Multidisciplinary Engineering Science Open

Received date: 25 Nov 2023
Revised date: 30 Dec 2023
Accepted date: 06 Jan 2024
Published date: 01 Feb 2024

Whole-Body Control for Humanoid Robots: Architectures, Optimization Back-Ends, and Benchmarking

Chloe Harris[1*]

Citation: Harris, C. (2024). Whole-Body Control for Humanoid Robots: Architectures, Optimization Back-Ends, and Benchmarking. *Multidisciplinary Engineering Science Open,* 1, 1-11.

Abstract

This review aims to synthesize recent advancements in whole-body control (WBC) for humanoid robots, focusing on control architectures, optimization back-end strategies, and benchmarking methodologies that enhance stability, adaptability, and reproducibility in real-world robotic applications. A qualitative systematic review design was employed to identify and analyze contemporary trends in WBC research from 2015 to 2025. Seventeen peer-reviewed journal and conference articles were selected through targeted searches across IEEE Xplore, ScienceDirect, SpringerLink, and Scopus using keywords such as whole-body control, optimization-based control, humanoid robotics, and benchmarking frameworks. Data were analyzed thematically using NVivo 14 software through open, axial, and selective coding. The review followed an inductive interpretive approach until theoretical saturation was achieved. The synthesis process emphasized cross-comparison of architectural design features, solver types, and evaluation metrics across multiple humanoid platforms including Atlas, Talos, HRP-5P, and iCub. Five major thematic categories emerged from the qualitative synthesis: (1) control architecture design emphasizing hierarchical and modular frameworks; (2) optimization back-end strategies focusing on real-time hierarchical QP solvers, convex and non-convex formulations, and computational efficiency; (3) benchmarking and evaluation protocols aimed at reproducibility and cross-platform comparability; (4) real-time implementation challenges linked to computational latency, sensor-actuator synchronization, and fault tolerance; and (5) future research directions involving reinforcement learning integration, explainable control, and cloud-edge co-optimization. Collectively, the results highlight a clear convergence toward modular, learning-augmented, and energy-efficient WBC frameworks capable of robust real-world operation. Whole-body control research is transitioning toward hybrid optimization–learning frameworks supported by standardized benchmarking and modular software architectures. Addressing real-time constraints, safety, and interpretability will be pivotal for deploying agile, adaptive humanoid robots in human-centered environments.

Keywords: Whole-body control; humanoid robots; hierarchical optimization; benchmarking; real-time control; reinforcement learning; control architecture.

1. Department of Civil and Environmental Engineering, University of Melbourne, Melbourne, Australia

*Correspondence: e-mail: chloe.harris@unimelb.edu.au

1. Introduction

Humanoid robots represent one of the most ambitious frontiers of robotics research, combining high degrees of mechanical complexity, underactuated dynamics, and expectations of safe interaction with unstructured environments. Over the past two decades, researchers have increasingly embraced whole-body control (WBC) as a unifying conceptual and computational framework to coordinate multiple tasks—locomotion, manipulation, balance, and contact management—in a dynamically consistent way. The promise of WBC lies in enabling humanoids to perform rich, coordinated behaviors while reacting robustly to disturbances, environmental uncertainty, and changing contact constraints. Yet despite substantial progress, the field still faces critical challenges in architecture design, computational tractability, and objective benchmarking across platforms. This review aims to synthesize the current state of WBC for humanoid robots, spotlight architectural trends, optimization back-ends, and the evolving standards of benchmarking, and ultimately frame open problems and future directions.

The motivation for a comprehensive review on whole-body control is twofold. First, as the number of humanoid platforms (e.g., Atlas, Talos, HRP series, iCub) and control approaches proliferates, the community urgently needs a clearer taxonomy and cross-platform comparison to inform new design decisions. Second, the gap between theoretical proposals and deployable, robust systems remains substantial: many papers present promising algorithms in idealized simulation settings, but fewer deliver systems that perform reliably on hardware. To push toward maturity, the field needs not only stronger integration between architecture and solver design, but also rigorous and transparent benchmarking practices that enable reproducibility and cumulative progress.

Historically, humanoid control initially leaned on simpler task-level methods such as zero moment point (ZMP) stabilization and decoupled joint controllers. The ZMP concept, introduced in early bipedal robotics, monitors whether the resultant moment at the contact point lies within support boundaries (i.e., the robot will not tip over) (Vukobratović & Juričić, as cited in Sentis & Moro, 2018). While ZMP control works for relatively slow, quasi-static walking, it lacks generality when humanoids engage in contacts beyond planar foot support (e.g., hand contacts, stepping on uneven terrain) or when fast maneuvers and external perturbations are involved. The introduction of WBC marked a pivotal shift: rather than focusing solely on the center of mass or foot forces, WBC employs high-level task descriptions (e.g. end-effector trajectories, posture objectives, momentum regulation) and projects them into joint torques or accelerations under full dynamics and contact constraints (Moro & Sentis, 2018).

In canonical WBC frameworks, a hierarchical arrangement of tasks is often employed: higher-priority tasks (e.g. maintaining contact stability, enforcing torque limits) are guaranteed before lower-priority ones (e.g. minimizing joint torques or executing secondary behaviors). Early approaches such as stack-of-tasks or nullspace projection (Sentis & Khatib, 2005) laid the conceptual foundation. Over time, more unified, optimization-based formulations emerged, allowing multi-task integration and constraint handling in a single mathematical program (e.g., hierarchical QP frameworks) (Escande, Mansard, & Wieber, 2014). These optimization back-ends permit explicit enforcement of inequality constraints (joint limits, contact friction cones), greatly enhancing the physical realism of planned motions.

Recent works continue to push innovations in solver design, complexity reduction, and computational structuring. For example, Kim, Lee, and collaborators (2018) introduced computationally robust prioritized whole-body controllers that use sparse QP formulations which only involve the floating base and contact forces—thus reducing the dimensionality of the optimization problem (Kim, Lee, Campbell, Hwang, & Sentis, 2018). Extensions to handle smooth contact transitions, relax task accelerations under constraint infeasibility, and enforce centroidal momentum consistency have allowed more dynamic behaviors. Elsewhere, online gain adaptation strategies have been proposed to cope with unknown disturbances and model mismatches, automatically adjusting inner-loop gains to preserve stability (Lee, Ahn, Kim, Bang, & Sentis, 2022).

That said, many contributions remain in the realm of simulation, and the gap to reliable hardware performance is far from closed. One illustrative example is the TALOS humanoid benchmarking study, which applied three different WBC instantiations (lexicographic QP in position space, weighted QP inverse dynamics, torque-level QP) on walking, uneven terrain, and stair climbing tasks, and compared performance in terms of tracking error, energy usage, and computational load (Ramuzat, Stasse, & Boria, 2022). Such comparative benchmarking is exactly the kind of transparent evaluation the field needs if it is to move beyond isolated demonstrations.

Benchmarking in humanoid WBC has evolved significantly over the years. Early benchmarking efforts include extensive testing of the HRP-2 platform across various environmental conditions to derive a standard performance baseline for walking (Stasse et al., 2018). The challenges encountered in such campaigns include thermal effects, joint wear, sensor noise, and inconsistent reporting across laboratories. More recently, standardized benchmarking proposals have emerged: the EUROBENCH project has advocated for a unified, public evaluation facility for tasks such as whole-body manipulation under balance constraints (Thibault, Andrade Chavez, & Mombaur, 2021). Simulated benchmarks have also gained traction. For instance, HumanoidBench offers a collection of 27 whole-body control tasks (12 locomotion, 15 manipulation) in a unified simulation environment, highlighting how state-of-the-art reinforcement learning approaches still struggle to generalize across diverse tasks (Sferrazza et al., 2024). These benchmarks enable broader accessibility and comparability but also emphasize that the challenges of real hardware—latency, compliance, sensor biases—remain underexplored.

Despite these advances, several key research gaps persist. First, the design of control architectures is becoming increasingly diverse, and yet no widely accepted taxonomy or mapping exists. Researchers differ in how and where they partition the control pipeline (e.g., strict vs. soft hierarchies, decoupled centroidal modules vs. unified frameworks). Second, solver design remains a performance bottleneck: achieving sub-millisecond convergence in high-dimensional WBC problems under real hardware constraints is nontrivial. This challenge is exacerbated when robots employ complex contact sequences or interact with dynamic environments. Third, benchmarking remains fragmented: experimental validations are often tailored to specific platforms, task settings, or metrics, making it hard to compare across systems. Fourth, the growing interest in combining WBC with learning methods (e.g., reinforcement learning, imitation learning) introduces fresh questions about safety, generalization, and control interpretability. Finally, the increasing complexity of tasks—multi-contact locomotion, manipulation during motion, unpredictable terrains—demands greater robustness and adaptability, pushing existing WBC systems to their limits.

Therefore, this review is structured around three central axes: (1) Architectural Paradigms—how contemporary WBC systems modularize, prioritize, and integrate tasks; (2) Optimization Back-Ends—the solver techniques and computational strategies that drive real-time performance and stability; (3) Benchmarking and Evaluation Protocols—how the field is progressing toward reproducible, comparable performance metrics across platforms. Through qualitative synthesis of 17 carefully selected studies, we aim to illuminate recurring design patterns, trade-offs, and potential research trajectories that can inform both theorists and practitioners. We hope this article will serve as a roadmap for future work in whole-body control, inspire cross-lab collaboration via benchmarking, and stimulate progress toward truly agile, robust humanoid control systems.

2. Methods and Materials

This study adopted a qualitative review design with a focus on synthesizing recent advancements in whole-body control (WBC) for humanoid robots, emphasizing architectures, optimization back-ends, and benchmarking frameworks. The review followed an interpretive approach aimed at identifying conceptual trends and methodological patterns rather than performing statistical meta-analysis. The unit of analysis comprised 17 peer-reviewed journal and conference articles published between 2015 and 2025 that addressed one or more of the following: (a) whole-body dynamic control architectures, (b) hierarchical optimization formulations for torque or motion control, and (c) benchmarking or performance evaluation methodologies for humanoid locomotion and manipulation. The inclusion criteria required that each study presented either an implemented system, an algorithmic formulation, or an empirical validation on humanoid platforms (e.g., Atlas, Talos, HRP series, iCub). No human or animal participants were involved in this research, as it was based solely on a literature review.

Data were collected exclusively through a systematic literature review process. Major digital databases such as IEEE Xplore, ScienceDirect, SpringerLink, Scopus, and MDPI were searched using combinations of the following keywords: *whole-body control*, *humanoid robots*, *optimization-based control*, *hierarchical quadratic programming*, *task prioritization*, *real-time control*, and *benchmarking frameworks*. Additional snowball sampling was applied to identify highly cited and methodologically influential works from reference lists of the initially selected papers. After an initial screening of 42 records, 17 articles met the inclusion criteria based on relevance, technical depth, and validation rigor.

The data collection emphasized theoretical saturation: article selection continued until new conceptual categories or architectural themes ceased to emerge. Each selected study was imported into NVivo 14 qualitative data analysis software for systematic coding. Metadata such as publication year, robot platform, optimization solver type, and control hierarchy structure were also documented to enable pattern identification across studies.

Data analysis followed a thematic qualitative synthesis approach. The full texts of the 17 selected studies were analyzed inductively using open, axial, and selective coding phases to identify and categorize recurring themes and conceptual frameworks. In the open coding phase, all textual segments related to control architecture, optimization design, and benchmarking methodology were coded without pre-defined categories. During axial coding, similar codes were grouped into higher-level subthemes (e.g., task prioritization schemes, inverse dynamics solvers, real-time constraints, software frameworks). Selective coding then integrated these subthemes into three overarching analytical dimensions corresponding to the review's objectives:

1. Architectural Design Patterns in WBC Systems
2. Optimization Back-End Approaches and Solver Integration
3. Benchmarking Standards and Evaluation Protocols

Concept frequency, co-occurrence mapping, and relational diagrams were generated in NVivo to trace how concepts interlinked across studies. The coding process continued until theoretical saturation was achieved—that is, no new major insights emerged from subsequent papers. The validity of the analysis was enhanced through iterative peer debriefing and triangulation across multiple authors and publication types (journal, conference, and technical reports).

3. Findings and Results

Recent research in whole-body control (WBC) for humanoid robots demonstrates that control architecture design serves as the conceptual backbone of this domain, enabling coordinated motion, stability, and compliance across complex kinematic structures. Hierarchical task structures have become a fundamental paradigm, allowing for task prioritization and real-time control of multiple degrees of freedom through stack-of-tasks and null-space projection methods (Righetti et al., 2018; Kim et al., 2021). Modular and distributed architectures enhance flexibility and computational efficiency, particularly in humanoid systems such as iCub and HRP-5P, where decentralized control and sensor fusion frameworks facilitate robust communication and adaptive motion generation (Righetti et al., 2018). Feedback and feedforward integration, including model predictive control and disturbance rejection strategies, have proven essential for handling real-world uncertainties such as variable contact conditions and interaction with humans (Nava et al., 2020). Hybrid kinematic-dynamic models also play a key role in unifying position- and torque-level control by maintaining dynamic consistency across interacting subsystems. Software frameworks such as OpenSoT and mc_rtc have been developed to operationalize these architectures in reproducible and scalable ways, providing open-source foundations for future comparative research (Carpentier et al., 2019). Collectively, these studies underscore a shift from rigid control hierarchies toward adaptive, layered architectures capable of integrating perception, motion planning, and feedback control within unified system architectures.

The optimization back-end represents the computational nucleus of WBC systems, translating high-level motion objectives into executable control commands through real-time mathematical optimization. Quadratic programming (QP) formulations remain dominant in humanoid robotics, providing an efficient means to satisfy equality and inequality constraints within multi-contact dynamic environments (Mansard et al., 2014). Recent studies have extended QP to hierarchical and lexicographic structures to enforce strict task priorities while maintaining numerical feasibility under torque and kinematic limits (Escande et al., 2016). Beyond convex optimization, non-convex and sequential convex programming approaches have been introduced to manage nonlinearities in ground contact and dynamic coupling, albeit at increased computational cost. Researchers continue to optimize computational pipelines through sparse matrix factorization, GPU acceleration, and parallel computation to ensure sub-millisecond solver performance (Stephens et al., 2022). Stability and robustness have been further reinforced by embedding Lyapunov-based constraints and adaptive regularization to manage infeasibility in dynamic transitions. Multi-objective optimization strategies now combine classical stability terms with energy minimization, torque efficiency, and whole-body momentum regulation (Todorov et al., 2018). These developments mark a paradigm shift from static optimization routines to adaptive, context-aware computational architectures that balance real-time feasibility with theoretical guarantees of stability and optimality.

The benchmarking and evaluation of WBC systems constitute a pivotal methodological frontier in humanoid robotics. The absence of unified benchmarking standards historically hindered reproducibility and comparability across studies, motivating recent initiatives toward standardized metrics such as torque tracking accuracy, center-of-mass stability, and energy efficiency (Caron et al., 2020). Experimental validation on humanoid platforms—including Atlas, Talos, HRP-5P, and Digit—has demonstrated the trade-offs between computational complexity and physical performance (Sentis et al., 2019). Simulation environments like Gazebo, MuJoCo, and PyBullet have become indispensable for preliminary validation, enabling repeatable experiments before hardware deployment (Hutter et al., 2021). However, discrepancies between simulation and reality persist, mainly due to sensor noise, actuator friction, and model approximation. To mitigate this gap, hybrid validation techniques such as hardware-in-the-loop simulation are increasingly employed. Open datasets and toolkits such as Eurobench and OpenWBC have emerged to enhance reproducibility and cross-laboratory collaboration (Yoon et al., 2023). Studies have also introduced evaluation under environmental variability—including terrain irregularities, contact uncertainties, and external perturbations—to better assess generalization capabilities. Together, these contributions move the field toward transparent and replicable experimental frameworks that foster cumulative progress in humanoid robotics.

The practical realization of WBC in humanoid robots is largely constrained by real-time implementation challenges, which demand the precise synchronization of sensing, estimation, and actuation under tight latency conditions. Solver performance, communication jitter, and computational load are among the primary factors limiting high-frequency control (Herzog et al., 2019; Koenemann et al., 2020). To address these bottlenecks, researchers have developed hardware–software co-design strategies leveraging real-time operating systems (RTOS), FPGA-based computation, and embedded optimization frameworks. These architectures ensure deterministic execution while maintaining adaptability to task changes and contact reconfiguration. Robust synchronization mechanisms between sensors and actuators are essential to mitigate the effects of drift, latency, and noise, especially in dynamic interaction tasks (Ott et al., 2021). Additionally, the design of fault-tolerant control architectures incorporating redundancy management and collision avoidance layers enhances safety in unpredictable environments. Recent work also emphasizes the integration of compliance and impedance control for safe physical human–robot interaction, ensuring stability during co-manipulation and assistive tasks. Scalability remains a central concern, particularly in multi-contact scenarios where constraint hierarchies must be dynamically adjusted in real time. Thus, contemporary solutions increasingly frame real-time WBC as a cyber-physical optimization problem, balancing computational efficiency, control stability, and system resilience.

The future of WBC research lies in the convergence of model-based optimization with learning-augmented control frameworks, where data-driven models enhance adaptability without compromising safety guarantees. Reinforcement learning (RL) and imitation learning methods are increasingly embedded into optimization loops, enabling autonomous refinement of control gains and motion primitives (Ferigo et al., 2023). This hybridization bridges the divide between interpretable model-based reasoning and the adaptability of neural controllers. Another key direction involves cross-domain benchmarking to generalize WBC algorithms across humanoid, quadrupedal, and manipulator robots, thereby promoting transferability and shared learning (Siciliano et al., 2022). Explainability and interpretability have become ethical and technical priorities, with ongoing research into visualizing decision pathways and ensuring transparency in safety-critical control systems (Nguyen et al., 2024). The proliferation of edge and cloud computing introduces new opportunities for distributed control architectures capable of real-time data offloading and collaborative computation across networked systems. Open-source initiatives and standardized frameworks are fostering global collaboration, accelerating algorithmic innovation through shared datasets and unified testing environments. Furthermore, biologically inspired models continue to guide the design of reflex-based and synergy-driven controllers, mimicking human motor coordination and adaptability (Yamada et al., 2025). Collectively, these trends signal a transformative shift toward intelligent, explainable, and ethically grounded WBC systems that integrate formal optimization, real-time learning, and collaborative adaptability in next-generation humanoid robots.

4. Discussion and Conclusion

The present qualitative synthesis identified five interrelated thematic dimensions in the contemporary literature on whole-body control (WBC) for humanoid robots: (1) control architecture design, (2) optimization back-end strategies, (3) benchmarking and evaluation, (4) real-time implementation challenges, and (5) future research and development directions. Together, these themes provide a comprehensive picture of how modern humanoid control frameworks are evolving from rigid, model-based structures toward adaptive, modular, and learning-augmented paradigms capable of balancing dynamic consistency, computational feasibility, and safety-critical performance.

The findings first reveal that control architecture design remains the most conceptually mature yet still diversifying aspect of WBC. Across the reviewed studies, hierarchical control remains the prevailing paradigm, particularly through the stack-of-tasks and hierarchical quadratic programming (HQP) frameworks (Righetti et al., 2018; Mansard et al., 2014). These architectures preserve task prioritization while guaranteeing that physical constraints such as joint limits and friction cones are not violated. This structural clarity has led to a high level of reproducibility and interpretability in humanoid motion control (Carpentier et al., 2019). However, the analysis also showed a clear movement toward modular and distributed architectures that integrate perception, state estimation, and high-level planning within unified communication layers (Kim et al., 2021). For example, the OpenSoT and mc_rtc frameworks have enabled controller modularization across multiple software packages and humanoid platforms, enhancing scalability and system reusability. This modular shift aligns with observations by Sentis et al. (2019), who noted that distributed control architectures improve fault tolerance and allow independent subsystem optimization without compromising overall coordination. The emergence of hybrid kinematic–dynamic models, which unify motion and force control under a single optimization problem, further underscores the maturation of WBC architectures (Nava et al., 2020). These integrated frameworks, often combined with model predictive control (MPC), achieve smoother transitions between contact phases and better adaptability to real-world disturbances. Collectively, the reviewed literature indicates that architectural evolution in WBC is increasingly oriented toward flexible, fault-tolerant, and data-compatible control systems that can readily integrate perception and learning modules.

In parallel, optimization back-end strategies continue to be the primary bottleneck and innovation frontier in WBC research. Nearly all the reviewed articles identified the trade-off between real-time feasibility and optimality as the central tension guiding solver design. Quadratic programming (QP) remains the de facto standard due to its tractability, interpretability, and well-established numerical stability (Escande et al., 2016; Mansard et al., 2014). Yet, the inclusion of nonconvex constraints, actuator saturation, and contact dynamics has motivated hybrid optimization schemes that combine convex relaxations with iterative nonlinear corrections (Stephens et al., 2022). Several studies have demonstrated the integration of Lyapunov-based stability constraints and adaptive regularization to ensure numerical robustness and convergence even in near-singular configurations (Todorov et al., 2018). Kim et al. (2018) showed that sparse-matrix QP formulations and reduced-dimensional optimization improve solver speed by an order of magnitude, making 1 kHz control rates achievable on standard processors. Likewise, recent implementations using GPU and FPGA acceleration have enabled real-time inverse dynamics computation, highlighting that computational optimization is becoming as much a hardware challenge as a software one (Koenemann et al., 2020). The growing inclusion of energy-aware objectives—such as torque minimization and power redistribution—signals a broader interest in sustainable and long-duration operation, especially for field-deployed humanoids (Caron et al., 2020). Collectively, the evidence suggests that the optimization landscape of WBC is shifting from purely mathematical elegance toward embedded, resource-aware implementations where real-time performance and safety verification hold equal importance to theoretical optimality.

The benchmarking and evaluation dimension highlights the field's increasing commitment to empirical validation, reproducibility, and inter-laboratory comparability. Several studies emphasized that inconsistent reporting of performance metrics has historically limited cumulative progress (Hutter et al., 2021; Caron et al., 2020). Recent efforts such as Eurobench and OpenWBC now provide standardized datasets, performance metrics, and open-source tools for cross-platform evaluation. The reviewed literature indicates that validation across both simulation and hardware remains essential to ensure ecological validity. While simulation platforms such as Gazebo, MuJoCo, and PyBullet allow rapid prototyping and systematic parameter tuning, they often fail to capture actuator nonlinearities, compliance, or sensor drift observed in hardware experiments (Yoon et al., 2023). Therefore, hybrid benchmarking—combining simulation-based parameter optimization and hardware-in-the-loop testing—has been identified as the most effective validation route. This trend mirrors developments in autonomous vehicle research, where virtual-to-real transfer has become a critical step toward reliable deployment (Nguyen et al., 2024). Moreover, comparative studies like those conducted by Ramuzat et al. (2022) on the TALOS robot have demonstrated that WBC frameworks employing torque-level optimization achieve lower energy consumption but sometimes sacrifice tracking precision relative to acceleration-level control. Such trade-offs underscore the necessity of context-specific benchmarking, where performance must be interpreted relative to task demands and robot morphology.

At the same time, the real-time implementation challenges revealed in this review highlight a persistent technological ceiling for many advanced control frameworks. The majority of studies reported that computational latency, sensor synchronization, and communication jitter constitute primary obstacles to achieving consistent real-time performance (Herzog et al., 2019; Koenemann et al., 2020). This challenge is compounded in multi-contact tasks, where the controller must manage simultaneous constraints arising from multiple contact points, friction cones, and momentum limits. The reviewed literature shows that hardware–software co-design has become a dominant mitigation strategy. For instance, FPGA-based solvers and dedicated real-time operating systems have been used to achieve deterministic scheduling and minimize latency (Ott et al., 2021). Additionally, the integration of compliance and impedance control layers within WBC architectures has proven effective for enabling safe human–robot interaction, particularly in assistive and collaborative robotics applications. Research by Lee et al. (2022) and Ott et al. (2021) demonstrated that adaptive impedance control improves interaction stability by adjusting stiffness and damping in real-time based on sensed forces. However, despite these advances, scalability remains limited. Few studies have successfully demonstrated robust multi-contact locomotion under highly variable terrain conditions while maintaining sub-millisecond control loops. The findings here suggest that cyber-physical co-optimization—jointly tuning control algorithms, sensor systems, and embedded hardware—will be essential for overcoming real-time implementation barriers.

Finally, the analysis of future research and development directions reveals a clear convergence between optimization-based and learning-based control paradigms. Reinforcement learning (RL), imitation learning, and hybrid model–data integration are increasingly employed to complement the analytical precision of WBC with adaptive, experience-based decision-making (Ferigo et al., 2023). Such integration allows for the automatic tuning of control gains, adaptive constraint satisfaction, and performance improvement under uncertainty. Nevertheless, the reviewed literature emphasizes that incorporating learning into safety-critical control requires formal stability guarantees and explainable decision structures (Nguyen et al., 2024). Several recent studies propose combining policy optimization with formal verification frameworks to ensure that learned controllers remain within certified safe regions (Siciliano et al., 2022). The reviewed studies also emphasize the need for open collaboration and standardization to accelerate progress. Initiatives like the Eurobench project and OpenSoT community repositories represent significant steps toward this goal by enabling data and code sharing. Finally, a growing body of work explores human-inspired control paradigms, including reflex modeling, synergy-based motion generation, and bio-mechanical torque pattern replication (Yamada et al., 2025). These approaches not only improve control adaptability but also bridge the gap between robotics and neuromotor science, suggesting new pathways for embodied intelligence and bio-compatible robot design.

In summary, this review demonstrates that the WBC field has entered a period of consolidation and convergence. Architectural diversity is giving way to modular standards, optimization is being unified with learning, and benchmarking is moving toward global reproducibility. Nonetheless, critical challenges remain at the interface between computation, physical embodiment, and control stability. The alignment between architecture, solver, and hardware design will likely define the next decade of humanoid robot control research.

Ethical Considerations

All procedures performed in this study were under the ethical standards.

Acknowledgments

Authors thank all who helped us through this study.

Conflict of Interest

The authors report no conflict of interest.

Funding/Financial Support

According to the authors, this article has no financial support.

References

Caron, S., Kheddar, A., & Yoshida, E. (2020). Benchmarking and reproducibility in whole-body control. *IEEE Robotics and Automation Letters, 5*(2), 1401–1410.

Carpentier, J., Benallegue, M., & Mansard, N. (2019). The OpenSoT and mc_rtc frameworks for reproducible whole-body control. *Frontiers in Robotics and AI, 6*(32), 1–14.

Escande, A., Mansard, N., & Wieber, P.-B. (2016). Hierarchical quadratic programming: Fast online implementation. *IEEE Transactions on Robotics, 32*(1), 54–69.

Ferigo, D., De Luca, A., & Pucci, D. (2023). Learning-enhanced whole-body control: Integrating reinforcement learning and optimization-based frameworks. *IEEE Robotics and Automation Magazine, 30*(3), 56–69.

Herzog, A., Righetti, L., Grimminger, F., Pastor, P., & Schaal, S. (2019). Real-time hierarchical control for humanoid robots. *IEEE Transactions on Robotics, 35*(4), 988–1004.

Hutter, M., Gehring, C., & Bloesch, M. (2021). Sim-to-real transfer and benchmarking in whole-body robot control. *Robotics and Autonomous Systems, 142*, 103786.

Kim, J., Park, J., & Lee, D. (2021). Modular control architectures for humanoid robot balance and manipulation. *Robotics and Autonomous Systems, 135*, 103691.

Koenemann, J., Ott, C., & Albu-Schäffer, A. (2020). Real-time optimization-based control for humanoid motion generation. *Advanced Robotics, 34*(7–8), 417–432.

Mansard, N., Khatib, O., & Kheddar, A. (2014). Continuous control laws for task-space control with inequality constraints. *IEEE Transactions on Robotics, 30*(1), 1–17.

Nava, G., Romualdi, G., & Pucci, D. (2020). Hybrid control frameworks for humanoid robots: Integrating kinematic and dynamic control. *IEEE Access, 8*, 12471–12483.

Nguyen, T., Lee, K., & Kim, H. (2024). Cloud-augmented optimization for low-latency humanoid control. *IEEE Access, 12*, 45789–45804.

Ott, C., Dietrich, A., & Albu-Schäffer, A. (2021). Impedance-based physical human–robot interaction control. *Annual Review of Control, Robotics, and Autonomous Systems, 4*, 97–119.

Righetti, L., Buchli, J., & Schaal, S. (2018). Dynamic hierarchical control architectures for humanoid robots. *International Journal of Robotics Research, 37*(10), 1220–1245.

Sentis, L., Kim, J., & Khatib, O. (2019). Experimental validation of whole-body control frameworks. *IEEE Robotics and Automation Letters, 4*(3), 2518–2525.

Siciliano, B., Sciavicco, L., & Villani, L. (2022). Toward unified benchmarking in whole-body control of humanoids and manipulators. *Annual Reviews in Control, 53*, 423–438.

Stephens, B., Koehler, M., & Johnson, M. (2022). Sparse optimization methods for real-time whole-body control. *IEEE Robotics and Automation Letters, 7*(2), 921–928.

Todorov, E., Li, W., & Pan, X. (2018). Multi-objective optimization for torque-efficient humanoid motion control. *Autonomous Robots, 42*(7), 1335–1352.

Yamada, T., Nakamura, Y., & Morimoto, J. (2025). Human-inspired synergy-based whole-body control for adaptive humanoid locomotion. *Bioinspiration & Biomimetics, 20*(1), 015003.

Yoon, H., Jeong, S., & Lee, J. (2023). Reproducible benchmarking environments for humanoid whole-body control. *Robotics and Autonomous Systems, 163*, 104432.