

---

# Redefining Robot Generalization Through Interactive Intelligence

---

Sharmita Dey

## Abstract

Recent advances in large-scale machine learning have produced high-capacity “foundation models” capable of adapting to a broad array of downstream tasks. While such models hold great promise for robotics, the prevailing paradigm still portrays robots as single, autonomous decision-makers, performing tasks like manipulation and navigation, with limited human involvement. However, a large class of real-world robotic systems, including wearable robotics (e.g., prostheses, orthoses, exoskeletons), teleoperation, and neural interfaces, are semiautonomous, and require ongoing interactive coordination with human partners, challenging single-agent assumptions. In this position paper, we argue that robot foundation models must evolve to an *interactive multi-agent* perspective in order to handle the complexities of real-time human-robot co-adaptation. We propose a generalizable, neuroscience-inspired architecture encompassing four modules: (1) a multimodal *sensing module* informed by sensorimotor integration principles, (2) an ad-hoc teamwork model reminiscent of joint-action frameworks in cognitive science, (3) a predictive world belief model grounded in internal model theories of motor control, and (4) a memory/feedback mechanism that echoes concepts of Hebbian and reinforcement-based plasticity. Although illustrated through the lens of *cyborg systems*, where wearable devices and human physiology are inseparably intertwined, the proposed framework is broadly applicable to robots operating in semi-autonomous or interactive contexts. By moving beyond single-agent designs, our position emphasizes how foundation models in robotics can achieve a more robust, personalized, and anticipatory level of performance.

## 1. Introduction

Over the past few years, the field of AI has been profoundly influenced by *foundation models*, which are large, high-capacity neural networks pre-trained on extensive and heterogeneous datasets (Brown et al., 2020; Achiam et al., 2023). These models, exemplified by large language models (LLMs) such as GPT-4, and large multimodal models (LMMs) like PaLM-E (Driess et al., 2023), offer a flexible interface for perception, reasoning, and action. In robotics, foundation models have been applied to unify diverse tasks, e.g., manipulation, navigation, or object recognition, under a single policy (Reed et al., 2022; Brohan et al., 2022), often operating in a *single-agent* paradigm where the robot acts autonomously, under minimal human involvement.

However, a large class of real-world robotics, particularly those involving continuous human collaboration, are inherently *multi-agent*. Applications such as teleoperation (Lichardopol, 2007; Okamura, 2004; Kofman et al., 2005), prosthetic devices (Best et al., 2023; Cimolatto et al., 2022; Quintero et al., 2017; Wen et al., 2019; Dey et al., 2020; Dey & Schilling, 2022), exoskeletons (Molteni et al., 2018; Rosen & Perry, 2007; Lo & Xie, 2012; Dey et al., 2023), neural interfaces (Jackson & Zimmermann, 2012; Donoghue, 2008; Schultz & Kuiken, 2011; Vogel et al., 2020), brain-computer interfaces (Wolpaw, 2013; Nicolas-Alonso & Gomez-Gil, 2012; McFarland & Wolpaw, 2017; Abiri et al., 2019), and other semi-autonomous systems (Suchan & Osterloh, 2023; Clark & Feng, 2015) require ongoing co-adaptation with humans or other participating agents in the environment, rather than isolated, one-shot instructions. In these contexts, the single-agent perspective encounters significant limitations: it struggles to interpret dynamic and evolving user states and fails to handle *non-stationary* factors such as shifting goals, user fatigue, and changing environmental conditions.

Human-interactive robotics and cyborgs, in particular, demand continuous, bidirectional feedback loops between the human user and the device. At each step, the device must integrate physiological signals (e.g., electromyography, joint kinematics), user preferences, and environmental cues to ensure safe and comfortable completion of user commands. Over time, both the human and the device need to learn to function as a coordinated pair. This complexity aligns more closely with

---

<sup>0</sup>Corresponding author: Sharmita Dey (contact.deysharmita@gmail.com)

*neuroscience-based* perspectives on sensorimotor control, which emphasize dynamic feedback loops, internal predictive models, and adaptive synergy between multiple interacting systems (e.g., brain, muscles, external supports) (Wolpert et al., 1995; Kawato, 1999). Such interactions necessitate a *multi-agent interaction modeling* approach, rather than the traditional model of single-agent autonomy, where both the user and the device function as separate decision-making entities.

Consequently, this paper takes a position: *future robot foundation models must adopt an interactive multi-agent framework*, explicitly modeling both the robot and its counterpart (human or environment) as actively adapting agents. We illustrate this perspective through wearable robotics or “cyborg systems,” where a user’s physiological signals (e.g., EMG, joint angles) continuously intertwine with the device’s actuation. By incorporating principles from neuroscience, cognitive science, and multi-agent systems, this position paper aims to establish a foundation for next-generation interactive robotics, such as cyborgs and wearable robots, enabling co-adaptation, comfort, and anticipatory control. We suggest that the widespread adoption of interactive, multi-agent paradigms in robotic foundation models will lead to fundamentally safer, more robust, and more user-centric performance, surpassing what is possible within the single-agent paradigm.

Throughout this paper, we use the term “cyborg systems” (or “wearable robotics”) to refer broadly to human–robot integrated device, such as prostheses, exoskeletons, brain-computer interfaces, and other assistive technologies that physically couple with the user’s body and neural signals. The paper is organized as follows: Section 2 reviews current robot foundation models, their single-agent limitations, and introduces multi-agent systems and ad-hoc teamwork. Section 3 explores the requirements of human-interactive scenarios and the shortcomings of single-agent approaches. Section 4 presents our position and proposes an interactive, neuroscience-inspired multi-agent architecture for cyborg systems. Section 5 extends this framework to other human-interactive robotic contexts. Section 7 addresses safety, ethical, and regulatory considerations. The paper concludes in Section 8.

## 2. Background

### 2.1. Robot Foundation Models and Their Single-Agent Roots

The advent of Large Language Models (LLMs) like GPT-4 (Achiam et al., 2023), LLaMA (Touvron et al., 2023), and Vision-Language Models (VLMs) like CLIP (Radford et al., 2021), BLIP (Li et al., 2022), BLIP-2 (Li et al., 2023) has significantly advanced robotics by enhancing perception, planning, and action generation capabilities. These models demonstrate exceptional abilities in understanding and generating multimodal data, which are crucial for complex robotic tasks (Driess et al., 2023; Brohan et al., 2022; Team et al., 2024). By leveraging the robust linguistic capabilities of LLMs, robots can interpret and execute tasks based on natural language commands, eliminating the need for complex programming interfaces. For instance, robots can parse instructions such as “Bring the red cup from the kitchen table” into structured subtasks involving object identification, navigation, and manipulation (Brown et al., 2020).

Despite the impressive capabilities of modern robot foundation models, such as Gato (Reed et al., 2022) and RT-1 (Brohan et al., 2022), their predominantly single-agent framework can limit performance in scenarios requiring tight coordination with humans or other agents. These models typically learn policies under the assumption that the robot operates largely on its own, taking in sensory inputs and issuing motor commands without ongoing, interactive feedback from a collaborator or user. Although they have achieved notable results on tasks like manipulation, navigation, and even some language grounding, key shortcomings emerge when real-time collaboration or continuous human guidance is essential.

#### 2.1.1. LIMITATIONS OF SINGLE-AGENT FOUNDATION MODELS IN ROBOTICS

**1) Inability to Handle Mid-Task User Corrections.** Gato (Reed et al., 2022), for example, was demonstrated on multiple discrete and continuous control tasks, ranging from Atari gameplay to real-world robotic arm manipulation. However, it was not designed to handle situations where a human might intervene mid-task with corrective feedback or dynamically changing instructions (e.g., “Wait, do not place the block there, hand it to me instead.”). As a single-agent learner, Gato follows its end-to-end policy after receiving an initial goal or observation. If the human’s intent shifts *during* task execution, Gato cannot seamlessly incorporate that feedback without externally resetting or retraining the policy. A *multi-agent* perspective, by contrast, would treat the user as a parallel decision-maker; the system would maintain a belief state about the user’s evolving instructions, thus adapting plans in real time rather than requiring full restarts.

**2) Lack of Human-Robot “Turn-Taking” in RT-1.** RT-1 (Brohan et al., 2022) demonstrated strong performance across a variety of robot manipulation tasks, incorporating visual inputs and token-based action outputs to execute pick-and-place

operations in real-world settings. While it accommodates sensorimotor fusion, the model expects minimal or one-shot user directives, akin to specifying a high-level goal. If a human spontaneously modifies the task (e.g., “Switch from picking an object to now opening a drawer.”), RT-1 has no built-in mechanism to *negotiate* or *re-evaluate* goals on the fly. A multi-agent framework would allow the robot to engage in turn-taking with the human, ask clarifying questions when it detects ambiguous instructions, and update its action policy based on the latest shared understanding.

**3) Missed Opportunities for Co-Adaptation.** Both Gato and RT-1 illustrate the single-agent assumption that the robot alone adapts its policy. In scenarios like teleoperation or assistive tasks, however, adaptation is a two-way street: the human also modifies their behavior in response to how the robot is acting, and vice versa. A single-agent viewpoint cannot fully leverage user posture shifts, subtle gestural cues, or real-time user feedback on comfort and safety. By contrast, a multi-agent approach (e.g., ad-hoc teamwork (Rahman et al., 2021; Mirsky et al., 2022)) would explicitly model how the human’s internal state may evolve, whether due to fatigue, changing preferences, or partial completion of sub-goals and modify the robot’s behavior accordingly. This co-adaptive loop can prevent errors (e.g., collisions, user frustration) that arise when the robot rigidly executes a policy absent mutual feedback.

**4) Overlooking Collaborative Goal Setting and Preference Tracking.** Another limitation is that single-agent foundation models rarely incorporate long-term *preference tracking* for an external user. While Gato and RT-1 excel at learning general policies from large datasets, they do not maintain a persistent model of a user’s personal constraints or historical preferences (e.g., “User typically prefers lighter grip force on fragile objects” or “User signals discomfort when the end-effector approaches from the left”). In a multi-agent framework, the robot would treat the user’s preferences as a dynamic factor, continuously updating its internal representation as tasks progress and new user feedback arises, thereby improving safety and user satisfaction in real-world contexts (Li et al., 2021).

## 2.2. Multi-Agent Systems and Ad-Hoc Teamwork

To address limitations of single-agent paradigms, especially in environments that demand complex, dynamic interactions and collaborative problem-solving, a section of research has advanced towards *multi-agent systems* (MAS) and *ad-hoc teamwork*, frameworks designed to facilitate effective collaboration among multiple entities. This section provides an overview of these paradigms, elucidating their capabilities and advantages over single-agent systems in a general context.

### 2.2.1. MULTI-AGENT SYSTEMS

*Multi-agent systems* consist of multiple interacting agents, each possessing individual goals, capabilities, and decision-making processes (Weiss, 1999; Stone & Veloso, 2000). These agents can be homogeneous or heterogeneous, cooperative or competitive, and operate within shared or overlapping environments. The primary distinction between MAS and single-agent systems lies in the ability to manage interdependencies and leverage collective intelligence to solve problems that are intractable for individual agents.

### 2.2.2. AD-HOC TEAMWORK

Building upon the foundation of MAS, *ad-hoc teamwork* (Rahman et al., 2021; Barrett, 2015; Rahman et al., 2021; Melo & Sardinha, 2016) enhances the ability of agents to collaborate effectively without prior coordination, control, or extensive communication. This capability is crucial in environments where agents must form temporary coalitions spontaneously to achieve common objectives, often under conditions of uncertainty and incomplete information (Mirsky et al., 2022).

### 2.2.3. PRINCIPLES OF AD-HOC TEAMWORK

- *Flexibility and Adaptation:* Agents must adapt to varying team compositions and roles, accommodating new members or the departure of existing ones without significant performance degradation (Rahman et al., 2021).
- *Implicit Communication:* Effective ad-hoc teamwork can rely on indirect cues and shared environmental information rather than explicit instructions, enabling seamless collaboration with minimal communication overhead (Barrett, 2015).
- *Shared Goals and Intentions:* Successful ad-hoc teams align their individual objectives with collective goals, ensuring that all agents work towards a common purpose despite originating from different starting points (Barrett, 2015).

#### 2.2.4. ADVANTAGES OF AD-HOC TEAMWORK OVER SINGLE-AGENT SYSTEMS

- *Dynamic Adaptation to Human Partners:* In human-robot interactions, ad-hoc teamwork enables robots to adjust their behavior based on real-time human actions and intentions (Rahman et al., 2021). This adaptability is essential for assistive devices and cyborg systems, where user needs can change rapidly (Mirsky et al., 2022).
- *Enhanced Collaboration in Unstructured Environments:* Robots operating in unpredictable settings (e.g., disaster zones or dynamic factory floors) benefit from the ability to form spontaneous coalitions and coordinate responses to new challenges (Barrett et al., 2017).
- *Improved User Experience:* By enabling more natural and intuitive interactions, often relying on implicit communication, ad-hoc teamwork boosts user satisfaction and trust. Robots are perceived as more reliable and competent partners when they fluidly adapt to human cues.
- *Adaptive Learning and Co-Evolution:* Agents in multi-agent systems learn from and adapt to one another over time. In human-in-the-loop settings, this means the robot and user can co-evolve, preventing the rigid, static policies typical of single-agent models.
- *Resilience to Change:* Decentralized or distributed control in multi-agent systems fosters robustness against environmental shifts or new team compositions. When one agent fails or a new user goal emerges, the team can reconfigure itself and maintain overall performance.
- *Facilitating Human-Robot Synergy:* By explicitly modeling the human’s actions, internal states, and likely next steps, multi-agent systems allow for smoother joint actions, reducing collisions, user discomfort, or misunderstandings about shared goals (Barrett et al., 2017).

By integrating multi-agent and ad-hoc teamwork capabilities, robotic foundation models can achieve a higher level of interaction and collaboration, addressing the fundamental shortcomings of single-agent models. This collaborative framework is essential for developing robots that can engage in real-time, adaptive interactions with humans and other agents, thereby enhancing their functionality and applicability in complex, dynamic environments.

### 3. The Human-Device Dyad and Cyborg Systems

A defining feature of human-interactive robots or cyborg systems (i.e., wearable robotics, including prostheses and exoskeletons) is the *fusion* of human physiology with artificial actuation. In these scenarios, the user (with biological muscles, joints, and neural control) and the robotic device (with actuators, sensors, and algorithms) act as two tightly coupled agents. This bi-directional relationship resonates with sensorimotor loops in neuroscience, where the brain sends motor commands to the musculoskeletal system and receives multisensory feedback to update its internal state (Wolpert & Kawato, 1998).

A single-agent perspective would treat high-level user commands as isolated or sporadic inputs, thereby neglecting the continuous interplay between user and device. This oversight can lead to abrupt control actions, delayed task transitions, or non-personalized responses, especially when the user’s biomechanical or cognitive state changes unexpectedly. Treating the human and the device as two *agents* with partially observable, evolving states opens the door to ad-hoc collaboration paradigms, wherein:

- The device *infers* the user’s intentions and biomechanical constraints from subtle signals like muscle activation patterns.
- The user, in turn, adjusts their motor strategies based on feedback from the device’s behavior.
- The device maintains and updates an internal model of the user dynamically, for more synergistic control.

Even in ostensibly “autonomous” robot applications, there can be hidden interaction partners, such as a human operator providing commands, or an environment whose states change in response to the robot’s actions. Integrating multi-agent interactions into foundation models equips robots with explicit representations of user states, fostering the ability to *predict* and *adapt* continuously based on teammates’ models of the world. This integration addresses the fundamental limitations of single-agent, autonomous controllers in human-interactive settings.

### 3.1. Alternative to Modern AI-Based Approaches: Finite-State Machines (FSMs)

Contrasting with modern AI-based controllers, many existing cyborg-like devices utilize *finite-state machines* (FSMs) to determine high-level behaviors such as *standing*, *walking*, or *ascending stairs* based on simple sensor thresholds (Tucker et al., 2015). While FSMs offer straightforward implementation and higher interpretability, they suffer from inherent limitations:

- *Reactive Design*: Transitions typically occur only after sensor thresholds are crossed, leading to potential lag or misclassification when dealing with rapid or subtle user motion changes.
- *Lack of Predictive Modeling*: FSMs generally do not maintain forward-looking estimates of user dynamics or intentions (e.g., anticipating a user’s transition from walking to running before sensor data fully reflects this change), nor do they store long-term user preferences.
- *Minimal Personalization*: Generic thresholds are usually uniform across all users or only slightly tuned, missing opportunities to develop personalized, long-term models of each user’s comfort and biomechanical idiosyncrasies.

**Shortfall:** These limitations overlook fundamental insights from neuroscience, such as the role of predictive and adaptive control in biological motor systems. For example, the concept of *effference copy* suggests that the central nervous system forwards internal predictions of movement outcomes to anticipate sensory feedback and modulate subsequent actions (Kawato, 1999). Similarly, single-agent foundation models lack deeper modeling of user trajectories and internal states, resulting in less fluid and adaptive interactions compared to multi-agent models that incorporate predictive and adaptive control mechanisms.

### 3.2. Non-Adaptive User Integration in Single-Agent Models

Single-agent models often process user inputs as external, episodic commands without maintaining a structured, continuously updated model of the user. This one-directional handling prompts policy decisions without an evolving *teammate model*, which is particularly problematic in wearable robotics where the user’s physiological and behavioral states are dynamic and critically impact device control. Specifically, single-agent, non-adaptive integration falls short in:

- *Tracking User State Trajectories*: Without an explicit internal model of the user’s evolving comfort level, fatigue, or motion intent, the device cannot proactively adjust its outputs in response to changes in the user’s gait or stance, potentially leading to overshooting or misalignment of forces.
- *Anticipating Transitions*: If the device treats user commands as static triggers rather than dynamic signals, it may lag behind sudden shifts in user intent, resulting in suboptimal or unsafe responses (e.g., delayed torque adjustments when transitioning to an incline for a prosthesis).
- *Capitalizing on Repeated Interactions*: The absence of memory regarding user-specific preferences or long-term progression hinders the ability to personalize interactions and improve user experience over time.

In essence, single-agent systems rarely treat the user as a *co-evolving entity* with its own sensorimotor goals. Wearable prosthetics and exoskeletons demand tightly coupled, bidirectional information flow, aligning more closely with a multi-agent perspective that includes explicit or learned models of human states and intentions (Rahman et al., 2021; Li et al., 2021; Mirsky et al., 2022). This paradigm shift is especially critical in safety-critical or high-comfort applications, where even minor latency or mismatches in human-device coordination can lead to falls, fatigue, or user frustration.

## 4. Position: Embracing an Interactive Multi-Agent Foundational Architecture Inspired by Neuroscience

We position that *future robot foundation models must adopt an interactive multi-agent framework*, especially for human-in-the-loop or semi-autonomous domains, one that recognizes the user and the robot (e.g., a cyborg prosthetic device) as two interacting agents. We suggest that advanced language or multimodal models should serve not simply as “language-to-action” converters as in current robot foundation models, but as *one component* (a sensing module) within a larger, hierarchical architecture. To concretize this vision, we propose a four-module, neuroscience-inspired approach:

- A *sensing module* that integrates language and multimodal inputs (e.g., EMG, camera data) into structured proposals, mirroring how biological systems merge sensory feedback with high-level goals (Pulvermüller, 2005).
- An *ad-hoc teamwork model* (Rahman et al., 2021; Li et al., 2021; Mirsky et al., 2022) that applies multi-agent collaboration principles, aligning with joint-action and shared intentionality theories in cognitive science (Tomasello & Rakoczy, 2003; Sebanz et al., 2006).
- A *predictive world belief model* that maintains an internal model of the user and/or collaborating agents’ states, enabling anticipatory actions, inspired by motor control theories on forward internal models and predictive coding (Wolpert et al., 1995; Wolpert, 1997; Wolpert & Kawato, 1998; Rao & Ballard, 1999; Millidge et al., 2021; Denham & Winkler, 2020).
- A *memory/feedback subsystem* that stores user-specific preferences and updates policies in a reinforcement-like manner, akin to the role of synaptic plasticity and reinforcement learning in shaping long-term sensorimotor adaptations (Dayan & Abbott, 2005)

#### 4.1. Module 1: Sensing via Language and Multimodal Inputs

**Neuroscientific Parallels.** Biological organisms integrate sensory cues from diverse modalities (visual, auditory, proprioceptive) to build coherent percepts and guide behavior (Stein, 1993). Similarly, language in humans is thought to interact with high-level planning networks in the brain, providing semantic context and task-relevant goals (Pulvermüller, 2005).

**Design Concept.** Our *Sensing Module* leverages language models (e.g., LLaMA (Touvron et al., 2023)) and multimodal models (e.g., PaLM-E (Driess et al., 2023)) to parse:

- *User Commands:* Natural language instructions, which may include explicit directives about speed or comfort, are incorporated into semantically meaningful embeddings.
- *Multimodal Inputs:* Visual, auditory, and biomechanical cues (e.g., from onboard cameras, IMUs, or EMG sensors), fused within a multimodal encoder (Driess et al., 2023) or CLIP-like encoders (Radford et al., 2021).

The Sensing Module outputs a high-level *proposal* about how to adjust the robot control parameters (e.g., torque, stiffness). In essence, it acts like a “multisensory integrator” that also factors in the user’s explicit linguistic preferences. This is reminiscent of how the central nervous system merges sensory feedback with high-level goals to plan motor commands (Fuster, 2002).

#### 4.2. Module 2: Ad-Hoc Teamwork Model

**Cognitive Science and Joint Action Parallels.** In cognitive science, joint action explores how individuals coordinate tasks by internally representing each other’s goals and actions (Sebanz et al., 2006). The brain also employs partial models of another person’s internal states in collaborative tasks, often referred to as *shared intentionality* (Bratman, 1992). When humans collaborate, they engage in *shared intentionality*, and mutual understanding and prediction of each other’s intentions, states, and possible actions (Bratman, 1992). Neuroscientific research further reveals that humans utilize a form of theory of mind (Baron-Cohen et al., 1994), allowing one to infer another’s mental states, thereby enabling synchronization in activities such as dancing, carrying a table together, or passing objects. These mechanisms underpin our ability to anticipate partners’ behavior, rapidly adapt to unexpected changes, and maintain coordinated trajectories.

**Design Concept.** Robot foundation models can adopt a similar approach, treating the user’s state (fatigue, intention, comfort preference) as an evolving variable to be *modeled*, not just a passive source of commands. This enables collaboration by predicting user behavior in advance, leading to more coordinated and synergistic interactions. Adapting these ideas, our *Ad-hoc Teamwork Model* views the device and user as two agents collaborating under imperfect information. The cumulative function of this module is:

- *Intent Inference:* The device (e.g., a leg prosthesis) cannot directly “read” the user’s mind but can *infer* the user’s short-term goals and motor patterns from EMG signals, gait kinematics, and language cues.

- *Belief State Maintenance*: Building on the theory of mind (Baron-Cohen et al., 1994), each agent maintains a dynamic belief state about the other. For the prosthesis, this includes anticipating the user’s comfort boundaries (e.g., acceptable torque levels or speed), preferred movement patterns, and likely next moves. By comparing actual outcomes to predictions at each timestep, the prosthesis refines its internal model, capturing latent variables such as fatigue onset or changes in user intent.
- *Refinement of Proposals*: Further, this module refines the high-level proposals from the *sensing module* by combining: 1) *User-centric constraints*, e.g., muscle fatigue or joint stress. 2) *Online collaboration strategies*: e.g., coordinating action output with the user’s shift in stance to achieve stable gait transitions.

In multi-agent reinforcement learning (MARL), agents may learn joint policies by sharing partial states (Matignon et al., 2012). Here, the user’s partial state must be inferred. The device, through repeated interactions, refines its predictive model of the user, aiming for *co-adaptation* rather than a unidirectional control scheme.

#### 4.3. Module 3: Predictive World Belief Model

**Internal Forward Models in Neuroscience.** A substantial body of research in motor neuroscience underscores the role of *internal forward models*, which predict the sensory outcomes of motor commands, adjusting subsequent behavior in anticipation of future states. (Wolpert & Kawato, 1998; Kawato, 1999). Predictive coding theories go further, proposing that the brain’s cortex continuously attempts to minimize prediction errors by updating these models (Friston, 2009).

**Design Concept.** Drawing inspiration from these frameworks, our *Predictive World Belief Model* maintains a probabilistic, forward-looking estimate of both the user’s internal states (muscle fatigue, comfort thresholds, intention shifts) and external conditions (terrain type, slope, potential obstacles). By modeling these dynamic processes over time, this module achieves:

- *Anticipatory control*: The device (e.g., a leg prosthesis) adjusts torque or joint impedance *before* the user’s foot meets a slippery or uneven surface.
- *Context-specific transitions*: It can predict that a user’s gait may shift from walking to running or from level-ground to stair ascent, adjusting control parameters accordingly.

Unlike simple threshold-based controllers or purely reactive controllers, a model-based approach predicts the probability distribution of future states. Bayesian strategies, along with predictive coding frameworks (Rao & Ballard, 1999; Friston, 2009), can be leveraged to continuously update such forward-model estimates.

#### 4.4. Module 4: Memory & Feedback for Refinement

**Long-Term Plasticity and Reinforcement in Biology.** Neuroscience literature emphasizes how synaptic plasticity through feedback-driven processes, including dopamine-mediated reinforcement signals, drive *long-term* changes in behavior (Gerstner & Kistler, 2002; Dayan & Abbott, 2005). For instance, repetitive practice consolidates motor memories that lead to progressive refinement in the brain’s motor representations and improved task efficiency.

**Design Concept.** Our *Memory Module* and *Feedback Mechanism* incorporate analogous processes:

- *Long-term preference storage*: The device (e.g., a leg prosthesis) stores user-specific torque settings, comfort ranges, and frequent command patterns, similar to how repeated exposure to a task solidifies neural pathways in motor learning.
- *Reinforcement-based updates*: The device can query the user (“Is this stiffness comfortable?”) or automatically evaluate signals of discomfort, adjusting internal parameters to optimize an objective function (e.g., minimal metabolic cost, subjective comfort).
- *Semantic memory bank*: Language-based preferences (“I like a softer ankle when walking on grass”) can be retained as textual or embedding-based knowledge, allowing the device to recall and apply them in the future.

Hence, the robotic cyborg evolves from a generic device to a personalized “partner,” continuously shaped by the user’s feedback and changing biomechanical needs.

#### 4.5. Implications for Training and Implementation

**Offline Pre-Training with Diverse Data.** Much like other foundation models, our system benefits from large-scale pre-training:

- *Multimodal corpora:* Recorded EMG, IMU, camera data across diverse tasks (walking, running, stair-climbing), combined with user commands in natural language.
- *Extensive annotation:* Labels for user comfort, stability, and environment factors (terrain type, obstacles) to facilitate robust supervised or self-supervised learning.

Challenges include data diversity (users vary widely in gait patterns, limb morphology, or assistance requirements) and the ethical dimensions of collecting sensitive physiological data.

**In-Situ Fine-Tuning and Personalization.** Aligning with the concept of online plasticity, the system should allow frequent updates based on real-time user feedback to maximize the user’s cumulative comfort or functional metrics (e.g., walking speed, and joint stability). Each device undergoes an online adaptation phase, deploying techniques from continual or reinforcement learning (Rolnick et al., 2019; Dayan & Abbott, 2005). The interacting robotic device thus refines parameters (e.g., stiffness range, torque profiles) while actively interacting with its user in real-world conditions, leveraging the Memory/Feedback mechanism.

**Safety and Interpretability.** Human-interactive robotic devices necessitate robust safety constraints and interpretability.

- *Reflex-level safeguards:* The framework should include a “low-level reflex layer,” analogous to spinal reflexes, (Wolpaw et al., 2002) to prevent unsafe actuator outputs.
- *Explainability tools:* Clinicians and users should be able to understand *why* the device chose specific parameters (e.g., a certain ankle torque), fostering trust and regulatory approval.

### 5. More Scenarios of Multi-Agent Coordination

While we emphasize user-device dyads such as cyborg systems, the same multi-agent logic can apply to developing foundation models for extended scenarios such as:

- *Collaborative Task Environments:* In settings where multiple robotic agents operate alongside human users, such as in industrial manufacturing, disaster response, or healthcare, multi-agent coordination enables seamless collaboration. For instance, in a manufacturing assembly line, multiple robots can dynamically adjust their roles and tasks based on real-time production demands and human worker inputs. Similarly, in disaster response scenarios, a team of drones and ground robots can coordinate their efforts to search for survivors, navigate hazardous terrains, and relay critical information, thereby enhancing the effectiveness and efficiency of the mission (Barrett et al., 2017).
- *Generalist Models for Smart Homes:* In human-interactive systems, such as personal assistants or smart home environments, multi-agent coordination allows for more personalized interactions. Multiple agents can manage different aspects of the environment (e.g., lighting, climate control, security) while collaboratively adapting to the user’s preferences and routines. This can lead to a more cohesive and intuitive user experience, where the system anticipates and responds to the user’s needs.
- *Generalist Synergy Models:* A single user might utilize a lower-limb prosthesis, an upper-limb orthosis, or additional wearable sensors for rehabilitation. Each device can maintain a local instance of the same foundation-model framework, with the potential for *coordinated synergy* if they share a global representation of the user’s state. For instance, an upper-limb exoskeleton might stiffen its elbow joint to assist balance when it detects that the lower-limb prosthesis is transitioning to a challenging slope. Such “whole-body integration” draws on the principle that motor coordination in the human body emerges from distributed neural systems working in tandem (Ivry & Spencer, 2004).



## 6. Safety, Ethical, and Regulatory Perspectives

### 6.1. Medical Accountability and Clinical Evidence

Adopting multi-agent foundation models in robotics, particularly within medical and assistive technologies, necessitates stringent compliance with regulatory standards. The integration of multiple autonomous agents introduces complexities in risk assessment and accountability. It is imperative to conduct comprehensive clinical trials and gather substantial evidence to validate the safety and efficacy of these AI-driven control strategies. Ensuring minimal risks of adverse events, such as falls or device malfunctions, and demonstrating reliable performance across diverse user populations is critical for regulatory approval and widespread adoption (Bender et al., 2021).

### 6.2. Privacy, Data Ownership, and Bias Mitigation

Multi-agent systems often require extensive collection and processing of physiological and contextual data, increasing privacy concerns. Implementing robust data protection measures, such as on-device processing, federated learning, and anonymized data management, is essential to safeguard user information. Additionally, large-scale models can inadvertently perpetuate biases present in their training data (Bender et al., 2021). In healthcare and assistive contexts, such biases could lead to unequal performance across different user groups. Thus, rigorous data curation, bias mitigation strategies, and domain-specific fine-tuning are imperative to ensure equitable and unbiased system performance.

## 7. Discussion for Future Research

### 7.1. Toward Adaptive, Personalized Prostheses at Scale

The proposed methodology heralds a shift from static, command-following prostheses to devices that actively *co-adapt* with their human partners. In practice, widespread adoption relies on:

- *Standardized Benchmarks*: Similar to the large-scale robotics benchmarks in manipulation, new testbeds for human-interactive robots, such as, prosthetic synergy are required. Tasks should capture real-world complexity, e.g., irregular outdoor environments, dynamic changes in user fatigue, and prolonged usage scenarios.
- *Open-Source Ecosystems*: Encouraging open data and model sharing can expedite progress. Much like large-scale image or speech datasets have revolutionized computer vision and NLP, similarly sized corpora of wearable robotics data could accelerate foundation-model development.
- *Clinical Partnerships*: Collaboration with clinicians, physical therapists, and user communities can ensure that the system’s objectives align with real-world functional outcomes (e.g., reduced risk of falls, improved metabolic efficiency, subjective comfort).

### 7.2. Active Learning and Human-in-the-Loop Refinement

Beyond passively receiving user feedback, future prosthetics could proactively query the user to reduce uncertainty: “Would you like increased ankle torque for this incline?” Such *active learning* parallels the process by which humans ask clarifying questions during joint action to refine shared tasks (Sebanz et al., 2006). In the context of foundation models, these interactions deepen the system’s semantic memory and enhance personalization.

## 8. Conclusion

We argue that future robot foundation models must be designed from a *multi-agent* perspective to meet the demands of interaction and non-stationarity. Rather than functioning as isolated, single-agent systems, these models should explicitly account for both the robot and its human counterpart (or broader environment) as actively adapting agents. By weaving together insights from neuroscience, cognitive science, and multi-agent systems, next-generation interactive robotics, exemplified by cyborgs and wearable devices, can move beyond one-shot instructions and rigid autonomy. Embracing this multi-agent framework has the potential to ultimately deliver interactive, more robust, and user-centric performance, offering capabilities that surpass the limitations of today’s single-agent paradigms.

## References

- Abiri, R., Borhani, S., Sellers, E. W., Jiang, Y., and Zhao, X. A comprehensive review of eeg-based brain–computer interface paradigms. *Journal of neural engineering*, 16(1):011001, 2019.
- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Baron-Cohen, S., Ring, H., Moriarty, J., Schmitz, B., Costa, D., and Ell, P. Recognition of mental state terms. *British Journal of Psychiatry*, 165(5):640–649, 1994.
- Barrett, S. *Making friends on the fly: advances in ad hoc teamwork*, volume 603. Springer, 2015.
- Barrett, S., Rosenfeld, A., Kraus, S., and Stone, P. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence*, 242:132–171, 2017.
- Bender, E. M., Gebru, T., McMillan-Major, A., and Shmitchell, S. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pp. 610–623, 2021.
- Best, T. K., Welker, C. G., Rouse, E. J., and Gregg, R. D. Data-driven variable impedance control of a powered knee–ankle prosthesis for adaptive speed and incline walking. *IEEE Transactions on Robotics*, 39(3):2151–2169, 2023.
- Bratman, M. E. Shared cooperative activity. *The philosophical review*, 101(2):327–341, 1992.
- Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- Brown, T., Mann, B., Ryder, N., et al. Language models are few-shot learners. *NeurIPS*, 2020.
- Cimolatto, A., Driessen, J. J., Mattos, L. S., De Momi, E., Laffranchi, M., and De Michieli, L. Emg-driven control in lower limb prostheses: A topic-based systematic review. *Journal of NeuroEngineering and Rehabilitation*, 19(1):43, 2022.
- Clark, H. and Feng, J. Semi-autonomous vehicles: Examining driver performance during the take-over. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, volume 59, pp. 781–785. SAGE Publications Sage CA: Los Angeles, CA, 2015.
- Dayan, P. and Abbott, L. F. *Theoretical neuroscience: computational and mathematical modeling of neural systems*. MIT press, 2005.
- Denham, S. L. and Winkler, I. Predictive coding in auditory perception: challenges and unresolved questions. *European Journal of Neuroscience*, 51(5):1151–1160, 2020.
- Dey, S. and Schilling, A. F. Data-driven gait-predictive model for anticipatory prosthesis control. In *2022 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–6. IEEE, 2022.
- Dey, S., Yoshida, T., and Schilling, A. F. Feasibility of training a random forest model with incomplete user-specific data for devising a control strategy for active biomimetic ankle. *Frontiers in Bioengineering and Biotechnology*, 8:855, 2020.
- Dey, S., De Schultz, N., and Schilling, A. F. Why hard code the bionic limbs when they can learn from humans? In *2023 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–6. IEEE, 2023.
- Donoghue, J. P. Bridging the brain to the world: a perspective on neural interface systems. *Neuron*, 60(3):511–521, 2008.
- Driess, D., Xia, F., Sajjadi, M. S., Lynch, C., Chowdhery, A., Ichter, B., Wahid, A., Tompson, J., Vuong, Q., Yu, T., et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- Friston, K. The free-energy principle: a rough guide to the brain? *Trends in cognitive sciences*, 13(7):293–301, 2009.
- Fuster, J. M. Frontal lobe and cognitive development. *Journal of neurocytology*, 31(3):373–385, 2002.

- Gerstner, W. and Kistler, W. M. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.
- Ivry, R. B. and Spencer, R. M. The neural representation of time. *Current opinion in neurobiology*, 14(2):225–232, 2004.
- Jackson, A. and Zimmermann, J. B. Neural interfaces for the brain and spinal cord—restoring motor function. *Nature Reviews Neurology*, 8(12):690–699, 2012.
- Kawato, M. Internal models for motor control and trajectory planning. *Current opinion in neurobiology*, 9(6):718–727, 1999.
- Kofman, J., Wu, X., Luu, T. J., and Verma, S. Teleoperation of a robot manipulator using a vision-based human-robot interface. *IEEE transactions on industrial electronics*, 52(5):1206–1219, 2005.
- Li, H., Ni, T., Agrawal, S., Jia, F., Raja, S., Gui, Y., Hughes, D., Lewis, M., and Sycara, K. Individualized mutual adaptation in human-agent teams. *IEEE Transactions on Human-Machine Systems*, 51(6):706–714, 2021.
- Li, J., Li, D., Xiong, C., and Hoi, S. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pp. 12888–12900. PMLR, 2022.
- Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023.
- Lichardopol, S. A survey on teleoperation. 2007.
- Lo, H. S. and Xie, S. Q. Exoskeleton robots for upper-limb rehabilitation: State of the art and future prospects. *Medical engineering & physics*, 34(3):261–268, 2012.
- Matignon, L., Laurent, G. J., and Le Fort-Piat, N. Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 27(1):1–31, 2012.
- McFarland, D. J. and Wolpaw, J. R. Eeg-based brain–computer interfaces. *current opinion in Biomedical Engineering*, 4: 194–200, 2017.
- Melo, F. S. and Sardinha, A. Ad hoc teamwork by learning teammates’ task. *Autonomous Agents and Multi-Agent Systems*, 30:175–219, 2016.
- Millidge, B., Seth, A., and Buckley, C. L. Predictive coding: a theoretical and experimental review. *arXiv preprint arXiv:2107.12979*, 2021.
- Mirsky, R., Carlucho, I., Rahman, A., Fosong, E., Macke, W., Sridharan, M., Stone, P., and Albrecht, S. V. A survey of ad hoc teamwork research. In *European conference on multi-agent systems*, pp. 275–293. Springer, 2022.
- Molteni, F., Gasperini, G., Cannaviello, G., and Guanziroli, E. Exoskeleton and end-effector robots for upper and lower limbs rehabilitation: narrative review. *PM&R*, 10(9):S174–S188, 2018.
- Nicolas-Alonso, L. F. and Gomez-Gil, J. Brain computer interfaces, a review. *sensors*, 12(2):1211–1279, 2012.
- Okamura, A. M. Methods for haptic feedback in teleoperated robot-assisted surgery. *Industrial Robot: An International Journal*, 31(6):499–508, 2004.
- Pulvermüller, F. Brain mechanisms linking language and action. *Nature reviews neuroscience*, 6(7):576–582, 2005.
- Quintero, D., Martin, A. E., and Gregg, R. D. Toward unified control of a powered prosthetic leg: A simulation study. *IEEE Transactions on Control Systems Technology*, 26(1):305–312, 2017.
- Radford, A. et al. Clip: Connecting text and images. In *ICML*, 2021.
- Rahman, M. A., Hopner, N., Christianos, F., and Albrecht, S. V. Towards open ad hoc teamwork using graph-based policy learning. In *International conference on machine learning*, pp. 8776–8786. PMLR, 2021.

- Rao, R. P. and Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature neuroscience*, 2(1):79–87, 1999.
- Reed, S., Zolna, K., Parisotto, E., Colmenarejo, S. G., Novikov, A., Barth-Maron, G., Gimenez, M., Sulsky, Y., Kay, J., Springenberg, J. T., et al. A generalist agent. *arXiv preprint arXiv:2205.06175*, 2022.
- Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., and Wayne, G. Experience replay for continual learning. *Advances in neural information processing systems*, 32, 2019.
- Rosen, J. and Perry, J. C. Upper limb powered exoskeleton. *International Journal of Humanoid Robotics*, 4(03):529–548, 2007.
- Schultz, A. E. and Kuiken, T. A. Neural interfaces for control of upper limb prostheses: the state of the art and future possibilities. *PM&R*, 3(1):55–67, 2011.
- Sebanz, N., Bekkering, H., and Knoblich, G. Joint action: bodies and minds moving together. *Trends in cognitive sciences*, 10(2):70–76, 2006.
- Stein, B. *The Merging of the Senses*. MIT Press, 1993.
- Stone, P. and Veloso, M. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8: 345–383, 2000.
- Suchan, J. and Osterloh, J.-P. Assessing drivers’ situation awareness in semi-autonomous vehicles: Asp based characterisations of driving dynamics for modelling scene interpretation and projection. *arXiv preprint arXiv:2308.15895*, 2023.
- Team, O. M., Ghosh, D., Walke, H., Pertsch, K., Black, K., Mees, O., Dasari, S., Hejna, J., Kreiman, T., Xu, C., et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- Tomasello, M. and Rakoczy, H. What makes human cognition unique? from individual to shared to collective intentionality. *Mind & language*, 18(2):121–147, 2003.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- Tucker, M. R., Olivier, J., Pagel, A., Bleuler, H., Bouri, M., Lambercy, O., Millán, J. d. R., Riener, R., Vallery, H., and Gassert, R. Control strategies for active lower extremity prosthetics and orthotics: a review. *Journal of neuroengineering and rehabilitation*, 12:1–30, 2015.
- Vogel, J., Hagengruber, A., Iskandar, M., Quere, G., Leipscher, U., Bustamante, S., Dietrich, A., Höppner, H., Leidner, D., and Albu-Schäffer, A. Edan: An emg-controlled daily assistant to help people with physical disabilities. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4183–4190. IEEE, 2020.
- Weiss, G. *Multiagent systems: a modern approach to distributed artificial intelligence*. MIT press, 1999.
- Wen, Y., Si, J., Brandt, A., Gao, X., and Huang, H. H. Online reinforcement learning control for the personalization of a robotic knee prosthesis. *IEEE transactions on cybernetics*, 50(6):2346–2356, 2019.
- Wolpaw, J. R. Brain–computer interfaces. In *Handbook of clinical neurology*, volume 110, pp. 67–74. Elsevier, 2013.
- Wolpaw, J. R., Birbaumer, N., McFarland, D. J., Pfurtscheller, G., and Vaughan, T. M. Brain–computer interfaces for communication and control. *Clinical neurophysiology*, 113(6):767–791, 2002.
- Wolpert, D. M. Computational approaches to motor control. *Trends in cognitive sciences*, 1(6):209–216, 1997.
- Wolpert, D. M. and Kawato, M. Multiple paired forward and inverse models for motor control. *Neural networks*, 11(7-8): 1317–1329, 1998.
- Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. An internal model for sensorimotor integration. *Science*, 269(5232): 1880–1882, 1995.