
Human-Robotic Prosthesis as Collaborating Agents for Symmetrical Walking

Ruofan Wu *
Arizona State University

Junmin Zhong *
Arizona State University

Brent Abraham Wallace
Arizona State University

Xiang Gao
Arizona State University

He Huang
North Carolina State University

Jennie Si †
Arizona State University

Abstract

This is the first attempt at considering human influence in the reinforcement learning control of a robotic lower limb prosthesis toward symmetrical walking in real world situations. We propose a collaborative multi-agent reinforcement learning (cMARL) solution framework for this highly complex and challenging human-prosthesis collaboration (HPC) problem. The design of an automatic controller of the robot within the HPC context is based on accessible physical features or measurements that are known to affect walking performance. Comparisons are made with the current state-of-the-art robot control designs, which are single-agent based, as well as existing MARL solution approaches tailored to the problem, including multi-agent deep deterministic policy gradient (MADDPG) and counterfactual multi-agent policy gradient (COMA). Results show that, when compared to these approaches, treating the human and robot as coupled agents and using an estimated human adaption in robot control design can achieve lower stage cost, peak error, and improved symmetry to ensure better human walking performance. Additionally, our approach accelerates learning of walking tasks and increases learning success rate. The proposed framework can potentially be further developed to examine how human and robotic lower limb prosthesis interact, an area that little is known about. Advancing cMARL toward real world applications such as HPC for normative walking sets a good example of how AI can positively impact on people’s lives.

1 Introduction

The concept, design, and applications of human-robot cooperation have advanced rapidly due to new demands in AI-enabled applications fueled by powerful deep learning and reinforcement learning algorithms [1, 2, 3]. Human-robot collaboration can take on a variety of forms depending on tasks to be solved, how information is shared [1, 2], and the nature of interaction [3]. Examples of complex collaborative tasks may include picking up or carrying objects together [4, 5], cooperating on a production line, in which cases a robot can learn to imitate human demonstrations [6, 7, 8]. Other application scenarios may include intermittent robotic correction of human driving, or vice versa [9]. In essence, most of these recent studies involve interactions between a human and a robot in a way such that there is either space between the agents, or there is time for predictive counter measures to interfere. By contrast, in the HPC problem the human and robot agents are physically coupled together, and there is often little time for the human agent to react to prevent from falling or injury.

*Equal Contribution.

†Corresponding author: si@asu.edu

As a result, innovative solution frameworks are needed to first address the design, analysis, and testing of complex control systems to restore locomotion for amputees. Based on these platforms, human-robot interactions can be studied, and fundamental issues such as user preference and prosthesis embodiment can be examined. The knowledge in turn, will further help design seamless automatic robot control systems. As such, challenges arising from controlling a wearable robotic lower limb prosthesis to meet the human user's needs are unique, and in some aspects, these challenges are greater than what have been studied [10, 11, 12, 13].

State-of-the-art automatic control of robotic prosthesis has been reported recently with successful human subject testing [14, 15, 16, 17]. These approaches are single agent-based reinforcement learning controls, specific tasks such as level ground and ramp walking were accomplished via designer prescribed or specified robotic joint movement profiles [14, 15]. A most recent progress in single-agent based reinforcement learning solution demonstrated that human amputee subject can perform level ground and ramp walking with the robot controller aimed at mimicking the intact joint movement [17]. Note, however, that all previous results have not directly considered human influence on the human-robot system performance in tuning the robotic prosthesis controller, a phenomena that has direct impact not only on restoring walking but also on human health [18].

Achieving normative walking is fundamentally a real time control problem that involves continuous states and continuous controls of the human-robot system. Continuous state and control problem has received great attention. Approaches based on (deep) reinforcement learning have shown promise to substantively address real world applications. Several algorithms, such as Deep Deterministic Policy Gradient (DDPG) [19], Proximal Policy Optimization (PPO) [20], Soft Actor-Critic (SAC) [21], and Twin Delayed DDPG (TD3) [22], have demonstrated success with solving complex control problems. For example, in simulated human locomotion control [23, 24], deep reinforcement learning solved over 20 independent control signals to facilitate a humanoid robot to achieve different walking tasks. Additional single agent-based continuous control has also demonstrated promise in engineering applications such as stabilization, tracking, and reconfiguring control of Apache helicopters [25, 26, 27], stabilization and control of large power grids [28, 29, 30], robotic manipulation and locomotion via MuJoCo and OpenAI Gym [31, 32], and wearable robots with human in the loop [14, 15, 16, 17].

While these works are encouraging toward solving realistic single-agent continuous control problems, it is not obvious how they can directly address the multi-agent human-robot normative walking problem as needed when we consider human influence in the robot control design. From a physical human-robot interaction (pHRI) perspective, robotic upper-limb control has undergone intense development, especially in the realms of patient rehabilitation [33] and industrial applications [34]. However, this type of pHRI problems differ fundamentally from the human-robot walking problem, as their control target usually consists of well-defined end points generated by a decoupled trajectory generation exosystem [35]. On the other hand, human-robot walking tasks are difficult to associate with an end point task goal due to tight dynamic coupling between the human and the robotic lower limb, and many factors can affect the human's performance goal. As such, even though multi-agent reinforcement learning (MARL) control is a natural candidate to address our HPC challenge, a feasible solution is yet to be developed.

2 Related Work and Challenges

Single agent RL for automatic control of a robotic limb. **Simulation.** Most state-of-the-art RL control design approaches to enable continuous human-robot walking are single-agent based. Important milestones have been achieved by two major classes of RL algorithms: actor-critic algorithms including direct heuristic dynamic programming (dHDP) [36, 37, 38, 39], and variants of policy iteration algorithms such as flexible policy iteration (FPI) [40]. Both types of single-agent based control algorithms were developed and demonstrated in simulated environments first. **Human Tests.** Then these algorithms were tested on human subjects walking with a robotic knee prosthesis [15, 16, 41]. Note that all of the above methodologies and tests used designer-prescribed robot joint movement profiles generated *a priori* for the specific subjects and walking tasks. In real-world use scenarios, joint movement profiles evolve dynamically in real-time to accommodate internal human walking objectives [14, 42]. A more recent single-agent RL control work [17] replaced designer-prescribed joint movement profiles with the intact joint motion. Note also that, none of the above results have directly taken into account human influence on human-robot walking in the control design, a novel contribution of this work. **Multi-agent reinforcement learning (MARL).**

A general cMARL problem usually cannot be transformed to an equivalent single-agent problem, as additional agent(s) and their control policies introduce uncertainties and/or non-stationarity to the environment [43]. Partial observability is common in cMARL problems, which are usually addressed by the decentralized POMDP framework. Centralized training decentralized execution (CTDE) is a popular paradigm to address multi-agent coordination problems. A central feature of most CTDE approaches is factorization of the joint state-action value function into individual utility functions. **MADDPG** [44] is a popular CTDE-based method to solve both cooperative and competitive MARL problems. **COMA** [45] is another popular MARL algorithm which utilizes a single centralized critic by using global state information and actions of all agents. We therefore use them in benchmark evaluations. **Shared Autonomy (SA)**. The human-prosthesis problem under our consideration falls into the area of physical human-robot interaction (pHRI) as the human and the robotic prosthesis are physically coupled at all times. It is not, however, the extensively studied pHRI problem archetype (e.g., cooperative object manipulation, human operating in a remote environment), wherein interactions are usually mediated by a third object. For similar reasons, our pHRI problem is different in several important aspects from the approaches to the existing shared autonomy, such as cross-training [46], bounded memory adaptation [47], predict and blend [48], and model-free RL [49]. The wearable exoskeleton control problem seems relevant to our HPC problem. Yet, there are still fundamental differences. Recent exoskeleton approach also takes into account human-exoskeleton interacting effects [50, 51]. However, these exoskeletons mainly focus on end point performance of a foot or lower limb joints where user intent is quantified by estimating human joint torque or interactive torque. Thus, human-exoskeleton system performance goal is for the robot to produce a well-defined joint or endpoint trajectory [52, 53, 54]. For HPC problem we consider, there is no clear end point goal as there is in most existing current pHRI problems, including robotic upper limb that has been intensively addressed in literature. Therefore, how to define a performance goal for our HPC problem is a challenge. We address this study from existing limited yet proved knowledge [14, 42, 55]. **Modeling Challenges**. The HPC walking problem involves two strongly coupled agents, the interacting dynamics of which are difficult or nearly impossible to describe by ordinary difference or differential equations in large part because there is no clear performance goal in HPC as in studied cases of shared autonomy. HPC may be affected by several factors, such as lower limb mechanics, inter-limb neuromechanical coupling, and physical structure of the human body including the lower limb [56, 57]. Physical and physiological differences in individual human subjects further complicates modeling and control design. Even though our knowledge of human motor control and motor learning in cases such as post-stroke or general loss of normative locomotion capabilities have expanded greatly [58], little is known about how an active robotic prosthesis, not a traditional stick-type passive prosthesis, affects human and vice versa in walking tasks. This is because that the human-robotic prosthesis interacts continuously, the intricate human neurocontrol circuits including sensing, perception and feedback control, which are necessary to facilitate normative walking, is disrupted after amputation. **Human Utility Challenges**. Because of the reasons above, also because of no clear end point target as performance goals as in most studied shared autonomy cases, seamless collaboration between a robotic lower limb prosthesis serially attached to its human user poses new issues that requires to be answered [10, 11]. In this study, we based on latest understanding on the HPC problem, to provide an innovative solution approach to account for human influence in HPC.

Contributions. Human-robot collaborative tasks will continue to play an increasingly vital role in modern life, and existing single- and multi-agent control frameworks have only covered a small subset of the problems. The contributions of this work are as follows. 1) We introduce an innovative approach to automatically control a robotic lower limb prosthesis by treating the human user as a collaborating agent. We thus address a new challenge in the domain of shared autonomy problems. 2) To solve the control problem, we introduce a new cMARL approach to solve our HPC problem, a problem that cannot be readily solved to satisfaction by existing MARL approaches such as COMA or MADDPG. 3) This is the very first attempt in the field of wearable lower limb robots that human influence is explicitly considered in the robot control design.

3 Fundamental control problem statement

In powered lower limb prosthesis, the finite state impedance control (FS-IC) framework most frequently serves to provide intrinsic control of a robotic joint, i.e., it is the built-in controller from robotic prosthesis manufacturers. While position control is common in industrial robots, when a human is affixed to a robotic joint, robot trajectory tracking using position control may preclude any

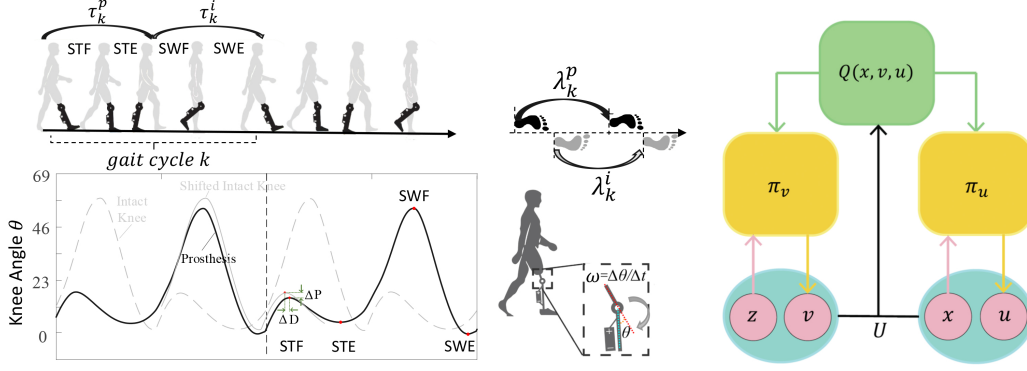


Figure 1: Left: human gait and prosthesis FS-IC characteristics. Right: cMARL solution approach to the HPC problem.

dynamic interaction of the robot with its human user and the environment [59]. This may cause an amputee to react to the awkwardness of the prosthesis rather than to interact constructively with it [60]. As FS-IC impedance control framework is considered to provide compliant control for human use of the robot, it is expected to generate stable and predictable human-robot walking behavior, and thus, it is adopted in this study which aims for its controller designs at real world applications.

The FS-IC treats a gait cycle, or a step, as four consecutive gait phases (Figure 1, Left): stance flexion (STF) as Phase 1, stance extension (STE) as Phase 2, swing flexion (SWF) as Phase 3, and swing extension (SWE) as Phase 4. In what follows, we design four individual controllers for each of the four phases. Since the four controllers are conceptualized similarly and designed using the same approach, for the sake of clarity we carry out subsequent discussion without explicitly referring to the specific phase numbers (but we emphasize that each of the four phases requires its own independently-trained controller). Additional details on FS-IC such as phase detection and transition are provided in Appendix A.3,

At the k -th gait cycle, a robot controller (solved by cMARL in this paper) is to determine three impedance parameters, $I_k = [K_k, B_k, (\theta_e)_k]^T \in \mathbb{R}^3$ under the FS-IC framework, representing stiffness K_k , damping coefficient B_k , and equilibrium position $(\theta_e)_k$, respectively. The prosthetic joint motor torque $T_k \in \mathbb{R}$ is then generated based on joint kinematics (knee joint angle θ and angular velocity ω) according to the following impedance control law,

$$T_k = K_k(\theta - (\theta_{e_k})) + B_k\omega. \quad (1)$$

The control problem formulation requires automatically determining 12 control inputs or impedance control parameters (3 for each of the 4 gait phases) for individual users aiming at walk in real world situations. The initial set of feasible baseline impedance control parameters I_0 can be obtained from manufacturers or rehabilitation clinics. Impedance updates take place according to

$$I_{k+1} = I_k + u_k, \quad (2)$$

where $u_k \in \mathbb{R}^3$ is to be determined from our proposed cMARL approach.

4 Method

The human user and the robotic limb are considered collaborating agents. Our cMARL solution approach toward human-robot symmetrical walking problem is formulated based on physical features and measurements that have been shown affecting human-robot walking performance, and they are respectively available to the human and the robot, but not necessarily to each other.

State and control variables. Refer to Figure 1. At the k -th gait cycle, let τ_k^i and λ_k^i , respectively, represent the stance time (time of foot on the ground) and step length (length between two consecutive steps when toe touching the ground) of the human intact leg. Similarly, τ_k^p and λ_k^p , respectively, for the prosthetic leg. Let $\Delta\tau_k$, $\Delta\lambda_k$ denote respectively the difference of stance time and step length,

$$\Delta\tau_k = \tau_k^p - \tau_k^i, \quad \Delta\lambda_k = \lambda_k^p - \lambda_k^i. \quad (3)$$

We thus define human state variable $z_k \in \mathbb{R}^2$ as $z_k = [\Delta\tau_k, \Delta\lambda_k]^T$.

Robot states include kinematic features determined from the knee movement profiles as well as step length and stance time (or z_k). For robot kinematic state variables, we extract gait kinematic features from knee motion profiles of both limbs (Figure 1). Let P_k^i and D_k^i represent the knee angle and time duration of the intact limb, and similarly, P_k^p and D_k^p for the prosthetic limb. Let ΔP_k and ΔD_k denote the error of peak knee angles and the error of duration time between the intact and the robotic knees, respectively,

$$\Delta P_k = P_k^p - P_k^i, \Delta D_k = D_k^p - D_k^i. \quad (4)$$

We then define robot state $x_k \in \mathbb{R}^4$ to include the following four variables, the first two of which are shared between the human and the robot, $x_k = [\Delta\tau_k, \Delta\lambda_k, \Delta P_k, \Delta D_k]^T$.

As such, human states are available to the robot, but a subset of the robot state (the kinematic features) is not available to the human. This reasonably reflects practical situations in real life applications.

The robot control policy is a state feedback law from robot state to robot impedance control parameters, namely, control u_k consists of increments to the impedance control parameters (Eq. 2),

$$u_k = [\Delta K_k, \Delta B_k, (\Delta\theta_e)_k]^T. \quad (5)$$

Determining human control policy which originates from human neurocontrol circuit is a daunting task, especially now that sensing, perception and feedback control circuits are interrupted after a lower limb amputation. While agents' policy solutions are solved from Bellman optimality-like equations, a human cannot interpret and implement an MARL solution via their complicated neural circuits. But this is not our intention. We instead use the human control variable as an estimated input to the robot controller design, i.e., the robot is informed by an estimated human influence when they share the same goal of symmetrical walking.

As the very first step to demonstrate this idea, we consider the human user of the wearable robot intentionally or voluntarily walk at a reference step length, which may also be from an instructional feedback. For real life scenarios, such behavioral cues can change over time, task, or environment. But it is important to validate such formulation of considering human influence in the robot control design. Toward this end, we consider a practical reference cue denoted as a desired step length λ_o . We let v_k represent an endogenous control signal of the human which is a function of human physical and mental states that involve activities ranging from neural level to muscular and joint level. Accordingly, we represent the endogenous human influence solved from the MARL design as a step length λ_k^d . We therefore define human control as,

$$v_k = \lambda_k^d - \lambda_o. \quad (6)$$

Including this estimated human control into the problem of robot control design, we take into account human influence on human-robot walking performance while they share the same symmetrical walking goal.

Symmetrical walking as shared task goal.

We consider the stage cost to be shared between the human and the robot.

$$U(x_k, v_k, u_k) = x_k^T R_x x_k + R_v v_k^2 + u_k^T R_u u_k + \mu h_k^2, \quad (7)$$

where $R_x \in \mathbb{R}^{4 \times 4}$ and $R_u \in \mathbb{R}^{3 \times 3}$ are positive semi-definite weighting matrices, R_v and μ are positive weighting constants, and $h_k = v_k - \bar{v}_k$. In the above, the difference between actual human step length λ_k^i and reference λ_o , denoted as $\bar{v}_k = \lambda_k^i - \lambda_o$, is considered practical and available.

In this formulation, the shared control objective is represented in two ways. First, the robot kinematic state variables in x_k are to match the intact knee, i.e., the robotic joint angle profile is to match that of the intact joint. Additionally, a walking symmetry measure is directly considered by the differences in step length and stance time between human and robot, commonly used gait symmetry measures [61, 62, 63]. The human is assumed to perceive the gait symmetry but do not have access to or understand robot kinematic data (peak knee deflection error ΔP_k and peak knee time error ΔD_k). The works [14, 64] show from test subject data that human control and adaptation has a direct influence on robot kinematics through dynamic interactions [59]. To account for this dynamic learning phenomenon, this framework penalizes the estimation error $h_k = v_k - \bar{v}_k = \lambda_k^d - \lambda_k^i$

between human perceived difference in step length. An unexpected or undesired human action will increase the cost represented by h_k , while a diminishing h_k indicates that human input that originated from the neuromuscular system matches actual human performance. Minimal controller energy expenditures are also included in the cost structure.

The human-robot control objective to be solve by MARL as a function of the state variables and the the controls is formulated as an infinite horizon, discounted cost $Q(x_k, v_k, u_k) = \sum_{j=k}^{\infty} \gamma^{j-k} U(x_j, v_j, u_j)$ where γ ($0 < \gamma < 1$) is the discount factor for the infinite-horizon problem.

Solutions to control actions. Let $v_k = \pi_v(z_k)$ and $u_k = \pi_u(x_k)$, where π_v and π_u are the human and the robot control policies solved from cMARL, respectively. Solving optimal control policies π_v^* and π_u^* requires solving the optimal Q function that satisfies Bellman optimality equation:

$$Q^*(x_k, v_k, u_k) = U(x_k, v_k, u_k) + \gamma Q^*(x_{k+1}, \pi_v^*(z_{k+1}), \pi_u^*(x_{k+1})). \quad (8)$$

When using a neural network-based actor-critic solution framework, we approximately solve the optimal control problem using the following iterative procedure (i is the iteration index),

$$Q_{i+1}(x_k, v_k, u_k) = U(x_k, v_k, u_k) + \gamma Q_i(x_{k+1}, \pi_{v_i}(z_{k+1}), \pi_{u_i}(x_{k+1})). \quad (9)$$

where $\pi_{v_i}(z_k) = \arg \min_{v_k} Q_i(x_k, v_k, u_k)$, and $\pi_{u_i}(x_k) = \arg \min_{u_k} Q_i(x_k, v_k, u_k)$, and $Q_i(x_k, v_k, u_k)$ are iterative actor policies and iterative Q value function.

During training, the actor and critic back-propagate their respective squared error to update their weights. The prediction error of actor $e_{a_v,k}, e_{a_u,k} \in \mathbb{R}$ is,

$$e_{a_v,k} = e_{a_u,k} = \frac{1}{2} (Q_i(x_k, v_k, u_k))^2. \quad (10)$$

The prediction error for the critic $e_{c,k}$ is formulated based on the Bellman error,

$$e_{c,k} = U + \gamma Q_i(x_{k+1}, \pi_{v_i}(z_{k+1}), \pi_{u_i}(x_{k+1})) - Q_{i+1}(x_k, v_k, u_k), \quad (11)$$

and the critic neural network is trained to minimize $e_{c,k} = \frac{1}{2} e_{c,k}^2$.

The optimal state-action cost-to-go function $Q^*(x_k, v_k, u_k)$ is approximated by a critic neural network which learns the Q function by minimizing the Bellman error on the shared cost signal, not on a local cost signal for either the human or the robot as the human and the robot are physically coupled. We use our established direct heuristic dynamic programming (dHDP) algorithm [36] to solve this approximation dynamic programming problem. The critic neural network is a three-layer MLP with 6 hidden units and uses linear activation function in the output layer. Therefore, we have the approximated cost to go value represented by:

$$\hat{Q}_i(x_k, u_k) = W_{c2,i} \varphi \left(W_{c1,i} [x_k^T, u_k^T]^T \right), \quad (12)$$

where $W_{c1,i} \in \mathbb{R}^{6 \times 8}$ denotes the weight matrix between the input layer and the hidden layer, and $W_{c2,i} \in \mathbb{R}^{1 \times 6}$ the weight matrix between the hidden layer and the output layer during the i th learning update. The weight updates of the hidden layer matrix W_{c2} are according to

$$\Delta W_{c2,i} = l_c \left[-\frac{\partial e_{c,k}}{\partial W_{c2}} \right], \quad (13)$$

and the weight updates of the input layer matrix W_{c1} are according to

$$\Delta W_{c1,i} = l_c \left[-\frac{\partial e_{c,k}}{\partial W_{c1}} \right], \quad (14)$$

where $l_c > 0$ is the learning rate of the critic network.

Similar to the critic network, the actor networks for the human and the robot, respectively are three-layer MLP with 6 hidden units with hyperbolic activation function in the output layer to bound the action output. The same SGD optimizer [36] can be applied to the actor networks as well.

The two actors, u and v , respectively are

$$\begin{aligned} u_k &= \varphi(W_{a_u2,i} * \varphi(W_{a_u1,i} x_k)), \\ v_k &= \varphi(W_{a_v2,i} * \varphi(W_{a_v1,i} z_k)), \end{aligned} \quad (15)$$

where $W_{a_u,1} \in \mathbb{R}^{6 \times 4}$, $W_{a_v,1} \in \mathbb{R}^{6 \times 2}$, $W_{a_u,2} \in \mathbb{R}^{3 \times 6}$ and $W_{a_v,2} \in \mathbb{R}^{1 \times 6}$ are the weight matrices, and $\varphi(\cdot)$ is the hyperbolic tangent activation function used in the hidden layer and the output layer. The weight updates of the hidden layer matrix $W_{a_{2,i}}$ are according to

$$\begin{aligned}\Delta W_{a_u,2,i} &= l_a \left[-\frac{\partial e_{a_u,k}}{\partial W_{a_u,2,i}} \right], \\ \Delta W_{a_v,2,i} &= l_a \left[-\frac{\partial e_{a_v,k}}{\partial W_{a_v,2,i}} \right].\end{aligned}\tag{16}$$

The weight updates of the input layer matrix $W_{a_{1,i}}$ are according to

$$\begin{aligned}\Delta W_{a_u,1,i} &= l_a \left[-\frac{\partial e_{a_u,k}}{\partial W_{a_u,1,i}} \right], \\ \Delta W_{a_v,1,i} &= l_a \left[-\frac{\partial e_{a_v,k}}{\partial W_{a_v,1,i}} \right],\end{aligned}\tag{17}$$

where $l_a > 0$ is the learning rate of the actor.

Further details on the cMARL automatic control solution are provided in Appendix D.

5 Experiment

We conduct design evaluations of the proposed cMARL for symmetrical walking using OpenSim, a well-established open source biomechanical modeling tool for conducting biomechanics research and motor control science [65]. The robot knee control is realized within an FS-IC framework. In real life, initial impedance parameters can be selected based on manufacturer and/or rehabilitation clinician’s recommendations. Similar care is given to OpenSim simulated walking yet in all evaluations, the initial impedance values vary from a large range of settings for fair examination. As in [37], we enforce realistic safety constraints to prevent the human from stumbling or falling. Impedance parameters are reset to initial impedance values if any of the state variables exceed the safety bounds. The safety protocols followed in this work are fully described in Appendix A.4.

To make the simulations reflective of real world conditions, sensor and actuator noise data extracted from real human experimental testing sessions is applied to all the simulations in this study. Appendix A.2 provides the complete procedure followed of extracting noise data from experiments involving human subjects and injecting it into all the simulations.

In this section, we provide results of a large set of simulation studies aiming at answering the following questions: 1) Does our cMARL solution framework provide better performance than state-of-the-art baselines, including MADDPG and COMA? 2) Does including human influence in robot control design accelerate learning and improve success rate of policy in comparison to single-agent based approach (wout/human)? 3) Is our cMARL applicable to different and realistic walking tasks? To provide answers, we show three sets of evaluations: benchmark, ablation and reliability. Benchmark and ablation evaluations are based on a level ground walking task with a pace of 1m/s. Reliability evaluations are based two new walking tasks: slope walking on a 11.5 degree ramp and level ground walking at an increased pace of 1.12m/s.

Performance Criteria. In order to ensure that amputee subjects walk safely and continuously, we consider several performance metrics: 1) As an optimal control problem, the objective is to minimize state regulation cost (smaller is better). 2) Peak knee error can directly reflect amputee safety, preventing falling and stumbling (smaller is better). 3) Symmetry in walking can prevent secondary injury (closer to 0 is better). 4) Fast learning in terms of fewer tuning steps is practically important to amputees (fewer steps is better). 5) High success rate boosts amputees’ confidence and hence walking performance (higher is better). Details of how the data was obtained can be found in Appendix E.2. The main evaluation results are presented in Table 1. The second value in each entry (i.e., after the \pm symbol) represents the standard deviation of the performance metrics.

Benchmark study. MADDPG and COMA differ from our proposed cMARL approach toward symmetry walking. MARL problems can vary greatly, same are expected of their solutions [66, 67].

To perform benchmark studies, we tailor MADDPG and COMA, respectively to our HPC problem. Details are provided in Appendix D.

Table 1: Performance of the implemented algorithms in terms of the five performance criteria. The results of the best-performing algorithm for each criterion are boldfaced

Performance	w/human	wout/human	COMA	MADDPG
Stage Cost	0.002 ± 0.001	0.008 ± 0.003	0.004 ± 0.002	0.38 ± 0.306
Peak Error	0.003 ± 0.001	0.007 ± 0.002	0.003 ± 0.001	0.025 ± 0.014
Symmetry	0.001 ± 0.001	0.006 ± 0.003	0.003 ± 0.002	0.022 ± 0.012
Converge Steps	97 ± 88.3	187.5 ± 105.2	136 ± 121.8	—
Success Rate	0.7	0.58	0.5	—




Figure 2: The total number of steps needed to reach convergence in training (left). The success rate during evaluation (right).

Figure 2 shows that our cMARL solution has the best convergence profile and success rate over the baselines. Further results on training and evaluation are shown in Figures 3 and 4 which compare the environment sample efficiency as well as algorithm performances. To answer the first question based on benchmarking, our cMARL solution outperforms the baselines both in terms of kinematic and symmetry measurements as shown in Figures 3 and 4, center and right panels for training and evaluation, respectively. Additionally, the left panels of Figures 2, 3 and 4 show that our cMARL solution outperforms benchmarks with at least 30% less environment samples.

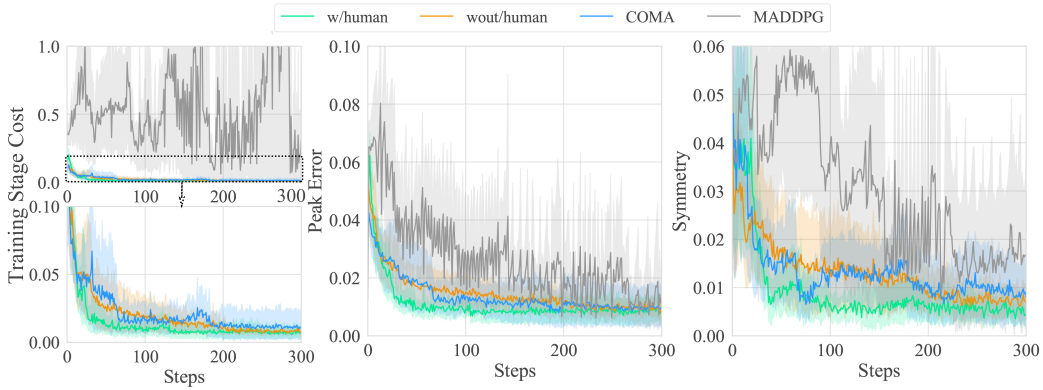


Figure 3: Learning curves of stage cost during training (left), peak angle error (middle) and symmetry of step length (right) for benchmark (w/human, COMA, MADDPG) and ablation (w/human, wout/human) studies. Each learning curve is averaged over 16 different random seeds and shaded by their respective 95% confidence interval.

Ablation study. To gain insights on how human control influences human-robot walking performance, an ablation study is carried out with direct human influenced terms removed, including h_k and v_k ,




Figure 4: Learning curves of averaged total cost (terms associated with human influence removed) during evaluation (left), peak angle error (middle) and symmetry of step length (right) for benchmark (w/human, COMA) and ablation (w/human, wout/human) studies. Each (smoothed) learning curve is averaged over 5 different random seeds and shaded by their respective 95% confidence interval.

from the problem formulation. This only leaves the kinematics and symmetry measurements in the stage cost. Similar to the benchmark study, the ablation study has a training session and evaluation session. For comparable results, the total cost in evaluation is obtained as $U = (x)^T (R_x) x + u^T R_u u$. Figures 3 and 4 show that with an estimated human control influence accounted for in the robot control design, our cMARL solution (the green curves) outperforms the one without it (orange curves). Treating human and robot as collaborating agents toward a shared performance goal, our cMARL solution approach achieves increased success rate and accelerated learning speed (Figure 2). This result makes sense as an estimated human control provides a predictive signal to the robot control which aims at duplicating the intact human joint movement. Additional information on the quality of estimated human control is given in Appendix B.2.

Reliability study. To make the proposed cMARL method practical and useful in real life, we setup two new walking tasks: (1) slope walking (11.5 degree ramp) and (2) walking at an increased pace (1.12m/s). They will result in different walking patterns from those used in the baseline study, and thus different knee joint profiles. For slope walking, knee flexion will be more pronounced during the stance phase since it walks inclined. In the case of faster pace, stance time will be compressed. To carry out the tests, the same training procedure is used as in benchmark and ablation studies.

Figure. 5 shows the performance of cMARL during slope walking and increased pace walking tasks. Performance of the two new tasks follow the same trend as that of the level ground walking at a nominal pace. These results again validate our design approach of using the intact knee movement trajectory as the target for the robotic knee to copy. By doing so, we have removed a major control design barrier in the way of performing different walking tasks by automatic control.




Figure 5: Learning curves of stage cost during training (Left), peak angle error (Middle) and symmetry of step length (Right) for different walking tasks (level ground walking at increased pace and ramp walking). Each learning curve is averaged over 16 different random seeds and shaded by their respective 95% confidence interval.

Limitations of This Study. Here we conduct an in-depth simulation-based analyses as an important first step to evaluate this novel cMARL solution to the HPC problem. Simulations are critically important as exploration of problem formulation, control algorithm design, and systematic evaluation are necessary to be performed prior to human experiment due to factors such as human fatigue, human safety, human loss of interest/confidence caused by repeated trial-and-error, time spent, and significant cost associated with testing amputee subjects. However, our framework is still to be tested in human experiment. Based on several important works in the literature [37, 40, 68, 69], extensive simulation studies followed by real life human test studies [41, 17, 70] have proven a highly successful development procedure for human/robot control. This will be the next step of this study.

For scenarios in which the terrain or task has changed significantly, a task planner will become necessary, making the problem a planning and control problem (as opposed to automatic control, the focus of this work). This expanded automatic control algorithm must be extensively verified in simulation and then in human tests before it can be integrated into a real-world planning framework for daily use cases. Such planning frameworks constitute intended future works.

6 Conclusion and discussion

1) In the US, approximately 1.7 million people live with limb loss. The amputee population is expected to double by 2050 as the population ages and incidence of dysvascular disease increases. As most lower limb amputees use prosthetic legs to restore basic bipedal locomotion, our solution to the prosthesis control problem can potentially help improve the function and quality of life of lower limb amputees. 2) In this work, we develop a novel cMARL framework towards systematically integrating the human and robot as collaborative agents to achieve normative walking toward solving real world problems. With reaching symmetric locomotion as shared control performance goal, we demonstrate improved walking performance. 3) Symmetry is selected as the shared goal for the collaborating agents because asymmetric walking has been linked to secondary health complications including back pain and osteoarthritis [71, 72, 73]. Although human-robot walking performance goals are difficult to systematically catalogue, additional considerations such as embodiment of the robot into the human will be considered in future works. 4) By breaking apart the shared cost for the human and the robot (cf. Section 5.1), the symmetrical walking task is treated and evaluated by MADDPG and COMA, respectively. Simulation results show that the factorization-based CTDE paradigm struggles to address the human-robot problem. The observed performance issues with factorization likely stem from the intrinsic coupling between the human and robot agents. 5) While ensuring human user safety has been carefully considered during control design, additional analysis of important properties such as convergence of learning, (sub)optimality of control policy, and human-robot closed-loop stability is still needed for this framework. Encouragingly, previous related works [74, 75, 76, 77] indicate that these theoretical results are very likely to be provable.

7 Acknowledgments and Disclosure of Funding

This research was supported in part by NSF under grants 1563921, 1808752 and 2211740 for Si; grants 156454, 1808898 and 2211739 for Huang. Zhikai Yao participated in early stage discussion involving the formulation of shared performance goal.

References

- [1] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [2] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- [3] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

- [4] Martin Lawitzky, Alexander Mörtl, and Sandra Hirche. Load sharing in human-robot cooperative manipulation. In *19th International Symposium in Robot and Human Interactive Communication*, pages 185–191. IEEE, 2010.
- [5] Antoine Bussy, Pierre Gergondet, Abderrahmane Kheddar, François Keith, and André Crosnier. 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*, pages 962–967. IEEE, 2012.
- [6] Hongbo Wang and Kazuhiro Kosuge. Control of a robot dancer for enhancing haptic human-robot interaction in waltz. *IEEE transactions on haptics*, 5(3):264–273, 2012.
- [7] David Vogt, Simon Stepputtis, Steve Grehl, Bernhard Jung, and Heni Ben Amor. A system for learning continuous human-robot interactions from human-human demonstrations. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2882–2889. IEEE, 2017.
- [8] Shuhei Ikemoto, Takashi Minato, and Hiroshi Ishiguro. Analysis of physical human–robot interaction for motor learning with physical help. *Applied Bionics and Biomechanics*, 5(4):213–223, 2008.
- [9] Biao Ma, Yulong Liu, Xiaoxiang Na, Yahui Liu, and Yiyong Yang. A shared steering controller design based on steer-by-wire system considering human-machine goal consistency. *Journal of the Franklin Institute*, 356(8):4397–4419, 2019.
- [10] Dylan P Losey, Craig G McDonald, Edoardo Battaglia, and Marcia K O’Malley. A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. *Applied Mechanics Reviews*, 70(1), 2018.
- [11] Thomas B Sheridan. Human–robot interaction: status and challenges. *Human factors*, 58(4):525–532, 2016.
- [12] Agostino De Santis, Bruno Siciliano, Alessandro De Luca, and Antonio Bicchi. An atlas of physical human–robot interaction. *Mechanism and Machine Theory*, 43(3):253–270, 2008.
- [13] Panagiota Tsarouchi, Sotiris Makris, and George Chryssolouris. Human–robot interaction review and challenges on task planning and programming. *International Journal of Computer Integrated Manufacturing*, 29(8):916–931, 2016.
- [14] Yue Wen, Minhan Li, Jennie Si, and He Huang. Wearer-prosthesis interaction for symmetrical gait: a study enabled by reinforcement learning prosthesis control. *IEEE transactions on neural systems and rehabilitation engineering*, 28(4):904–913, 2020.
- [15] Minhan Li, Yue Wen, Xiang Gao, Jennie Si, and He Huang. Toward expedited impedance tuning of a robotic prosthesis for personalized gait assistance by reinforcement learning control. *IEEE Transactions on Robotics*, 38(1):407–420, 2022.
- [16] Minhan Li, Xiang Gao, Yue Wen, Jennie Si, and He Helen Huang. Offline policy iteration based reinforcement learning controller for online robotic knee prosthesis parameter tuning. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2831–2837. IEEE, 2019.
- [17] Ruofan Wu, Minhan Li, Zhikai Yao, Wentao Liu, Jennie Si, and He Huang. Reinforcement learning impedance control of a robotic prosthesis to coordinate with human intact knee motion. *IEEE Robotics and Automation Letters*, 7(3):7014–7020, 2022.
- [18] Kenton R Kaufman, Serena Frittoli, and Carlo A Frigo. Gait asymmetry of transfemoral amputees using mechanical and microprocessor-controlled prosthetic knees. *Clinical Biomechanics*, 27(5):460–465, 2012.
- [19] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.

- [20] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [21] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- [22] Scott Fujimoto, Herke Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, pages 1587–1596. PMLR, 2018.
- [23] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018.
- [24] Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer Hicks, Sergey Levine, Christopher G Atkeson, and Scott L Delp. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *Journal of neuroengineering and rehabilitation*, 18(1):1–17, 2021.
- [25] Russell Enns and Jennie Si. Apache helicopter stabilization using neural dynamic programming. *Journal of guidance, control, and dynamics*, 25(1):19–25, 2002.
- [26] Russell Enns and Jennie Si. Helicopter trimming and tracking control using direct neural dynamic programming. *IEEE transactions on neural networks*, 14(4):929–939, 2003.
- [27] Russell Enns and Jennie Si. Helicopter flight-control reconfiguration for main rotor actuator failures. *Journal of guidance, control, and dynamics*, 26(4):572–584, 2003.
- [28] Chao Lu, Jennie Si, and Xiaorong Xie. Direct heuristic dynamic programming for damping oscillations in a large power system. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 38(4):1008–1013, 2008.
- [29] Wentao Guo, Feng Liu, Jennie Si, Dawei He, Ronald Harley, and Shengwei Mei. Approximate dynamic programming based supplementary reactive power control for dfig wind farm to enhance power system stability. *Neurocomputing*, 170:417–427, 2015.
- [30] Jingyi Zhang, Yonghong Luo, Boya Wang, Chao Lu, Jennie Si, and Jie Song. Deep reinforcement learning for load shedding against short-term voltage instability in large power systems. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [31] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, Pieter Abbeel, and Wojciech Zaremba. Hindsight experience replay. *arXiv preprint arXiv:1707.01495*, 2017.
- [32] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [33] Qingsong Ai, Zemin Liu, Wei Meng, Quan Liu, and Sheng Q Xie. Machine learning in robot assisted upper limb rehabilitation: A focused review. *IEEE Transactions on Cognitive and Developmental Systems*, 2021.
- [34] Dmitry Kalashnikov, Jacob Varley, Yevgen Chebotar, Benjamin Swanson, Rico Jonschkowski, Chelsea Finn, Sergey Levine, and Karol Hausman. Mt-opt: Continuous multi-task robotic reinforcement learning at scale. *arXiv preprint arXiv:2104.08212*, 2021.
- [35] A Isidori. *Nonlinear Control Systems: An Introduction*. Springer Verlag, Berlin, Germany, 1985.
- [36] Jennie Si and Yu-Tsung Wang. Online learning control by association and reinforcement. *IEEE Transactions on Neural networks*, 12(2):264–276, 2001.

- [37] Yue Wen, Jennie Si, Xiang Gao, Stephanie Huang, and He Helen Huang. A new powered lower limb prosthesis control framework based on adaptive dynamic programming. *IEEE transactions on neural networks and learning systems*, 28(9):2215–2220, 2016.
- [38] Ruofan Wu, Zhikai Yao, Jennie Si, and He Helen Huang. Robotic knee tracking control to mimic the intact human knee profile based on actor-critic reinforcement learning. *IEEE/CAA Journal of Automatica Sinica*, 9(1):19–30, 2021.
- [39] Xiang Gao, Jennie Si, Yue Wen, Minhan Li, and He Helen Huang. Knowledge-guided reinforcement learning control for robotic lower limb prosthesis. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 754–760. IEEE, 2020.
- [40] Xiang Gao, Jennie Si, Yue Wen, Minhan Li, and He Huang. Reinforcement learning control of robotic knee with human-in-the-loop by flexible policy iteration. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10):5873–5887, 2022.
- [41] Yue Wen, Jennie Si, Andrea Brandt, Xiang Gao, and He Huang. Online reinforcement learning control for the personalization of a robotic knee prosthesis. *IEEE transactions on cybernetics*, 2019.
- [42] Wentao Liu, Junmin Zhong, Ruofan Wu, Bretta L Fylstra, Jennie Si, and He Helen Huang. Inferring human-robot performance objectives during locomotion using inverse reinforcement learning and inverse optimal control. *IEEE Robotics and Automation Letters*, 2022.
- [43] Laetitia Matignon, Guillaume J Laurent, and Nadine Le Fort-Piat. Independent reinforcement learners in cooperative markov games: a survey regarding coordination problems. *The Knowledge Engineering Review*, 27(1):1–31, 2012.
- [44] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems*, 30, 2017.
- [45] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [46] Stefanos Nikolaidis and Julie Shah. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 33–40. IEEE, 2013.
- [47] Stefanos Nikolaidis, Yu Xiang Zhu, David Hsu, and Siddhartha Srinivasa. Human-robot mutual adaptation in shared autonomy. In *2017 12th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 294–302. IEEE, 2017.
- [48] Shervin Javdani, Siddhartha S Srinivasa, and J Andrew Bagnell. Shared autonomy via hindsight optimization. *Robotics science and systems: online proceedings*, 2015, 2015.
- [49] Siddharth Reddy, Anca D Dragan, and Sergey Levine. Shared autonomy via deep reinforcement learning. *arXiv preprint arXiv:1802.01744*, 2018.
- [50] Seyed Farokh Atashzar, Mahya Shahbazi, and Rajni V Patel. Haptics-enabled interactive neurorhabilitation mechatronics: classification, functionality, challenges and ongoing research. *Mechatronics*, 57:1–19, 2019.
- [51] S Farokh Atashzar, Hsien-Yung Huang, Fulvia Del Duca, Etienne Burdet, and Dario Farina. Energetic passivity decoding of human hip joint for physical human-robot interaction. *IEEE Robotics and Automation Letters*, 5(4):5953–5960, 2020.
- [52] Tingfang Yan, Marco Cempini, Calogero Maria Oddo, and Nicola Vitiello. Review of assistive strategies in powered lower-limb orthoses and exoskeletons. *Robotics and Autonomous Systems*, 64:120–136, 2015.

- [53] Shraddha Srivastava, Pei-Chun Kao, Seok Hun Kim, Paul Stegall, Damiano Zanon, Jill S Higginson, Sunil K Agrawal, and John P Scholz. Assist-as-needed robot-aided gait training improves walking function in individuals following stroke. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 23(6):956–963, 2014.
- [54] Andres Martinez, Brian Lawson, Christina Durrrough, and Michael Goldfarb. A velocity-field-based controller for assisting leg movement during walking with a bilateral hip and knee lower limb exoskeleton. *IEEE Transactions on Robotics*, 35(2):307–316, 2018.
- [55] Wentao Liu, Ruofan Wu, Jennie Si, and He Huang. A new robotic knee impedance control parameter optimization method facilitated by inverse reinforcement learning. *IEEE Robotics and Automation Letters*, 7(4):10882–10889, 2022.
- [56] Charles T Leonard, RL Craik, and CA Oatis. The neurophysiology of human locomotion. In *Gait Analysis: Theory and Application*. Mosby-Year, 1995.
- [57] Gentaro Taga, Yoko Yamaguchi, and Hiroshi Shimizu. Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological cybernetics*, 65(3):147–159, 1991.
- [58] Ryan T Roemmich and Amy J Bastian. Closing the loop: from motor neuroscience to neurorehabilitation. *Annual review of neuroscience*, 41:415–429, 2018.
- [59] Neville Hogan. Impedance control: An approach to manipulation: Part i—theory. 1985.
- [60] Frank Sup, Amit Bohara, and Michael Goldfarb. Design and control of a powered transfemoral prosthesis. *The International journal of robotics research*, 27(2):263–273, 2008.
- [61] Slavka Viteckova, Patrik Kutilek, Zdenek Svoboda, Radim Krupicka, Jan Kauler, and Zoltan Szabo. Gait symmetry measures: A review of current and prospective methods. *Biomedical Signal Processing and Control*, 42:89–100, 2018.
- [62] Antonio Nardone, Marco Godi, Margherita Grasso, Simone Guglielmetti, and Marco Schieppati. Stabilometry is a predictor of gait performance in chronic hemiparetic stroke patients. *Gait & Posture*, 30(1):5–10, 2009.
- [63] Chris Mizelle, Mary Rodgers, and Larry Forrester. Bilateral foot center of pressure measures predict hemiparetic gait velocity. *Gait & Posture*, 24(3):356–363, 2006.
- [64] Yue Wen, Andrea Brandt, Jennie Si, and He Helen Huang. Automatically customizing a powered knee prosthesis with human in the loop using adaptive dynamic programming. In *2017 International Symposium on Wearable Robotics and Rehabilitation (WeRob)*, pages 1–2. IEEE, 2017.
- [65] Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering*, 54(11):1940–1950, 2007.
- [66] Afshin OroojlooyJadid and Davood Hajinezhad. A review of cooperative multi-agent deep reinforcement learning. *arXiv preprint arXiv:1908.03963*, 2019.
- [67] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. *Handbook of Reinforcement Learning and Control*, pages 321–384, 2021.
- [68] Yue Wen, Ming Liu, Jennie Si, and He Helen Huang. Adaptive control of powered transfemoral prostheses based on adaptive dynamic programming. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 5071–5074. IEEE, 2016.
- [69] Ruofan Wu, Zhikai Yao, Jennie Si, and He Helen Huang. Robotic knee tracking control to mimic the intact human knee profile based on actor-critic reinforcement learning. *IEEE/CAA Journal of Automatica Sinica*, 9(1):19–30, 2022.

- [70] Minhan Li, Yue Wen, Xiang Gao, Jennie Si, and He Huang. Toward expedited impedance tuning of a robotic prosthesis for personalized gait assistance by reinforcement learning control. *IEEE Transactions on Robotics*, 2021.
- [71] M Jason Highsmith, Casey R Andrews, Claire Millman, Ashley Fuller, Jason T Kahle, Tyler D Klenow, Katherine L Lewis, Rachel C Bradley, and John J Orriola. Gait training interventions for lower extremity amputees: a systematic literature review. *Technology & Innovation*, 18(2-3):99–113, 2016.
- [72] Alberto Esquenazi. Gait analysis in lower-limb amputation and prosthetic rehabilitation. *Physical Medicine and Rehabilitation Clinics*, 25(1):153–167, 2014.
- [73] George NS Marinakis. Interlimb symmetry of traumatic unilateral transtibial amputees wearing two different prosthetic feet in the early rehabilitation stage. *Journal of rehabilitation research and development*, 41(4):581–590, 2004.
- [74] Feng Liu, Jian Sun, Jennie Si, Wentao Guo, and Shengwei Mei. A boundedness result for the direct heuristic dynamic programming. *Neural Networks*, 32:229–235, 2012.
- [75] Yury Sokolov, Robert Kozma, Ludmilla D Werbos, and Paul J Werbos. Complete stability analysis of a heuristic approximate dynamic programming control design. *Automatica*, 59:9–18, 2015.
- [76] Xiang Gao. *Data-Efficient Reinforcement Learning Control of Robotic Lower-Limb Prosthesis With Human in the Loop*. PhD thesis, Arizona State University, 2020.
- [77] Qingtao Zhao, Jennie Si, and Jian Sun. Online reinforcement learning control by direct heuristic dynamic programming: From time-driven to event-driven. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [78] SimTK. Tutorial 1 - intro to musculoskeletal modeling.
- [79] Daniel Jacobs. From the ground up: Building a passive dynamic walker model, 2014.
- [80] Mrn P Kadaba, HK Ramakrishnan, and ME Wootten. Measurement of lower extremity kinematics during level walking. *Journal of orthopaedic research*, 8(3):383–392, 1990.
- [81] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.
- [82] Peng Peng, Ying Wen, Yaodong Yang, Quan Yuan, Zhenkun Tang, Haitao Long, and Jun Wang. Multiagent bidirectionally-coordinated nets: Emergence of human-level coordination in learning to play starcraft combat games. *arXiv preprint arXiv:1703.10069*, 2017.
- [83] Jakob Foerster, Nantas Nardelli, Gregory Farquhar, Triantafyllos Afouras, Philip HS Torr, Pushmeet Kohli, and Shimon Whiteson. Stabilising experience replay for deep multi-agent reinforcement learning. In *International conference on machine learning*, pages 1146–1155. PMLR, 2017.
- [84] Juanjuan Zhang, Pieter Fiers, Kirby A Witte, Rachel W Jackson, Katherine L Poggensee, Christopher G Atkeson, and Steven H Collins. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science*, 356(6344):1280–1284, 2017.
- [85] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.
- [86] Jonathan Kofman, Xianghai Wu, Timothy J Luu, and Siddharth Verma. Teleoperation of a robot manipulator using a vision-based human-robot interface. *IEEE transactions on industrial electronics*, 52(5):1206–1219, 2005.