar III: Human Models

The first two pillars enable the robot to control its body and infer its partner's goals. However, to elevate interactions from merely functional to truly safe, efficient, and intuitive, the robot needs a deeper comprehension of the entity it is interacting with. This requires a computational model of the human, essentially a framework for understanding the physical principles that govern human movement, capabilities, and limitations [163]. This third pillar goes beyond treating humans as generic sources of intent signals. Instead, it models them as complex cognitive biomechanical systems [164]. The remainder of this section reviews human models across various categories, including musculoskeletal and physics-based bodies, reflex and motor-control descriptions, data-driven motion priors, and cognitive/decision abstractions.

## 5.1 Musculoskeletal and Physics-Based Models

To effectively reason about a human partner beyond mere surface kinematics, pHHI could benefit from models that expose the body's internal mechanics, ranging from muscles to whole-body dynamics [165, 166]. At the high-fidelity level, OpenSim offers reusable musculoskeletal models, estimation tools, and validation pipelines for neuromuscular control and movement analysis [167]. Notably, comprehensive full-body models, such as a gait model, incorporate numerous skeletal Degrees of Freedom (DOFs) and extensive sets of muscle actuators. These models generate joint kinetics and muscle activations that align with experimental data, establishing them as standard references for analysis and controller-in-loop studies [168].

Conversely, at a more abstract but faster level, rigid-body human models with joint limits enable real-time whole-body simulation for virtual reality, teleoperation, and control-theoretic analysis [144, 169]. This makes them practical interfaces for robot planners and safety filters [170]. These physics frameworks seamlessly connect to humanoid balance. Empirical evidence shows that humans regulate whole-body angular momentum during walking, which provides principled targets and constraints for momentum-aware collaboration and shared load handling [171]. Complementary interactive locomotion studies, where two people physically pair while walking, illustrate how coupling forces, cadence, and phase coordination develop, offering datasets and dynamical patterns that can be emulated or anticipated in human-robot interactions [172].

For pHHI, the benefits of these models are twofold. First, musculoskeletal models can predict unmeasured internal variables, such as muscle forces and joint contact loads. This enables robots to plan contact behaviors that minimize discomfort or risk and to incorporate ergonomic costs directly into their assistance policies. Second, physics-based human dynamics offer physically plausible motion priors under perturbations, which enhance short-horizon posture forecasts and inform balance and momentum controllers with human-consistent constraints, such as allowable angular momentum changes. The main challenges in this area are personalization and latency. Accurately identifying an individual's bone geometry, soft tissue, and muscle properties requires a labor-intensive process, making it challenging to tailor models efficiently. Even when such models are developed, forward musculoskeletal simulations are rarely fast enough to run in real time, further limiting their practical use.

## 5.2 Reflex and Motor-Control Models

Pure biomechanics alone cannot fully explain how people adapt during contact. However, neuromechanical models help bridge that gap by linking reflex pathways and task-level feedback to observed motion. Studies have demonstrated that compact muscle-reflex loops, which are tuned to legged mechanics, can accurately reproduce human walking kinematics and EMG patterns. This provides evidence that reflexes not only stabilize gait but also redistribute effort across different phases [173]. Reflex gains are also dependent on the task at hand. For instance, humans tend to enhance "position-like" feedback when focusing on positional goals and shift toward "force-like" feedback for tasks involving force tracking [174]. This suggests that our reflexes can adapt flexibly in response to context. In the upper limb, endpoint stiffness is described by a stiffness ellipsoid, whose orientation depends on posture and whose size is influenced by co-contraction. This creates reduced-order maps that link configuration and activation to interactive compliance [175–177]. Together, these models allow us to interpret posture, phase, and activation as actionable insights into how a partner might respond during contact.

In pHHI, these insights translate directly into compliance scheduling. If the posture and phase suggest a low-stiffness human limb or whole body, the robot should absorb more disturbances and adjust its admittance accordingly. Conversely, if the human is bracing (indicating higher co-contraction), the robot can share more of the load or increase its impedance without causing unexpected reactions. These cues can also guide momentum and whole-body controllers, ensuring that interaction forces are predictable and maintaining stable balance. However, the challenges are expected: there is significant variability both between different individuals and within the same individual, and estimating internal states (like co-contraction) in real-time without wearables is difficult.

## 5.3 Data-Driven Human Motion Models

A complementary approach to understanding human motion involves learning directly from data, evolving from early statistical models to modern deep generative methods. Classic techniques fittingly employ low-dimensional manifolds and Markovian transitions based on motion capture data, resulting in straightforward predictors and synthesizers of likely motion patterns [178]. Contemporary models, however, enhance both expressivity and robustness. For instance, diffusion models can generate diverse and natural motion sequences based on text or action prompts [179]. Additionally, HuMoR [184], a conditional variational prior, effectively stabilizes pose tracking and prediction even in the presence of visual disturbances and occlusions. DeepMimic [180], for example, learns controllers that replicate a broad spectrum of human skills in simulation while maintaining credibility despite disturbances. Beyond isolated movements, human-scene interaction models incorporate contacts and affordances, enabling the learning of behaviors limited by environmental constraints that respect object dynamics [181, 185, 186].

Recent research has begun to merge statistical methods with physics-based approaches to enhance both temporal consistency and contact realism. This integration is particularly promising for interaction forecasting, where contact plays a critical role in determining dynamics. Finally, behavioral studies show that humans often seek to minimize task-specific costs, which supports the use of learned or handcrafted objective structures to improve prediction and planning [182].

### 5.4 Cognitive and Decision Models

Beyond the physical execution of movement, a comprehensive human model must also account for the cognitive processes that drive behavior. Cognitive and decision models aim to capture how humans perceive, reason about, and make choices within collaborative tasks [164]. This includes modeling how a human mentally represents a task, often as a hierarchy of goals and sub-goals. By maintaining a belief over the human's current position in this task plan, the robot can better predict not just the next movement, but the next logical step in the overall process.

Furthermore, these models can incorporate principles from decision theory to understand how humans make choices under uncertainty or trade off between competing objectives, such as speed and accuracy. The emerging frontier in this domain is the use of LLMs as a proxy for human cognitive reasoning [183]. By leveraging the world knowledge and common-sense reasoning embedded in LLMs, a robot can infer a human's high-level plan from sparse instructions or contextual cues, enabling a more proactive and semantically aware form of collaboration. The primary challenge is grounding these abstract cognitive models in the robot's real-time sensory data and ensuring their reasoning is both safe and contextually appropriate.

Within a deployable pHHI stack, cognitive and decision models should not operate independently from control. Instead, they must provide explicit signals that directly influence planning and safety. For instance, inferred role or goal hypotheses serve as priors for intent belief and determine planner modes, such as lead versus follow or assist versus hold. Cognitive states, including fatigue, attention, and confusion, adjust cost weights and permissible levels of assistance. Additionally, discrete decisions, such as confirm, assist, or stop, initiate mode switches that the safety-certified whole-body controller can continuously enforce. This interface-based integration allows cognitive processes to enhance collaboration while maintaining safety constraints.

### 5.5 Pillar III: Insights and Discussion

Pillar III translates human factors into controller-ready constraints and costs, turning ergonomics into a live interface the robot can trust. It exports feasibility sets, such as joint-load limits and admissible wrenches, along with comfort/effort costs and task targets, so the robot can share work while staying within human limits. In our framework, the pillar is anchored by ergonomics, fidelity, and personalization, and it complements control (Pillar I) and intent (Pillar II) by supplying constraints and objectives that remain consistent across the stack. Table 4 summarizes the section organization by relevant model families, what they export to control, the inputs they require, and the strengths and limitations.

Choosing the right level of fidelity is central to making this interface useful. High-fidelity musculoskeletal models (e.g., OpenSim-style bodies) expose internal variables, such as muscle forces and joint contact loads, which are invaluable for ergonomic costs and safe load sharing. However, they can be slow and sensitive to modeling assumptions. Equally important, the human model must compute in real-time so that the exported feasibility sets and costs accurately reflect the person's shifting state. Relying on outdated estimates reduces the quality of assistance and can violate comfort or safety limits, so the human model must update at control-relevant rates while intensive inference operates in the background. In contrast, rigid-body human models run in real time and pair naturally with balance-aware planners, although with less physiological detail. A pragmatic strategy is to "budget" fidelity: deploy physics-grounded limits and targets online (e.g., feasible wrench envelopes, centroidal momentum targets), while reserving full musculoskeletal inference for offline calibration or slower background updates. This division captures the benefits of load sharing while acknowledging the risks associated with model mismatch and forward-simulation latency.

Because people vary both across users and over time, personalization must shift from population averages to

**Table 4** Computational Human Models for pHHI (Pillar III): exports to control, required inputs, strengths, limitations, and representative references

| Model family (Pillar III category) | Exports to controller (constraints / costs / targets) | Required inputs & parameters | Strengths / when to use | Limitations & risks | Rep. refs |
|---|---|---|---|---|---|
| Musculoskeletal & physics-based bodies | Joint-load limits; feasible wrench/contact sets; centroidal momentum targets; effort/comfort costs | Link properties; contact assumptions; human morphology; physics model configuration | Physically grounded limits / targets; informative for load sharing and balance-aware co-manipulation | Personalization and latency; model mismatch; forward simulations can be slow | [167–169, 171, 172] |
| Reflex & motor-control (neuromechanical) | Endpoint stiffness ellipsoids; impedance / feedback targets; task-dependent "position-like/force-like" gains | Posture / phase; activation/co-contraction cues; reduced-order stiffness maps | Directly informs compliance scheduling and whole-body shaping for safe, predictable interaction | Inter-/intra-person variability; challenging real-time estimation without wearables | [173–176] |
| Data-driven motion priors | Fast predictors of human motion/forces; priors for costs/constraints and short-horizon targets | Datasets/logs (pose, wrench, video); trained generative/sequence models | Real-time surrogates that compress complex dynamics; multi-modal forecasts | Distribution shift; need calibrated uncertainty for safe use | [178–181] |
| Cognitive/decision abstractions | Preference-weighted costs; "fairness" and comfort/effort-sharing objectives | User preferences, task goals, comfort envelopes, and summaries passed as constraints/weights | Aligns robot behavior with user comfort and task goals; supports arbitration at the planning level | Ambiguous trade-offs; requires careful calibration of preferences | [163, 164, 182, 183] |

individual profiles. To support this, per-user calibration of upstream signals (EMG, IMU, gaze) enhances downstream predictions and should be a routine practice, whether achieved through few-shot normalization or lightweight online adaptation. In turn, structural personalization exposes individualized feasibility and cost parameters to the controller: endpoint stiffness ellipsoids, joint-load limits, and effort/comfort terms that reflect each partner's capabilities. With these in place, Pillar I controllers can re-weight assistance and impedance on the fly, aligning help with the person's current state rather than an average model.

# 6 Discussion, Challenges, and Future Directions for pHHI

This review has examined the state of the art in pHHI through three foundational pillars: (i) humanoid modeling and control, (ii) human intent estimation and prediction, and (iii) computational models of human behavior. This approach attempts to tackle three challenging problems. First, the control of a floating, hybrid, multi-contact body in real time. Second, the inference of a human goal that is inherently latent and time-varying, all under conditions of partial observability. Finally, the reasoning about the human partner not merely as a source of signals, but as a biomechanical system with its own limitations, preferences, and sensitivities to risk. While significant advancements have been made within each of these pillars, the field has yet to achieve the seamless, intuitive, and robust interaction envisioned for the widespread deployment of humanoid robots in human-centric environments. This section discusses the primary challenges hindering progress, explores the reasons behind the difficulties in creating a common framework, and outlines future directions and research paths that could lead to breakthroughs in the field.

## 6.1 Principal Identified Challenges in pHHI

The development of seamless and robust pHHI is hindered by a set of deeply interconnected challenges that span multiple temporal, abstract, and interaction scales. These are not isolated algorithmic gaps but rather a reflection of the fundamental difficulty of creating predictable, safe, and effective physical collaboration between a robot and an adaptive, uncertain human partner. The following sections elaborate on the principal identified challenges that the field should currently overcome.

### 6.1.1 Time-Scale Management for Real-Time Operations

pHHI requires the integration of subsystems that operate on significantly different time scales. High-frequency control loops for torque and impedance must operate at rates of thousands of hertz, while mid-level planners for locomotion and multi-contact maneuvers operate at rates ranging from tens to a few hundreds of hertz. In contrast, intent estimators typically provide information at slower rates, such as those from video or wearable sensors, and important human biomechanical factors, such as fatigue and comfort, change over timescales of seconds to minutes. A significant challenge is designing multi-rate architectures that can maintain stability and effectiveness when high-level estimates are received intermittently or with varying latencies. An effective system must intelligently determine which variables can be reliably controlled at specific rates and implement well-considered fallback strategies. These strategies may include tightening safety constraints or softening impedance in response to delayed predictions, ensuring that the robot's actions remain safe and predictable for the human partner.

### 6.1.2 Dynamic Abstraction and the Curse of Dimensionality

The current model-based frameworks for humanoids involved in pHHI follow a hierarchical approach, dividing it into two tasks: high-level planning using abstracted models and low-level controllers that execute these high-level plans on the humanoid. The motivation for such abstracted models comes from the computationally vast state-decision space, which includes body pose, contact modes, footstep sequences, and grasp wrenches. These abstractions are often overly simplified. For example, authors in [28] extend the LIP models to account for interaction, but this typically comes at the cost of fidelity. As a result, notable discrepancies arise between the abstracted model and the actual system dynamics. This often results in motion plans that are inefficient for the task. For instance, they may struggle to achieve compliance with the co-transported object and fail to fully exploit the humanoid's capabilities. For model-based controllers, there is thus a pressing need for abstractions of the coupled dynamics in pHHI that are both computationally efficient and more faithful to the underlying system. Data-driven and learning-based approaches offer promising avenues for developing such abstractions often referred to as *latent representations*, but they remain largely unexplored. Once established, these abstractions can be integrated into model-based methods, such as MPPI, or used within model-based reinforcement learning frameworks.

### 6.1.3 Heterogeneity of Interactions and Autonomy Modes

The scope of pHHI is extensive, encompassing both direct and indirect interactions, as well as shared autonomy in complex tasks. Each type of interaction presents its own unique challenges regarding observability, potential failure modes, and social expectations. Therefore, a common framework should focus on standardizing the interfaces between components, defining terms such as “intent belief," “human feasibility model," and “safety invariant", while also allowing for mode-specific implementations of sensors, predictors, and constraints.

#### 6.1.4 Assurance, Interpretability, and the Sim-to-Real Gap for Learned Components

Data-driven methods have become essential in all three pillars of pHHI. However, interactions that involve contact can be challenging, and infrequent prediction errors can lead to significant consequences. Providing end-to-end safety assurances when learned modules are involved remains an unresolved issue. The sim-to-real gap further complicates this problem, as simulators often struggle to accurately model the intricacies of contact physics and the variability of human behavior. Finally, the interpretability of these learned modules is crucial for humans to trust and effectively collaborate with robots.

#### 6.1.5 Missing Datasets for Collaborative Tasks

Progress in pHHI is significantly limited by the absence of large, standardized datasets that capture whole-body, bidirectional human–robot interaction. Most existing datasets focus on vision or kinematics for individual agents, providing limited contact information, sparse or noisy force and torque signals, minimal intent annotation, restricted task diversity, and homogeneous participant groups. These limitations hinder fair benchmarking and impede the development of learning methods that require comprehensive supervision to model intent, predict contact transitions, ensure safety and comfort, and generalize across different partners and tasks. An effective dataset should be multimodal and time-synchronized, incorporating motion capture or IMUs, joint states, interaction wrenches, tactile or pressure data, gaze, voice, EMG (where relevant), and environment geometry. It should also include explicit task and goal labels, partner intent (such as scripts, self-reports, or latent annotations), and cover core pHHI tasks, including handover, co-manipulation, balance assistance, locomotion support, and collaborative assembly. Additionally, the dataset should record safety and comfort outcomes as well as rare events such as near-collisions and slips to facilitate risk-aware learning. Without these comprehensive datasets, methods remain tailored to specific cases and are difficult to compare.

We recommend the following criteria for dataset and data collection reporting: (1) clearly define tasks, environment geometry, and initial conditions; (2) describe participant diversity, including anthropometrics, strength, mobility, and inclusion criteria including age range, diagnosis/condition; (3) specify the sensor suite, sampling rates, and time synchronization; (4) outline the labeling protocol, including goal, phase, role, and uncertainty annotation; (5) detail train and test splits across participants, tasks, and environments, including out-of-distribution tests; (6) report safety and comfort outcomes, including rare events such as near-collisions, slips, and near-falls; and (7) provide baseline controllers and use a standardized reporting template to ensure fair comparison.

#### 6.1.6 Metric for collaboration: Stability, Efficiency, and Safety

Most pHHI papers compare systems using task-specific surrogates, such as stability via ZMP/CoP or capture-point margins, and “time in contact constraints”. Efficiency metrics include task time, path/pose error, success rate, interaction effort, assistance ratio, and sometimes human cost. Safety and comfort are measured by contact force peaks/rates, slip or over-torque counts, minimum separation, collision counts, jerk, and user studies. These metrics allow basic comparisons but are not standardized, often mask trade-offs, and rarely assess intent alignment, tail risks, co-adaptation, or user diversity. Stronger evaluation should: i) report multi-objective fronts (stability, efficiency, safety); ii) include intent-aware metrics (intent-prediction accuracy or lag, mutual information between robot and human goal); iii) use risk-sensitive statistics (CVaR of forces, minimum stability margin, barrier-function violations); iv) quantify interaction work and role-allocation smoothness; and v) standardize tasks/datasets with disturbances, and anthropometric diversity. This approach will help clarify where each method provides added value.

We recommend reporting improvements relative to a reactive baseline, such as an impedance-only follower without intent preview. In addition to task success and completion time, the following metrics should be included: $i$) peak interaction wrench and impulse ($\int \|w\| dt$), $ii$) human effort proxy, such as EMG integral when available, estimated joint effort, or validated subjective workload, $iii$) anticipation lead time, defined as the time the robot's supportive action precedes a human phase or force transition, $iv$) response latency, measured as the interval from intent change to policy update, $v$) stability and safety, including minimum stability margin and the number of safety-filter interventions, and $vi$) fluency, assessed by interaction work and role-switch smoothness. We consider a method more proactive if it reduces peak wrench or effort and demonstrates positive anticipation lead time without increasing safety interventions or near-failure events. Crucially, this suite of metrics is selected to ensure that evaluations rigorously capture system robustness against the specific failure modes and operational risks categorized in Table 5.

**Table 5** Failure modes and fallback policies in pHHI

| Failure mode (detector / symptom) | Fallback policy (recommended response) |
|---|---|
| Intent misclassification Rising belief entropy; contradictory cues | Reduce assistance gain; revert to compliant follower behavior; optionally request clarification depending on the application. |
| Sensing occlusion / dropout Missing pose frames; degraded perception | Slow down; tighten safety margins; rely more on haptics/IMU; avoid aggressive proactive actions. |
| Loss of contact Sudden wrench/tactile drop | Stop applying cooperative forces; stabilize stance; re-establish a safe grasp/contact before resuming. |
| Human–robot misalignment / slip Increasing tangential shear; slip indicators | Reduce tangential forces; re-align frames; replan footsteps and/or hand pose. |
| Model mismatch / payload change Residual spikes in tracking; unexpected CoM shifts | Adapt impedance; update estimates; if persistent, transition to safe hold. |

## 6.2 Pathways Towards Seamless pHHI: A Modular Approach

Addressing the challenges outlined earlier requires a shift in approach from developing individual components in isolation. Achieving seamless interaction depends on creating unified frameworks that integrate the pillars. We propose an interaction-centric architecture that clearly defines the roles and interfaces, or "contracts," between these pillars, enabling them to work together as a cohesive whole. This architecture is conceptualized as a multi-layered system, where each layer builds upon the guarantees and information provided by the layer beneath it.

### 6.2.1 The Foundation: Safety-Certified Controller (Pillar I)

At the lowest level, the framework should be built upon a fast controller that guarantees both the robot's physical integrity and the safety of humans at all times. This layer must be more than just a trajectory tracking module, it should serve as the ultimate authority for ensuring safe physical interactions.

**Core Functionality** The foundational controller should comprise a hierarchical WBC that respects physical constraints such as friction cones and actuation limits. This must be augmented with adaptive impedance or admittance control at the interaction contact points, allowing the robot to be both firm and compliant as needed [48, 52, 53, 69, 157, 158].

**Interface Contract** This layer is essential since it enforces a set of formally verified invariant sets that operate independently of high-level commands. This "safety shield" utilizes tools such as CBFs and stability criteria based on the ZMP or the DCM. These mechanisms ensure that the robot never enters an unrecoverable or dangerous state [79, 115, 116, 130, 162, 183].

**Uncertainty Handling and Fallback Switching** The uncertainty-to-safety interface requires the safety layer to process uncertainty outputs from the intent and human-model modules, such as belief entropy, ensemble disagreement, and predictive variance. These uncertainty metrics are mapped to conservative control actions through three primary mechanisms: (i) **constraint tightening**, which increases stability margins while reducing friction utilization and imposing conservative step limits; (ii) **assistance gain scheduling**, which attenuates feedforward assistance and increases damping or compliance as uncertainty rises; and (iii) **mode switching**, which utilizes a state machine to transition the system progressively from proactive to predictive, reactive-compliant, and finally hold/stop modes. These transitions are triggered by sustained low-confidence events (e.g., confidence remaining below a threshold for a set duration) or by critical faults such as perception dropouts and loss of contact.

### 6.2.2 The Core: Dynamic Human Feasibility and Cost Modeling (Pillar III)

At the core of the architecture should be a dynamic and continuously updated computational model of the human partner. This model should not be static but provide a real-time representation that personalizes the interaction.

**Core Functionality** This layer estimates and represents the human's physical state, including biomechanical properties, physiological status (e.g., fatigue), and ergonomic comfort.

**Interface Contract (Human Feasibility & Cost)** The system provides the higher-level planner with specific limits and preferences tailored to each individual. This includes estimated endpoint stiffness ellipsoids, joint load limits, and cost functions that assess effort or discomfort. These parameters are utilized to align the robot's behavior with the capabilities and comfort level of the human user [168, 173, 179, 184, 185, 187].

#### 6.2.3 The Brain: Predictive and Uncertainty-Aware Intent Estimation (Pillar II)

This cognitive layer interprets human actions and predicts future intentions, aiming to transition the robot from a reactive to a proactive stance.

**Core Functionality** It should integrate various sensory data, such as human pose, gaze, and interaction tools, to infer human goals. This can be accomplished using models like RNNs, Transformers, or Probabilistic Movement Primitives [167, 170, 172, 188, 189] A practical intent-belief module should: (1) align multimodal streams in time; (2) extract features such as pose derivatives, contact wrench statistics, gaze events, and EMG/IMU fatigue indices; (3) run a belief estimator, such as a Bayesian filter or ensemble Transformer, to generate goal, phase, predicted wrench envelopes, and calibrated uncertainty; (4) publish an explicit interface contract specifying belief state, horizon, confidence, and update rate; and (5) manage downstream assistance, enabling proactive actions when confidence is high and applying conservative impedance and tighter constraints when confidence is low.

**Interface Contract (Intent Belief)** It generates a streaming, calibrated belief about human intent. This is not a single point estimate, but rather a probability distribution over future trajectories, desired force envelopes, and discrete goals. This calibrated uncertainty is crucial for making informed, robust decisions.

#### 6.2.4 The Apex: Belief-Aware Shared Autonomy and Planning

At the highest level, a planner or arbiter integrates information from the lower layers to make decisions. It orchestrates collaboration, striking a balance between task efficiency and human comfort and safety.

**Core Functionality** This layer utilizes belief-aware planners, such as MPC or MPPI control, to co-design robot actions, including footsteps and grasp forces [49, 156, 159].

**Interface Contract (Action Policy)** The process begins with the Intent Belief and Human Feasibility and Cost as inputs, which are used to generate an optimal action policy. This policy is then relayed to the Foundation layer, serving as a reference for execution while adhering to safety standards at all times. Additionally, this layer is responsible for managing shared autonomy, negotiating control, and ensuring that the robot's actions are understandable to its human partner.

#### 6.2.5 End-to-End Example: Co-carrying a Table Through a Doorway (Multi-Rate and Uncertainty Handling)

In this task, a human and a humanoid robot jointly transport a bulky object, necessitating continuous force exchange, fine micro-adjustments, and occasional discrete decisions such as stopping, rotating, or stepping. The system operates at multiple signal rates: torque and impedance control are executed at kilohertz frequencies; whole-body control or quadratic programming (WBC/QP) operates at 200 to 1000 Hz; footstep and object motion planning update at 20 to 100 Hz; and intent belief updates occur at 10 to 30 Hz, based on vision and gaze, with intermittent higher-confidence cues from contact wrenches and tactile sensing. The intent module generates a calibrated belief over discrete goals (e.g., go-straight, rotate-left, rotate-right, stop) and produces a distribution over short-horizon trajectories and wrench envelopes, explicitly representing uncertainty through measures such as entropy and variance. A human model supplies person-specific feasibility and ergonomic cost metrics, including admissible wrenches and postures, comfort-weighted effort, tangential shear limits, and preferred relative pose. A belief-aware planner, such as Model Predictive Control (MPC) or Model Predictive Path Integral (MPPI), computes a receding-horizon action policy that integrates footsteps, hand and object motion, and nominal impedance targets, optimizing task progress while accounting for feasibility and risk costs under uncertainty. The safety-certified controller enforces balance and contact constraints, including center of mass and stability margins, friction cones, and joint or torque limits, by filtering unsafe commands. When uncertainty increases or state estimates are delayed, the controller tightens safety margins and reduces impedance to ensure predictable compliance. If confidence remains below a specified threshold for N cycles, the system transitions to a conservative compliant follower or hold mode until intent estimates stabilize.

#### 6.2.6 Qualitative Comparison with Alternative System Architectures

Modular interface-based designs offer distinct trade-offs compared to other approaches. Pure reactive impedance schemes are simple but have difficulty with long-term coordination. Optimization-only stacks ensure feasibility but often lack effective intent handling. End-to-end learned policies are expressive but require robust safety measures and large datasets. The proposed interface-based architecture seeks to combine interpretability through clear module responsibilities, safety via certified constraint enforcement, and adaptability by incorporating uncertainty-aware learning and prediction. It also provides explicit pathways for failure and fallback.

## 6.3 Pathways Towards Seamless pHHI: A Unified Approach

The most effective collaboration in pHHI will emerge from combining the three pillars into a unified framework. While modular approaches offer advantages such as interpretability and reduced computational cost, they also suffer from coupling issues and a lack of mutual influence between components. To achieve seamless and efficient collaboration, a unified framework is needed that integrates these elements holistically. Such unifications could be directly or indirectly inspired by the methods presented in the three pillars. We propose dividing this effort into two major areas of focus: Intent-Aware Human Models and Human-Aware Humanoid Control.

### 6.3.1 Intent-Aware Human Models

There have been primarily two approaches to integrating human intent with human models. (a) Embedding intent within the human model: In this paradigm, intent is directly incorporated into the model of human dynamics, resulting in a unified framework that predicts human states based on knowledge of the underlying goals. This leads to more coherent predictions, as motion, force, and intent are generated from a single representation. (b) Treating intent as a separate module: Here, intent is estimated independently and then used to condition the human model, or alternatively, separate human models are constructed for different intents or task categories. This modular design offers flexibility, enabling the combination of different inference and modeling methods. However, it introduces challenges of *temporal misalignment* (intent updates and biomechanical predictions may operate on different timescales) and *coupling issues* (conflicts or inconsistencies can arise when the outputs of intent inference do not align with the physical predictions of the human model). For example, a human may transiently generate forces inconsistent with their long-term intent, leading to contradictory signals across the two modules. A more unified approach, where intent and human modeling are integrated into a single predictive framework, would require richer datasets and more complex models, but would avoid these integration pitfalls. A high-fidelity model would be able to capture not only human intended goals and biomechanical aspects but also the interdependence between them. A key requirement, however, is that such models maintain both high fidelity and real-time implementability. Advancing both approaches, while maintaining real-time implementability and high-fidelity human representation, remains an important direction for future work.

### 6.3.2 Human-Aware Humanoid Control

Current literature on pHHI often treats the human either as a disturbance or as an oversimplified controller [24, 28]. Such formulations fail to capture the full state and behavior of the human, which restricts the humanoid's ability to optimize for human efficiency or adapt to individual interaction patterns. This can lead to inefficient and less natural collaboration. To address this gap, future work should focus on integrating intent-aware human models that strike a balance between fidelity and real-time applicability into humanoid control paradigms. Embedding a state–action model of the human directly within the control loop would enable human-aware control, allowing the humanoid to anticipate human responses to its actions. In turn, this would support more compliant, effective, and efficient collaborations. Moreover, optimization costs or reward functions could be shaped to explicitly account for human efficiency and comfort. These models could either embed intention directly within the model or rely on a separate intent-prediction module. Promising approaches in this direction include deep reinforcement learning (DRL) and model-based methods such as MPPI that incorporate learned models of the coupled human–humanoid system. Approaches analogous to combining terrain representations with locomotion policies in deep reinforcement learning [190] could similarly be employed to integrate human models with humanoid control.

## 6.4 Pathways Towards Seamless pHHI: Hybrid Approach

While fully modular pipelines provide simplicity and computational tractability, they suffer from coupling issues, temporal misalignment, and the inability to capture mutual influences between intent, human state, and control. Conversely, fully unified frameworks promise seamless integration by embedding intent, biomechanics, and control within a single predictive–optimization structure, but are often computationally prohibitive for real-time human–robot collaboration.

At the lower level, humanoid control is integrated with a lightweight human model that captures only fast-changing aspects of interaction. These include instantaneous interaction forces, short-horizon velocity intentions, and small postural corrections. The lightweight nature of the human model ensures that updates remain computationally feasible. This level is designed to react quickly to transient changes, prevent unsafe forces, and maintain smooth collaboration.

At the higher level, a higher-fidelity human model is employed to estimate slow-changing intents and states. Examples include overall task goals, fatigue, strength asymmetries, ergonomic preferences, and comfort envelopes. The higher layer produces long-horizon intent beliefs and ergonomic summaries, which are passed down as references, constraints, and weights to the lower level for enforcement. In long-duration co-carrying tasks, fatigue can reduce a person's ability to communicate intent through force. A low-level model might misinterpret these weaker or noisy signals as a change in intent, causing the robot to overreact. With a two-layer hybrid framework, the high-level model tracks fatigue and provides corrections, allowing the robot to recognize that reduced force is due to fatigue rather than new intent. This enables the robot to remain efficient and aligned with the task goal while reducing human strain.

More generally, the framework could be extended to as many layers as needed, depending on how fast or slow different human states evolve. Such a hybrid framework retains the benefits of unification while remaining computationally and resource efficient, making it a practical pathway for robust human–robot collaboration in resource-constrained settings.

## 6.5 Future Directions for pHHI

A significant advancement in pHHI will arise from systems that can anticipate needs, personalize experiences, and generalize knowledge, all while ensuring safety and respecting human privacy. These systems should be able to reason about the consequences of their actions, adapt in real-time to unique partners, and transfer their skills seamlessly across different tasks and environments.

### 6.5.1 Proactive and Predictive Interaction Beyond Intent Labels

Tomorrow's pHHI should not only anticipate human intentions but also predict the outcomes of a robot executing a proposed plan. This requires action-conditioned prediction models that integrate pose, gaze, and force data to forecast short-term human motion and forces under hypothetical robot actions. These models should provide both trajectories and uncertainty estimates to support belief-aware planners. Two key capabilities are essential for this approach. First, consequence-aware forecasting helps distinguish whether a planned acceleration, footstep, or grip change will lead to assistance or resistance. Second, intentional versus accidental motion discrimination is necessary. This involves utilizing multimodal signals, such as gaze fixation, pre-activation patterns, and sudden transitions in force to differentiate purposeful movements from slips, startle reflexes, and to modulate assistance based on these cues appropriately.

### 6.5.2 Personalization and Co-Adaptation: Physical and Psychological

A robot must be able to customize its assistance based on the individual's morphology, strength, stiffness, fatigue, and preferences. High-fidelity musculoskeletal and neuromechanical models provide essential variables, such as joint loads, endpoint stiffness ellipsoids, and momentum heuristics. However, these models need to be simplified into fast, person-specific surrogates that can be created through brief, non-intrusive calibration methods. Online adaptation is necessary to update ergonomic parameters, such as reachable and comfortable postures and wrench-rate limits, as fatigue changes over time. Additionally, personalization has a psychological component; trust and comfort depend on a robot's legibility and predictability. Future systems should aim to encode user-defined preferences (including pace, assertiveness, and haptic firmness) along with implicit indicators (such as micro-withdrawals and grip variations) as constraints and costs. Furthermore, they should be capable of managing role exchange (i.e., deciding who leads or follows) when confidence in their intent or comfort diminishes.

### 6.5.3 Privacy by Design in Contact-Rich Interaction

The rich multimodal sensing required for advanced pHHI, including video, gaze tracking, and inertial wearables, raises legitimate concerns about privacy. A forward-looking approach must prioritize privacy by design, emphasizing data minimization and on-device inference. This includes favoring robot-side force and tactile sensing over high-fidelity imagery when in close contact, performing intent estimation at the edge, and ensuring that protocols are transparent, giving users explicit control over what data is stored and shared. Recommended mitigations include data minimization, specifically, storing abstract features rather than raw video or audio, and the enforcement of strict retention periods. Systems should default to on-device inference and offer configurable privacy modes, such as disabling RGB streams in favor of depth silhouettes or contact-only sensing. To enable personalization without centralizing raw data, we recommend federated learning and privacy-preserving aggregation methods like differential privacy. Finally, to validate these improvements, researchers should report a **privacy-utility curve** that quantifies the trade-off between performance (e.g., intent accuracy, collaboration fluency) and privacy settings, while explicitly documenting all data storage protocols and retention limits to ensure transparency.

## 7 Conclusion

This review has explored the current state of pHHI through the key pillars of modeling and control of humanoids, human intent estimation and prediction, and human models. While significant advancements have been made in each of these areas, the main obstacle to achieving seamless and robust interaction is the ongoing fragmentation among them. The challenges related to safety, predictability, and adaptation are not confined to a single pillar; rather, they are systemic issues that require an integrated solution.

We propose that the way forward involves creating unified architectural frameworks that closely connect predictive intent estimation, personalized human models, and certifiably safe whole-body control. By designing systems where these components inform and constrain one another in real-time, we can transition from reactive robots to proactive and efficient partners. Future research focused on this integration will be essential for developing humanoids that can anticipate human needs, provide personalized assistance, and generalize their skills, ultimately enabling them to become true collaborators in our daily lives.

**Author Contributions** G.A.C. and S.S.K. performed the literature search and data analysis and drafted the paper. P.A. had the idea for the article and critically revised the work. P.A. oversaw the project and provided funding and supervision.

**Funding** This material is based upon work supported by the National Science Foundation under Grants No. 2018905, 250047, 2415093.

**Data Availability** No datasets were generated or analysed during the current study.

## Declarations

**Ethical Approval** Not applicable.

**Conicts of Interest** The authors have no conicts of interest to declare that are relevant to the content of this article.

**Competing Interests** The authors declare no competing interests.



## References

1. Neumann, C.: Humanoid robots offer both disruption and promise. Here's why. (2025). https://www.weforum.org/stories/2025/06/humanoid-robots-offer-disruption-and-promise/. Accessed 09 Aug 2025
2. Sakagami, Y., Watanabe, R., Aoyama, C., Matsunaga, S., Higaki, N., Fujimura, K.: The intelligent ASIMO: System overview and integration. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, vol. 3, pp. 2478–2483 (2002). IEEE
3. Guizzo, E.: By leaps and bounds: An exclusive look at how boston dynamics is redefining robot agility. IEEE Spectr. **56**(12), 34–39 (2019)
4. Unitree Robotics: G1 Humanoid Robot. (n.d.). https://www.unitree.com/g1. Accessed 09 Aug 2025
5. Agility Robotics: Agility Robotics. (n.d.). https://www.agilityrobotics.com/. Accessed 09 Aug 2025
6. Mukherjee, S., Baral, M.M., Pal, S.K., Chittipaka, V., Roy, R., Alam, K.: Humanoid robot in healthcare: a systematic review and future research directions. In: 2022 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COM-IT-CON), vol. 1, pp. 822–826 (2022). IEEE
7. McGinn, C., Cullinan, M., Holland, D., Kelly, K.: Towards the design of a new humanoid robot for domestic applications. In: 2014 IEEE International Conference on Technologies for Practical Robot Applications (TePRA), pp. 1–6 (2014). IEEE
8. Settimi, A., Pavan, C., Varricchio, V., Ferrati, M., Mingo Hoffman, E., Rocchi, A., Melo, K., Tsagarakis, N.G., Bicchi, A.: A modular approach for remote operation of humanoid robots in search and rescue scenarios. In: International Workshop on Modelling and Simulation for Autonomous Systems, pp. 192–205 (2014). Springer
9. Diftler, M.A., Mehling, J.S., Abdallah, M.E., Radford, N.A., Bridgwater, L.B., Sanders, A.M., Askew, R.S., Linn, D.M., Yamokoski, J.D., Permenter, F., et al.: Robonaut 2-the first humanoid robot in space. In: 2011 IEEE International Conference on Robotics and Automation, pp. 2178–2183 (2011). IEEE
10. Tong, Y., Liu, H., Zhang, Z.: Advancements in humanoid robots: A comprehensive review and future prospects. IEEE/CAA Journal of Automatica Sinica. **11**(2), 301–328 (2024)
11. Mikolajczyk, T., Mikołajewska, E., Al-Shuka, H.F., Malinowski, T., Kłodowski, A., Pimenov, D.Y., Paczkowski, T., Hu, F., Giasin, K., Mikołajewski, D., et al.: Recent advances in bipedal walking robots: Review of gait, drive, sensors and control systems. Sensors. **22**(12), 4440 (2022)
12. Zhang, H., Wu, J., Fan, J., An, Y., Jin, X., Cui, D., Yang, Y.: A review of fall coping strategies for humanoid robots. J. Bionic Eng. **22**(2), 480–512 (2025)
13. Laschi, C., Dario, P., Carrozza, M.C., Guglielmelli, E., Teti, G., Taddeucci, D., Leoni, F., Massa, B., Zecca, M., Lazzarini, R., et al.: Grasping and manipulation in humanoid robotics. Scuola Superiore Sant Anna, Italia (2000)
14. Farajtabar, M., Charbonneau, M.: The path towards contact-based physical human-robot interaction. Robot. Auton. Syst. **182**, 104829 (2024)
15. Lefèvre, H., Chaki, T., Kawakami, T., Tanguy, A., Yoshiike, T., Kheddar, A.: Humanoid-human sit-to-stand-to-sit assistance. IEEE Robotics and Automation Letters. (2024)
16. Agravante, D.J., Cherubini, A., Sherikov, A., Wieber, P.-B., Kheddar, A.: Human-humanoid collaborative carrying. IEEE Trans. Rob. **35**(4), 833–846 (2019)
17. Darvish, K., Penco, L., Ramos, J., Cisneros, R., Pratt, J., Yoshida, E., Ivaldi, S., Pucci, D.: Teleoperation of humanoid robots: A survey. IEEE Trans. Rob. **39**(3), 1706–1727 (2023)

18. Losey, D.P., McDonald, C.G., Battaglia, E., O'Malley, M.K.: A review of intent detection, arbitration, and communication aspects of shared control for physical human-robot interaction. Appl. Mech. Rev. **70**(1), 010804 (2018)
19. Ficht, G., Behnke, S.: Bipedal humanoid hardware design: A technology review. Current Robotics Reports. **2**(2), 201–210 (2021)
20. Vianello, L., Penco, L., Gomes, W., You, Y., Anzalone, S.M., Maurice, P., Thomas, V., Ivaldi, S.: Human-humanoid interaction and cooperation: a review. Current Robotics Reports. **2**(4), 441–454 (2021)
21. Monje, C., Pierro, P., Balaguer, C.: A new approach on human-robot collaboration with humanoid robot rh-2. Robotica **29**(6), 949–957 (2011)
22. Evrard, P., Kheddar, A.: Homotopy-based controller for physical human-robot interaction. In: RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication, pp. 1–6 (2009). IEEE
23. Agravante, D.J., Sherikov, A., Wieber, P.-B., Cherubini, A., Kheddar, A.: Walking pattern generators designed for physical collaboration. In: 2016 IEEE International Conference on Robotics and Automation (ICRA), pp. 1573–1578 (2016). IEEE
24. Stasse, O., Evrard, P., Perrin, N., Mansard, N., Kheddar, A.: Fast foot prints re-planning and motion generation during walking in physical human-humanoid interaction. In: 2009 9th IEEE-RAS International Conference on Humanoid Robots, pp. 284–289 (2009). IEEE
25. Agravante, D.J., Cherubini, A., Bussy, A., Gergondet, P., Kheddar, A.: Collaborative human-humanoid carrying using vision and haptic sensing. In: 2014 IEEE International Conference on Robotics and Automation (ICRA), pp. 607–612 (2014). IEEE
26. Rahem, R., Wong, C.Y., Suleiman, W.: Human-humanoid robot cooperative load transportation: model-based control approach. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 8306–8312 (2022). IEEE
27. Stückler, J., Behnke, S.: Following human guidance to cooperatively carry a large object. In: 2011 11th IEEE-RAS International Conference on Humanoid Robots, pp. 218–223 (2011). IEEE
28. Kumbhar, S.S., Artemiadis, P.: MPC-QP-based control framework for compliant behavior of humanoid robots in physical collaboration with humans. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), pp. 14714–14720 (2025). https://doi.org/10.1109/ICRA55743.2025.11128037
29. Bussy, A., Gergondet, P., Kheddar, A., Keith, F., Crosnier, A.: Proactive behavior of a humanoid robot in a haptic transportation task with a human partner. In: 2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication, pp. 962–967 (2012). IEEE
30. Maroger, I., Stasse, O., Watier, B.: From the study of table trajectories during collaborative carriages toward pro-active human-robot table handling tasks. In: 2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids), pp. 911–918 (2022). IEEE
31. López, A.M., Vaillant, J., Keith, F., Fraisse, P., Kheddar, A.: Compliant control of a humanoid robot helping a person stand up from a seated position. In: 2014 IEEE-RAS International Conference on Humanoid Robots, pp. 817–822 (2014). IEEE
32. Rapetti, L., Sartore, C., Elobaid, M., Tirupachuri, Y., Draicchio, F., Kawakami, T., Yoshiike, T., Pucci, D.: A control approach for human-robot ergonomic payload lifting. In: 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 7504–7510 (2023). IEEE
33. Yang, P., Yan, H., Yang, B., Li, J., Li, K., Leng, Y., Fu, C.: A centaur system for assisting human walking with load carriage. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 5242–5248 (2022). IEEE
34. Song, P., Yu, Y., Zhang, X.: A tutorial survey and comparison of impedance control on robotic manipulation. Robotica **37**(5), 801–836 (2019)
35. Li, Z., Chu, X., Hu, X., Zhang, Z., Li, N., Li, J.: Variable stiffness methods for robots: A review. Smart Mater. Struct. **33**(6), 063002 (2024)
36. Hyon, S.-H., Hale, J.G., Cheng, G.: Full-body compliant human-humanoid interaction: balancing in the presence of unknown external forces. IEEE Trans. Rob. **23**(5), 884–898 (2007)
37. Bussy, A., Kheddar, A., Crosnier, A., Keith, F.: Human-humanoid haptic joint object transportation case study. In: 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3633–3638 (2012). IEEE
38. Otani, K., Bouyarmane, K., Ivaldi, S.: Generating assistive humanoid motions for co-manipulation tasks with a multi-robot quadratic program controller. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 3107–3113 (2018). IEEE
39. Romano, F., Nava, G., Azad, M., Čamernik, J., Dafarra, S., Dermy, O., Latella, C., Lazzaroni, M., Lober, R., Lorenzini, M., et al.: The codyco project achievements and beyond: Toward human aware whole-body controllers for physical human robot interaction. IEEE Robotics and Automation Letters. **3**(1), 516–523 (2017)
40. Bethala, G.C.R., Huang, H., Pudasaini, N., Ali, A.M., Yuan, S., Wen, C., Tzes, A., Fang, Y.: H2-COMPACT: Human-humanoid co-manipulation via adaptive contact trajectory policies. (2025). arXiv preprint arXiv:2505.17627
41. Evrard, P., Gribovskaya, E., Calinon, S., Billard, A., Kheddar, A.: Teaching physical collaborative tasks: object-lifting case study with a humanoid. In: 2009 9th IEEE-RAS International Conference on Humanoid Robots, pp. 399–404 (2009). IEEE
42. Lee, D., Ott, C., Nakamura, Y., Hirzinger, G.: Physical human robot interaction in imitation learning. In: 2011 IEEE International Conference on Robotics and Automation, pp. 3439–3440 (2011). IEEE
43. Ikemoto, S., Amor, H.B., Minato, T., Ishiguro, H., Jung, B.: Physical interaction learning: Behavior adaptation in cooperative human-robot tasks involving physical contact. In: RO-MAN 2009-The 18th IEEE International Symposium on Robot and Human Interactive Communication, pp. 504–509 (2009). IEEE
44. Brecelj, T., Petrič, T.: Utilizing a phase state system for reliable physical assistance in human-humanoid robot collaboration. In: 2023 21st International Conference on Advanced Robotics (ICAR), pp. 258–263 (2023). IEEE
45. Kajita, S., Kanehiro, F., Kaneko, K., Fujiwara, K., Harada, K., Yokoi, K., Hirukawa, H.: Resolved momentum control: Humanoid motion planning based on the linear and angular momentum. In: Proceedings 2003 Ieee/rsj International Conference on Intelligent Robots and Systems (iros 2003)(cat. No. 03ch37453), vol. 2, pp. 1644–1650 (2003). IEEE
46. Orin, D.E., Goswami, A., Lee, S.-H.: Centroidal dynamics of a humanoid robot. Auton. Robot. **35**, 161–176 (2013)
47. Kuindersma, S., Deits, R., Fallon, M., Valenzuela, A., Dai, H., Permenter, F., Koolen, T., Marion, P., Tedrake, R.: Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. Auton. Robot. **40**, 429–455 (2016)
48. Dai, H., Valenzuela, A., Tedrake, R.: Whole-body motion planning with centroidal dynamics and full kinematics. In: 2014 IEEE-RAS International Conference on Humanoid Robots, pp. 295–302 (2014). IEEE
49. Ponton, B., Herzog, A., Schaal, S., Righetti, L.: A convex model of humanoid momentum dynamics for multi-contact motion generation. In: 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids), pp. 842–849 (2016). IEEE

50. Wieber, P.-B.: Trajectory free linear model predictive control for stable walking in the presence of strong perturbations. In: 2006 6th IEEE-RAS International Conference on Humanoid Robots, pp. 137–142 (2006). IEEE
51. Kajita, S., Tani, K.: Study of dynamic biped locomotion on rugged terrain-derivation and application of the linear inverted pendulum mode. In: Proceedings. 1991 IEEE International Conference on Robotics and Automation, pp. 1405–1406 (1991). IEEE Computer Society
52. Caron, S., Escande, A., Lanari, L., Mallein, B.: Capturability-based analysis, optimization and control of 3d bipedal walking. (2018). arXiv preprint arXiv:1801.07022
53. Caron, S., Mallein, B.: Balance control using both ZMP and COM height variations: A convex boundedness approach. In: 2018 IEEE International Conference on Robotics and Automation (ICRA), pp. 1779–1784 (2018). IEEE
54. Hirukawa, H., Hattori, S., Harada, K., Kajita, S., Kaneko, K., Kanehiro, F., Fujiwara, K., Morisawa, M.: A universal stability criterion of the foot contact of legged robots-adios ZMP. In: Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006., pp. 1976–1983 (2006). IEEE
55. Shahbazi, M., Babuška, R., Lopes, G.A.: Unified modeling and control of walking and running on the spring-loaded inverted pendulum. IEEE Trans. Rob. **32**(5), 1178–1195 (2016)
56. Geyer, H., Saranli, U.: Gait based on the spring-loaded inverted pendulum. In: Humanoid Robotics: A Reference, pp. 923–947. Springer, Dordrecht (2019)
57. Koolen, T., De Boer, T., Rebula, J., Goswami, A., Pratt, J.: Capturability-based analysis and control of legged locomotion, part 1: Theory and application to three simple gait models. The Int J Robot Res. **31**(9), 1094–1113 (2012)
58. Englsberger, J., Ott, C., Roa, M.A., Albu-Schäffer, A., Hirzinger, G.: Bipedal walking control based on capture point dynamics. In: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 4420–4427 (2011). IEEE
59. Englsberger, J., Mesesan, G., Ott, C.: Smooth trajectory generation and push-recovery based on divergent component of motion. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4560–4567 (2017). IEEE
60. Radosavovic, I., Xiao, T., Zhang, B., Darrell, T., Malik, J., Sreenath, K.: Learning humanoid locomotion with transformers. CoRR. (2023)
61. Lai, H., Zhang, W., He, X., Yu, C., Tian, Z., Yu, Y., Wang, J.: Sim-to-real transfer for quadrupedal locomotion via terrain transformer. In: 2023 IEEE International Conference on Robotics and Automation (ICRA), pp. 5141–5147 (2023). IEEE
62. Dugar, P., Shrestha, A., Yu, F., Marum, B., Fern, A.: Learning multi-modal whole-body control for real-world humanoid robots. arXiv preprint arXiv:2408.07295. (2024)
63. Paredes, V.C., Hereid, A.: Resolved motion control for 3d underactuated bipedal walking using linear inverted pendulum dynamics and neural adaptation. In: 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 6761–6767 (2022). IEEE
64. Dao, J., Duan, H., Fern, A.: Sim-to-real learning for humanoid box loco-manipulation. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 16930–16936 (2024). IEEE
65. Xie, Z., Tseng, J., Starke, S., Panne, M., Liu, C.K.: Hierarchical planning and control for box loco-manipulation. Proceedings of the ACM on Computer Graphics and Interactive Techniques. **6**(3), 1–18 (2023)
66. Cheng, X., Ji, Y., Chen, J., Yang, R., Yang, G., Wang, X.: Expressive whole-body control for humanoid robots (2024). arXiv preprint arXiv:2402.16796
67. Ji, M., Peng, X., Liu, F., Li, J., Yang, G., Cheng, X., Wang, X.: Exbody2: Advanced expressive humanoid whole-body control (2024). arXiv preprint arXiv:2412.13196
68. Liu, F., Gu, Z., Cai, Y., Zhou, Z., Jung, H., Jang, J., Zhao, S., Ha, S., Chen, Y., Xu, D., et al.: Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation (2024). arXiv preprint arXiv:2409.20514
69. Fu, Z., Zhao, Q., Wu, Q., Wetzstein, G., Finn, C.: Humanplus: Humanoid shadowing and imitation from humans (2024). arXiv preprint arXiv:2406.10454
70. Khatib, O.: A unified approach for motion and force control of robot manipulators: The operational space formulation. IEEE J. Robot. Autom. **3**(1), 43–53 (1987)
71. Sentis, L., Khatib, O.: Control of free-floating humanoid robots through task prioritization. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2005)
72. Sentis, L., Khatib, O.: A whole-body control framework for humanoids operating in human environments. In: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA) (2006)
73. Park, J., Khatib, O.: Robot multiple contact control. Robotica **26**(5), 661–673 (2008)
74. Sentis, L.: Compliant control of multicontact and center-of-mass behaviors in humanoid robots. IEEE Transactions on Robotics (2010)
75. Nozawa, S., Ueda, R., Kakiuchi, Y., Okada, K., Inaba, M.: A full-body motion control method for a humanoid robot based on on-line estimation of the operational force of an object with an unknown weight. In: 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 2684–2691 (2010). IEEE
76. Dietrich, A., Wimböck, T., Albu-Schäffer, A.: Dynamic whole-body mobile manipulation with a torque-controlled humanoid robot via impedance control laws. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (2011)
77. Dietrich, A., Wimböck, T., Albu-Schäffer, A., Hirzinger, G.: Reactive whole-body control: Dynamic mobile manipulation using a large number of degrees of freedom. IEEE Transactions on Robotics. (2012)
78. Hsu, S.-C., Xu, X., Ames, A.D.: Control barrier function based quadratic programs with application to bipedal robotic walking. In: 2015 American Control Conference (ACC), pp. 4542–4548 (2015). IEEE
79. Sartore, C., Rando, M., Romualdi, G., Molinari, C., Rosasco, L., Pucci, D.: Automatic gain tuning for humanoid robots walking architectures using gradient-free optimization techniques. In: 2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids), pp. 996–1003 (2024). IEEE
80. Murooka, M., Chappellet, K., Tanguy, A., Benallegue, M., Kumagai, I., Morisawa, M., Kanehiro, F., Kheddar, A.: Humanoid loco-manipulations pattern generation and stabilization control. IEEE Robotics and Automation Letters. **6**(3), 5597–5604 (2021)
81. Audren, H., Vaillant, J., Kheddar, A., Escande, A., Kaneko, K., Yoshida, E.: Model preview control in multi-contact motion-application to a humanoid robot. In: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 4030–4035 (2014). IEEE
82. Harada, K., Kajita, S., Kanehiro, F., Fujiwara, K., Kaneko, K., Yokoi, K., Hirukawa, H.: Real-time planning of humanoid robot's gait for force-controlled manipulation. IEEE/ASME Trans. Mechatron. **12**(1), 53–62 (2007)
83. Fahmi, S., Mastalli, C., Focchi, M., Semini, C.: Passive whole-body control for quadruped robots: Experimental validation over challenging terrain. IEEE Robotics and Automation Letters. **4**(3), 2553–2560 (2019)

84. Diedam, H., Dimitrov, D., Wieber, P.-B., Mombaur, K., Diehl, M.: Online walking gait generation with adaptive foot positioning through linear model predictive control. In: 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 1121–1126 (2008). IEEE
85. Herdt, A., Diedam, H., Wieber, P.-B., Dimitrov, D., Mombaur, K., Diehl, M.: Online walking motion generation with automatic footstep placement. Adv. Robot. **24**(5–6), 719–737 (2010)
86. Kumbhar, S.S., Kulkarni, A.M., Poulakakis, I.: Finite-step capturability and recursive feasibility for bipedal walking in constrained regions. In: 2025 IEEE International Conference on Robotics and Automation (ICRA), pp. 397–403 (2025). https://doi.org/10.1109/ICRA55743.2025.11128831
87. Geyer, H., Seyfarth, A., Blickhan, R.: Compliant leg behaviour explains basic dynamics of walking and running. Proceedings of the Royal Society B: Biological Sciences. **273**(1603), 2861–2867 (2006)
88. Xiong, X., Ames, A.D.: Dynamic and versatile humanoid walking via embedding 3d actuated slip model with hybrid lip based stepping. IEEE Robot. Autom. Lett. **5**(4), 6286–6293 (2020)
89. Khatib, O.: A unified approach for motion and force control of robot manipulators: The operational space formulation. IEEE J. Robot. Autom. **3**(1), 43–53 (1987). https://doi.org/10.1109/JRA.1987.1087068
90. Baerlocher, P., Boulic, R.: Task-priority formulations for the kinematic control of highly redundant articulated structures. In: Proceedings. 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems. Innovations in Theory, Practice and Applications (Cat. No. 98CH36190), vol. 1, pp. 323–329 (1998). IEEE
91. Mansard, N., Chaumette, F.: Task sequencing for high-level sensor-based control. IEEE Trans. Rob. **23**(1), 60–72 (2007). https://doi.org/10.1109/TRO.2006.889487
92. Escande, A., Mansard, N., Wieber, P.-B.: Hierarchical quadratic programming: Fast online humanoid-robot motion generation. Int. J. Robot. Res. **33**(7), 1006–1028 (2014). https://doi.org/10.1177/0278364914521306
93. Sentis, L., Khatib, O.: A whole-body control framework for humanoids operating in human environments. In: Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006., Orlando, FL, USA (2006)
94. Saab, L., Ramos, O.E., Keith, F., Mansard, N., Souères, P., Fourquet, J.-Y.: Dynamic whole-body motion generation under rigid contacts and other unilateral constraints. IEEE Trans. Rob. **29**(2), 346–362 (2013). https://doi.org/10.1109/TRO.2012.2234351
95. Herzog, A., Rotella, N., Mason, S., Grimminger, F., Schaal, S., Righetti, L.: Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid. Auton. Robot. **40**(3), 473–491 (2016). https://doi.org/10.1007/s10514-015-9476-6
96. Alvarez-Padilla, J., Zhang, J.Z., Kwok, S., Dolan, J.M., Manchester, Z.: Real-time whole-body control of legged robots with model-predictive path integral control (2024). arXiv preprint arXiv:2409.10469
97. Xue, H., Pan, C., Yi, Z., Qu, G., Shi, G.: Full-order sampling-based MPC for torque-level locomotion control via diffusion-style annealing (2024). arXiv preprint arXiv:2409.15610
98. Tassa, Y., Erez, T., Todorov, E.: Synthesis and stabilization of complex behaviors through online trajectory optimization. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 4906–4913 (2012). https://doi.org/10.1109/IROS.2012.6386025
99. Tassa, Y., Mansard, N., Todorov, E.: Control-limited differential dynamic programming. In: IEEE International Conference on Robotics and Automation (ICRA), pp. 1168–1175 (2014). https://doi.org/10.1109/ICRA.2014.6907001
100. Budhiraja, R., Carpentier, J., Mastalli, C., Mansard, N.: Differential dynamic programming for multi-phase rigid contact dynamics. In: IEEE-RAS International Conference on Humanoid Robots (Humanoids), pp. 1–8 (2018). https://doi.org/10.1109/HUMANOIDS.2018.8624925
101. Li, H., Wensing, P.M.: Hybrid systems differential dynamic programming for whole-body motion planning of legged robots. IEEE Robot. Autom. Lett. **5**(3), 4556–4563 (2020). https://doi.org/10.1109/LRA.2020.3007475
102. Schwarke, C., Klemm, V., Boon, M., Bjelonic, M., Hutter, M.: Curiosity-driven learning of joint locomotion and manipulation tasks. In: Proceedings of the 7th Conference on Robot Learning, vol. 229, pp. 2594–2610 (2023). PMLR
103. Bao, L., Humphreys, J., Peng, T., Zhou, C.: Deep Reinforcement Learning for Bipedal Locomotion: A Brief Survey (2025). arXiv:2404.17070
104. Siekmann, J., Valluri, S., Dao, J., Bermillo, L., Duan, H., Fern, A., Hurst, J.: Learning memory-based control for human-scale bipedal locomotion. In: Robotics Science and Systems (2020)
105. Siekmann, J., Godse, Y., Fern, A., Hurst, J.: Sim-to-real learning of all common bipedal gaits via periodic reward composition. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 7309–7315 (2021). https://doi.org/10.1109/ICRA48506.2021.9561814
106. Xie, Z., Berseth, G., Clary, P., Hurst, J., Panne, M.: Feedback control for cassie with deep reinforcement learning. In: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 1241–1246 (2018). https://doi.org/10.1109/IROS.2018.8593722
107. Xie, Z., Clary, P., Dao, J., Morais, P., Hurst, J., Panne, M.: Learning locomotion skills for cassie: Iterative design and sim-to-real. In: Conference on Robot Learning, pp. 317–329 (2020). PMLR
108. Li, Z., Cheng, X., Peng, X.B., Abbeel, P., Levine, S., Berseth, G., Sreenath, K.: Reinforcement learning for robust parameterized locomotion control of bipedal robots. In: 2021 IEEE International Conference on Robotics and Automation (ICRA), pp. 2811–2817 (2021). https://doi.org/10.1109/ICRA48506.2021.9560769
109. Li, J., Ye, L., Cheng, Y., Liu, H., Liang, B.: Agile and versatile bipedal robot tracking control through reinforcement learning. arXiv preprint arXiv:2404.08246. (2024)
110. Singh, R.P., Benallegue, M., Morisawa, M., Cisneros, R., Kanehiro, F.: Learning bipedal walking on planned footsteps for humanoid robots. In: 2022 IEEE-RAS 21st International Conference on Humanoid Robots (Humanoids), pp. 686–693 (2022). https://doi.org/10.1109/Humanoids53995.2022.10000067
111. Levine, S., Koltun, V.: Guided policy search. In: International Conference on Machine Learning, pp. 1–9 (2013). PMLR
112. Gangapurwala, S., Mitchell, A., Havoutis, I.: Guided constrained policy optimization for dynamic quadrupedal robot locomotion. IEEE Robot. Autom. Lett. **5**(2), 3642–3649 (2020). https://doi.org/10.1109/LRA.2020.2979656
113. Liu, Z., Lu, X., Liu, W., Qi, W., Su, H.: Human-robot collaboration through a multi-scale graph convolution neural network with temporal attention. IEEE Robot. Autom. Lett. **9**(3), 2248–2255 (2024)
114. Wu, F., Gu, Z., Wu, H., Wu, A., Zhao, Y.: Infer and adapt: Bipedal locomotion reward learning from demonstrations via inverse reinforcement learning. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 16243–16250 (2024). IEEE
115. Kumar, K.N., Essa, I., Ha, S.: Words into action: Learning diverse humanoid robot behaviors using language guided iterative motion refinement (2023). arXiv preprint arXiv:2310.06226
116. Jiang, Z., Xie, Y., Li, J., Yuan, Y., Zhu, Y., Zhu, Y.: Harmon: Whole-body motion generation of humanoid robots from language descriptions (2024). arXiv preprint arXiv:2410.12773

117. Yao, Y., He, W., Gu, C., Du, J., Tan, F., Zhu, Z., Lu, J.: Anybipe: An end-to-end framework for training and deploying bipedal robots guided by large language models (2024). arXiv preprint arXiv:2409.08904
118. Sun, J., Zhang, Q., Duan, Y., Jiang, X., Cheng, C., Xu, R.: Prompt, plan, perform: LLM-based humanoid control via quantized imitation learning. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 16236–16242 (2024). IEEE
119. He, T., Luo, Z., He, X., Xiao, W., Zhang, C., Zhang, W., Kitani, K., Liu, C., Shi, G.: Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning (2024). arXiv preprint arXiv:2406.08858
120. Robey, A., Ravichandran, Z., Kumar, V., Hassani, H., Pappas, G.J.: Jailbreaking LLM-controlled robots (2024). arXiv preprint arXiv:2410.13691
121. Zhang, H., Zhu, C., Wang, X., Zhou, Z., Hu, S., Zhang, L.Y.: Badrobot: Jailbreaking LLM-based embodied AI in the physical world **3** (2024). arXiv preprint arXiv:2407.20242
122. Westervelt, E.R., Grizzle, J.W., Koditschek, D.E.: Hybrid zero dynamics of planar biped walkers. IEEE Trans. Autom. Control **48**(1), 42–56 (2003)
123. Westervelt, E.R., Grizzle, J.W., Chevallereau, C., Choi, J.H., Morris, B.: Feedback Control of Dynamic Bipedal Robot Locomotion. CRC Press, Boca Raton (2018)
124. Grizzle, J.W., Chevallereau, C., Sinnet, R.W., Ames, A.D.: Models, feedback control, and open problems of 3D bipedal robotic walking. Automatica **50**(8), 1955–1988 (2014)
125. Zhang, Y., Doyle, T.: Integrating intention-based systems in human-robot interaction: A scoping review of sensors, algorithms, and trust. Frontiers in Robotics and AI. **10**, 1233328 (2023)
126. Hoffman, G., Bhattacharjee, T., Nikolaidis, S.: Inferring human intent and predicting human action in human–robot collaboration. Ann. Rev. Control Robot. Autonomous Syst. **7** (2023)
127. Martini, E., Parekh, H., Peng, S., Bombieri, N., Figueroa, N.: A robust filter for marker-less multi-person tracking in human-robot interaction scenarios. In: 2024 33rd IEEE International Conference on Robot and Human Interactive Communication (ROMAN), pp. 424–429 (2024). IEEE
128. Martinez, J., Hossain, R., Romero, J., Little, J.J.: A simple yet effective baseline for 3D human pose estimation. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 2640–2649 (2017)
129. Yu, R., Park, H., Lee, J.: Human dynamics from monocular video with dynamic camera movements. ACM Transactions on Graphics (TOG). **40**(6), 1–14 (2021)
130. Lanini, J., Razavi, H., Urain, J., Ijspeert, A.: Human intention detection as a multiclass classification problem: Application in physical human-robot interaction while walking. IEEE Robot. Autom. Lett. **3**(4), 4171–4178 (2018)
131. Lyu, K., Chen, H., Liu, Z., Zhang, B., Wang, R.: 3D human motion prediction: A survey. Neurocomputing **489**, 345–365 (2022)
132. Usman, M., Zhong, J.: Skeleton-based motion prediction: A survey. Front. Comput. Neurosci. **16**, 1051222 (2022)
133. Strabala, K., Lee, M.K., Dragan, A., Forlizzi, J., Srinivasa, S.S., Cakmak, M., Micelli, V.: Toward seamless human-robot handovers. Journal of Human-Robot Interaction. **2**(1), 112–132 (2013)
134. Admoni, H., Dragan, A., Srinivasa, S.S., Scassellati, B.: Deliberate delays during robot-to-human handovers improve compliance with gaze communication. In: Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction, pp. 49–56 (2014)
135. Mainprice, J., Berenson, D.: Human-robot collaborative manipulation planning using early prediction of human motion. In: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 299–306 (2013). IEEE
136. Townsend, E.C., Mielke, E.A., Wingate, D., Killpack, M.D.: Estimating human intent for physical human-robot co-manipulation (2017). arXiv preprint arXiv:1705.10851
137. Al-Yacoub, A., Zhao, Y., Eaton, W., Goh, Y.M., Lohse, N.: Improving human robot collaboration through force/torque based learning for object manipulation. Robotics and Computer-Integrated Manufacturing. **69**, 102111 (2021)
138. Shao, Y.S., Li, T., Keyvanian, S., Chaudhari, P., Kumar, V., Figueroa, N.: Constraint-aware intent estimation for dynamic human-robot object co-manipulation (2024). arXiv preprint arXiv:2409.00215
139. Erden, M.S., Tomiyama, T.: Human-intent detection and physically interactive control of a robot without force sensors. IEEE Trans. Rob. **26**(2), 370–382 (2010)
140. Li, Y., Ge, S.S.: Human-robot collaboration based on motion intention estimation. IEEE/ASME Trans. Mechatron. **19**(3), 1007–1014 (2013)
141. Belcamino, V., Takase, M., Kilina, M., Carfì, A., Mastrogiovanni, F., Shimada, A., Shimizu, S.: Gaze-based intention recognition for human-robot collaboration. In: Proceedings of the 2024 International Conference on Advanced Visual Interfaces, pp. 1–5 (2024)
142. Taborri, J., Palermo, E., Rossi, S., Cappa, P.: Gait partitioning methods: A systematic review. Sensors. **16**(1), 66 (2016)
143. Xue, H., Herzog, R., Berger, T.M., Bäumer, T., Weissbach, A., Rueckert, E.: Using probabilistic movement primitives in analyzing human motion differences under transcranial current stimulation. Frontiers in Robotics and AI. **8**, 721890 (2021)
144. Su, H., Qi, W., Chen, J., Yang, C., Sandoval, J., Laribi, M.A.: Recent advancements in multimodal human-robot interaction. Front. Neurorobot. **17**, 1084000 (2023)
145. Qiu, S., Zhao, H., Jiang, N., Wang, Z., Liu, L., An, Y., Zhao, H., Miao, X., Liu, R., Fortino, G.: Multi-sensor information fusion based on machine learning for real applications in human activity recognition: State-of-the-art and research challenges. Information Fusion. **80**, 241–265 (2022)
146. Wang, T., Zheng, P., Li, S., Wang, L.: Multimodal human-robot interaction for human-centric smart manufacturing: a survey. Adv. Intell. Syst. **6**(3), 2300359 (2024)
147. Wong, C.Y., Vergez, L., Suleiman, W.: Vision-and tactile-based continuous multimodal intention and attention recognition for safer physical human–robot interaction. IEEE Trans. Autom. Sci. Eng. (2023). https://doi.org/10.1109/TASE.2023.3248496
148. Callens, T., Have, T., Van Rossom, S., De Schutter, J., Aertbeliën, E.: A framework for recognition and prediction of human motions in human-robot collaboration using probabilistic motion models. IEEE Robot. Autom. Lett. **5**(4), 5151–5158 (2020)
149. Wang, Z., Mülling, K., Deisenroth, M.P., Ben Amor, H., Vogt, D., Schölkopf, B., Peters, J.: Probabilistic movement modeling for intention inference in human-robot interaction. Int. J. Robot. Res. **32**(7), 841–858 (2013)
150. Maeda, G.J., Neumann, G., Ewerton, M., Lioutikov, R., Kroemer, O., Peters, J.: Probabilistic movement primitives for coordination of multiple human-robot collaborative tasks. Auton. Robot. **41**(3), 593–612 (2017)
151. Yang, E., Mavrogiannis, C.: Implicit communication in human-robot collaborative transport (2025). arXiv preprint arXiv:2502.03346
152. Huang, Z., Liu, H., Wu, J., Lv, C.: Conditional predictive behavior planning with inverse reinforcement learning for human-like autonomous driving. IEEE Trans. Intell. Transp. Syst. **24**(7), 7244–7258 (2023)
153. Losey, D.P., Bajcsy, A., O'Malley, M.K., Dragan, A.D.: Physical interaction as communication: Learning robot objectives online from human corrections. Int. J. Robot. Res. **41**(1), 20–44 (2022)

154. Liu, C., Li, X., Li, Q., Xue, Y., Liu, H., Gao, Y.: Robot recognizing humans intention and interacting with humans based on a multi-task model combining ST-GCN-LSTM model and YOLO model. Neurocomputing **430**, 174–184 (2021)
155. Gao, X., Yan, L., Wang, G., Gerada, C.: Hybrid recurrent neural network architecture-based intention recognition for human-robot collaboration. IEEE Transactions on Cybernetics. **53**(3), 1578–1586 (2021)
156. Zhang, X., Tian, S., Liang, X., Zheng, M., Behdad, S.: Early prediction of human intention for human–robot collaboration using transformer network. J. Comput. Inf. Sci. Eng. **24**(5) (2024)
157. Domínguez-Vidal, J.E., Sanfeliu, A.: Exploring transformers and visual transformers for force prediction in human-robot collaborative transportation tasks. In: 2024 IEEE International Conference on Robotics and Automation (ICRA), pp. 3191–3197 (2024). IEEE
158. Domínguez-Vidal, J., Sanfeliu, A.: Force and velocity prediction in human-robot collaborative transportation tasks through video retentive networks. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 9307–9313 (2024). IEEE
159. Zhang, H., Zhuang, Z., Wang, X., Yang, X., Zhang, Y.: Mop-former: Motion-primitive transformer for wearable-sensor activity recognition (2025). arXiv preprint arXiv:2505.20744
160. Xiang, W., Haoteng, Y., Wang, H., Jin, X.: Socialcvae: Predicting pedestrian trajectory via interaction conditioned latents. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 38, pp. 6216–6224 (2024)
161. Barsoum, E., Kender, J., Liu, Z.: Hp-gan: Probabilistic 3D human motion prediction via gan. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 1418–1427 (2018)
162. Dermy, O., Chaveroche, M., Colas, F., Charpillet, F., Ivaldi, S.: Prediction of human whole-body movements with ae-promps. In: 2018 IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids), pp. 572–579 (2018). IEEE
163. Fang, C., Peternel, L., Seth, A., Sartori, M., Mombaur, K., Yoshida, E.: Human modeling in physical human-robot interaction: A brief survey. IEEE Robotics and Automation Letters. **8**(9), 5799–5806 (2023)
164. Hiatt, L.M., Narber, C., Bekele, E., Khemlani, S.S., Trafton, J.G.: Human modeling for human-robot collaboration. Int. J. Robot. Res. **36**(5–7), 580–596 (2017)
165. Maurice, P., Padois, V., Measson, Y., Bidaud, P.: Digital Human Modeling for Collaborative Robotics. In: DHM and Posturography, (2019). https://hal.science/hal-02389726
166. Magistris, G.D., Micaelli, A., Savin, J., Gaudez, C., Marsot, J.: Dynamic digital human models for ergonomic analysis based on humanoid robotics techniques. International Journal of the Digital Human. **1**(1), 81–109 (2015)
167. Seth, A., Hicks, J.L., Uchida, T.K., Habib, A., Dembia, C.L., Dunne, J.J., Ong, C.F., DeMers, M.S., Rajagopal, A., Millard, M., et al.: Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. PLoS Comput. Biol. **14**(7), 1006223 (2018)
168. Rajagopal, A., Dembia, C.L., DeMers, M.S., Delp, D.D., Hicks, J.L., Delp, S.L.: Full-body musculoskeletal model for muscle-driven simulation of human gait. IEEE Trans. Biomed. Eng. **63**(10), 2068–2079 (2016)
169. Prendergast, J.M., Balvert, S., Driessen, T., Seth, A., Peternel, L.: Biomechanics aware collaborative robot system for delivery of safe physical therapy in shoulder rehabilitation. IEEE Robotics and Automation Letters. **6**(4), 7177–7184 (2021)
170. Liu, L., Cooper, J.L., Ballard, D.H.: Computational modeling: Human dynamic model. Front. Neurorobot. **15**, 723428 (2021)
171. Herr, H., Popovic, M.: Angular momentum in human walking. J. Exp. Biol. **211**(4), 467–481 (2008)
172. Lanini, J., Duburcq, A., Razavi, H., Le Goff, C.G., Ijspeert, A.J.: Interactive locomotion: Investigation and modeling of physically-paired humans while walking. PLoS ONE **12**(9), 0179989 (2017)
173. Geyer, H., Herr, H.: A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities. IEEE Trans. Neural Syst. Rehabil. Eng. **18**(3), 263–273 (2010)
174. Mugge, W., Abbink, D.A., Schouten, A.C., Dewald, J.P., Van Der Helm, F.C.: A rigorous model of reflex function indicates that position and force feedback are flexibly tuned to position and force tasks. Exp. Brain Res. **200**(3), 325–340 (2010)
175. Burdet, E., Osu, R., Franklin, D., Yoshioka, T., Milner, T., Kawato, M.: A method for measuring endpoint stiffness during multi-joint arm movements. J. Biomech. **33**(12), 1705–1709 (2000)
176. Ajoudani, A., Fang, C., Tsagarakis, N., Bicchi, A.: Reduced-complexity representation of the human arm active endpoint stiffness for supervisory control of remote manipulation. Int. J. Robot. Res. **37**(1), 155–167 (2018)
177. Patel, H., O'Neill, G., Artemiadis, P.: On the effect of muscular cocontraction on the 3-D human arm impedance. IEEE Trans. Biomed. Eng. **61**(10), 2602–2608 (2014)
178. Bowden, R.: Learning statistical models of human motion. In: IEEE Workshop on Human Modeling, Analysis and Synthesis, CVPR, vol. 2000 (2000)
179. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-Or, D., Bermano, A.H.: Human motion diffusion model (2022). arXiv preprint arXiv:2209.14916
180. Peng, X.B., Abbeel, P., Levine, S., Panne, M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. ACM Transactions On Graphics (TOG). **37**(4), 1–14 (2018)
181. Hassan, M., Guo, Y., Wang, T., Black, M., Fidler, S., Peng, X.B.: Synthesizing physical character-scene interactions. In: ACM SIGGRAPH 2023 Conference Proceedings, pp. 1–9 (2023)
182. Liu, L., Ballard, D.: Humans use minimum cost movements in a whole-body task. Sci. Rep. **11**(1), 20081 (2021)
183. Li, S., Wang, Z., Yan, Z., Gao, Y., Jiang, H., Zheng, P.: Large language model for humanoid cognition in proactive human-robot collaboration. In: 2024 IEEE 20th International Conference on Automation Science and Engineering (CASE), pp. 540–545 (2024). IEEE
184. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3D human motion model for robust pose estimation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 11488–11499 (2021)
185. Xiao, Z., Wang, T., Wang, J., Cao, J., Zhang, W., Dai, B., Lin, D., Pang, J.: Unified human-scene interaction via prompted chain-of-contacts (2023). arXiv preprint arXiv:2309.07918
186. Wang, Y., Lin, J., Zeng, A., Luo, Z., Zhang, J., Zhang, L.: Physhoi: Physics-based imitation of dynamic human-object interaction (2023). arXiv preprint arXiv:2312.04393
187. Monari, E., Avallone, G., Valori, M., Agostini, L., Chen, Y., Palazzi, E., Vertechy, R.: Physical ergonomics monitoring in human-robot collaboration: A standard-based approach for hand-guiding applications. Machines. **12**(4), 231 (2024)
188. Keemink, A.Q., Kooij, H., Stienen, A.H.: Admittance control for physical human-robot interaction. Int. J. Robot. Res. **37**(11), 1421–1444 (2018)
189. Koppula, H.S., Saxena, A.: Anticipating human activities using object affordances for reactive robotic response. IEEE Trans. Pattern Anal. Mach. Intell. **38**(1), 14–29 (2015)
190. Gangapurwala, S., Geisert, M., Orsolino, R., Fallon, M., Havoutis, I.: Rloc: Terrain-aware legged locomotion using reinforcement

learning and optimal control. IEEE Trans. Rob. **38**(5), 2908–2927 (2022)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Gustavo A. Cardona** received the B.S. degree in Electrical Engineering and the M.S. degree in Industrial Engineering from the Universidad Nacional de Colombia, Bogotá, Colombia, and the Ph.D. degree in Mechanical Engineering from Lehigh University, Bethlehem, PA, USA, in 2025. His doctoral work on formal methods for multi-robot planning under uncertainty received the 2025 Outstanding Dissertation–Mechanical Engineering & Mechanics and the 2025 Elizabeth V. Stout Dissertation Award. He was also a research intern at Mitsubishi Electric Research Laboratories in 2024. He is currently a Postdoctoral Researcher at the University of Delaware, working on human-humanoid interaction. His research interests include formal methods, control theory, human-robot interaction, and machine learning.

**Shubham S. Kumbhar** received his dual degree (B.Tech. and M.Tech.) in Mechanical Engineering from the Indian Institute of Technology Bombay (IIT Bombay), India, in 2022. He is currently pursuing a Ph.D. degree in Mechanical Engineering at the University of Delaware. His research focuses on control frameworks for physical human-humanoid collaboration and shared autonomy in human-robot interaction tasks.

**Panagiotis Artemiadis** received the B.S. degree, and the M.S. and Ph.D. degrees in Mechanical Engineering from the National Technical University of Athens, Athens, Greece, in 2003 and 2009, respectively. His doctoral work on control methodologies for neuro-robotic systems received the Sarafi Award for the Best Ph.D. Thesis in 2008–2009. He was a Postdoctoral Research Associate at the Massachusetts Institute of Technology from 2009 to 2011 , and served as an Assistant and Associate Professor at Arizona State University from 2011 to 2019. He is currently a Professor of Mechanical Engineering and the Founder and Director of the Human-Oriented Robotics and Control (HORC) Lab at the University of Delaware. His research interests include neuro-robotics, human-robot interaction, rehabilitation robotics, and brain-machine interfaces.