

Human Preference-Based Learning for High-dimensional Optimization of Exoskeleton Walking Gaits

Maegan Tucker¹, Myra Cheng², Ellen Novoseller², Richard Cheng¹,
Yisong Yue², Joel W. Burdick^{1,2}, and Aaron D. Ames^{1,2}

Abstract—Optimizing lower-body exoskeleton walking gaits for user comfort requires understanding users’ preferences over a high-dimensional gait parameter space. However, existing preference-based learning methods have only explored low-dimensional domains due to computational limitations. To learn user preferences in high dimensions, this work presents LINECOSPAR, a human-in-the-loop preference-based framework that enables optimization over many parameters by iteratively exploring one-dimensional subspaces. Additionally, this work identifies gait attributes that characterize broader preferences across users. In simulations and human trials, we empirically verify that LINECOSPAR is a sample-efficient approach for high-dimensional preference optimization. Our analysis of the experimental data reveals a correspondence between human preferences and objective measures of dynamics, while also highlighting differences in the utility functions underlying individual users’ gait preferences. This result has implications for exoskeleton gait synthesis, an active field with applications to clinical use and patient rehabilitation.

I. INTRODUCTION

Human-in-the-loop online learning techniques have demonstrated significant potential in human-robot interaction tasks [1]–[3], such as improving the performance of robotic assistive devices. In particular, online learning from human feedback can help to optimize walking gaits for lower-body exoskeletons [4]–[6], which are placed over existing limbs to assist mobility-impaired individuals.

This work focuses on optimizing walking gaits for individual user comfort using the Atalante lower-body exoskeleton developed by Wandercraft. We use a pre-computed gait library, which generates gaits offline using optimization-based techniques from nonlinear dynamics and control [7]–[9]. Gaits are specified by parameters ranging from centers of pressure to step dimensions (step length, width, etc.).

Optimizing gait parameters for each exoskeleton user serves two purposes. First, it enables gait personalization to maximize each user’s comfort. Second, the relationships among different users’ preferences, in particular their optimal gaits, may provide insight into the properties of universally-preferred gaits. While some gait optimization approaches rely on numeric metrics such as the user’s metabolic expenditure [5], there are no metrics that have established




Fig. 1. Atalante Exoskeleton. The exoskeleton has 12 actuated joints. The experiments explore six exoskeleton gait parameters: step length, step duration, step width, maximum step height, pelvis roll, and pelvis pitch.

correspondences with user comfort. For example, metabolic expenditure is not an appropriate metric as the exoskeleton does not require the user to expend effort towards walking. A quantitative understanding of human preferences could help generate new gait profiles to improve the existing gait library, which represents a small fraction of the rich space of human walking behaviors. This motivates optimizing over the high-dimensional space of exoskeleton gaits to characterize the utility functions governing users’ gait preferences.

We rely on users’ pairwise preferences to learn exoskeleton gaits that optimize user comfort, as several studies have shown that for subjective human feedback, preferences are more reliable than numerical scores [10]–[12]. While interactive preference learning methods have been applied to robotics [4], [6], existing online preference learning methods are restricted to low-dimensional domains due to computational limitations; for example, previous work on preference-based exoskeleton gait optimization either learns over at most two dimensions [4] or utilizes domain knowledge to narrow the search space before performing online learning [6].

We present LINECOSPAR, a high-dimensional human preference-based learning approach that integrates existing techniques for preference learning [4] and high-dimensional optimization [13] into a unified framework. LINECOSPAR relies on preference feedback to iteratively explore one-dimensional subspaces. We demonstrate in simulation that LINECOSPAR exhibits sample-efficient convergence to user-preferred actions in high-dimensional spaces. The algorithm is then deployed experimentally to optimize exoskeleton

*This work was supported by NSF NRI award 1724464, NSF Graduate Research Fellowship No. DGE1745301, the Caltech Big Ideas Fund, and the ZEITLIN Fund.

This work was conducted under IRB No. 16-0693.

¹Authors are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125

²Authors are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125

walking over six gait parameters for six able-bodied subjects.

Using the preferred gaits identified by LINECOSPAR in the human experiments, we examine the connections among user-preferred gaits to understand what makes some gaits preferable to others. An analysis using the Zero Moment Point [14] reveals that users’ preferences correspond to quantitative metrics of dynamicity. We observe that while most users’ preferences are consistent with a metric that prioritizes dynamic stability, one user’s preferences are explained by metrics that favor static stability. Based on this analysis, we suggest metrics that should be considered in the gait generation process, such that future exoskeleton gait designs can draw from regions of the gait trajectory space that prioritize user comfort.

II. THE LEARNING ALGORITHM

The LINECOSPAR algorithm (Alg. 1) learns a Bayesian model over the user’s preferences in a high-dimensional space. To learn from preferences, we adopt the dueling bandit setting [15]–[17], in which the algorithm selects actions and receives relative preferences between them. The procedure, based on Thompson sampling, iterates through: 1) updating a Bayesian posterior over the actions’ utilities given the data, 2) sampling utility functions from the posterior, 3) executing the actions that maximize the sampled utility functions, and 4) observing preferences among the executed actions.

Drawing inspiration from the LINEBO algorithm [13], LINECOSPAR exploits low-dimensional structure in the search space by sequentially considering one-dimensional subspaces from which to sample actions. This allows the algorithm to maintain its Bayesian preference relation function over a subset of the action space in each iteration. LINECOSPAR builds upon COSPAR, which finds user-preferred parameters across one and two dimensions [4]. Compared to COSPAR, LINECOSPAR learns the model posterior much more efficiently and can be scaled to higher dimensions.

This section provides background on existing approaches and then describes the LINECOSPAR algorithm, including 1) defining the posterior updating procedure, 2) achieving high-dimensional learning, and 3) incorporating Thompson sampling and coactive feedback.

A. Background

Preference-Based Learning. We learn users’ preferred exoskeleton gaits through their relative preferences, which are more reliable than subjective numerical feedback [4], [10]–[12]. To maximize sample efficiency, we adopt the mixed-initiative approach of COSPAR [4], which learns from both pairwise preference and coactive feedback. In coactive learning [18], [19], after each time the algorithm selects an action, the user identifies an improved action. Under both feedback types, the exoskeleton user tests various gaits to specify preferences and suggest gait modifications. COSPAR effectively identifies user-preferred gait parameters across one and two dimensions. However, COSPAR is intractable in larger action spaces, as it jointly maintains and samples from a posterior

Algorithm 1 LINECOSPAR

```

1: procedure LINECOSPAR(Utility prior parameters;  $m$  = granularity of discretization)
2:    $\mathcal{D} = \emptyset, \mathcal{W} = \emptyset$   $\triangleright$   $\mathcal{D}$ : preference data,  $\mathcal{W}$ : actions in  $\mathcal{D}$ 
3:   Set  $\mathbf{p}_1, \mathbf{a}_0$  to uniformly-random actions
4:   for  $t = 1, 2, \dots, T$  do
5:      $\mathcal{L}_t$  = random line through  $\mathbf{p}_t$ , discretized via  $m$ 
6:      $\mathcal{V}_t = \mathcal{L}_t \cup \mathcal{W}$   $\triangleright$  Points over which to update posterior
7:      $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$  = posterior over points in  $\mathcal{V}_t$ , given  $\mathcal{D}$ 
8:     Sample utility function  $f_t \sim \mathcal{N}(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ 
9:     Execute action  $\mathbf{a}_t = \operatorname{argmax}_{\mathbf{a} \in \mathcal{V}_t} f_t(\mathbf{a})$ 
10:    Add pairwise preference between  $\mathbf{a}_t$  and  $\mathbf{a}_{t-1}$  to  $\mathcal{D}$ 
11:    Add coactive feedback  $\mathbf{a}'_t$  to  $\mathcal{D}$ 
12:    Set  $\mathcal{W} = \mathcal{W} \cup \{\mathbf{a}_t\} \cup \{\mathbf{a}'_t\}$   $\triangleright$  Update actions in  $\mathcal{D}$ 
13:    Set  $\mathbf{p}_{t+1} = \operatorname{argmax}_{\mathbf{a} \in \mathcal{V}_t} \mu_t(\mathbf{a})$ 
14:  end for
15: end procedure

```

over every action, causing the computational complexity to increase exponentially in the action space dimension.

High-Dimensional Bayesian Optimization. Bayesian optimization is a powerful approach for optimizing expensive-to-evaluate black-box functions. It maintains a model posterior over the unknown function, and cycles through a) using the posterior to acquire actions at which to query the function, b) querying the function, and c) updating the posterior using the obtained data. This procedure is challenging in high-dimensional search spaces due to the computational cost of the acquisition step (a), which often requires solving a non-convex optimization problem over the search space, and maintaining the posterior in the update step (c), which can require manipulating matrices that grow exponentially with the action space’s dimension. Dimensionality reduction techniques are therefore an area of active interest. Solutions vary from optimizing variable subsets (DROPOUTBO) [20] to projecting into lower-dimensional spaces (REMBO) [21] to sequentially optimizing over one-dimensional subspaces (LINEBO) [13]. We draw upon the approach of LINEBO because of its state-of-the-art performance in high-dimensional spaces. Furthermore, it is especially sample-efficient in spaces with underlying low-dimensional structure. In the case of exoskeleton walking, this low-dimensional structure may appear as linear relationships between two gait parameters in the user’s utility function, i.e., users who prefer short step lengths also prefer short step durations.

B. The LINECOSPAR Algorithm

Modeling Utilities Using Pairwise Preference Data. LINECOSPAR uses pairwise comparisons to learn a Bayesian model posterior over the relative utilities of actions (i.e., gait parameter combination) to the user based upon the Gaussian process preference model in [22]. We use Gaussian process learning, as it enables us to model a Bayesian posterior over a class of smooth, non-parametric functions.

Let $\mathcal{A} \subset \mathbb{R}^d$ be the set of possible actions. In iteration t of the algorithm, we consider a subset of the actions $\mathcal{V}_t \subset \mathcal{A}$, with cardinality V_t (we will define \mathcal{V}_t later). We assume that each action $\mathbf{a} \in \mathcal{A}$ has a latent utility

to the user, denoted as $f(\mathbf{a})$. Throughout the learning process, LINECOSPAR stores a dataset of all user feedback, $\mathcal{D} = \{\mathbf{a}_{k_1} \succ \mathbf{a}_{k_2} \mid k = 1, \dots, N\}$, consisting of N preferences, where $\mathbf{a}_{k_1} \succ \mathbf{a}_{k_2}$ indicates that the user prefers action \mathbf{a}_{k_1} to action \mathbf{a}_{k_2} . The preference data \mathcal{D} is used to update the posterior utilities of the actions in \mathcal{V}_t . Defining $\mathbf{f} = [f(\mathbf{a}_{t_1}), f(\mathbf{a}_{t_2}), \dots, f(\mathbf{a}_{t_{V_t}})]^T \in \mathbb{R}^{V_t}$, where \mathbf{a}_{t_i} is the i^{th} action in \mathcal{V}_t , the utilities \mathbf{f} have posterior:

$$\mathcal{P}(\mathbf{f}|\mathcal{D}) \propto \mathcal{P}(\mathcal{D}|\mathbf{f})\mathcal{P}(\mathbf{f}). \quad (1)$$

In each iteration t , we define a Gaussian process prior over the utilities \mathbf{f} of actions in \mathcal{V}_t :

$$\mathcal{P}(\mathbf{f}) = \frac{1}{(2\pi)^{V_t/2} |\Sigma_t^{\text{pr}}|^{1/2}} \exp\left(-\frac{1}{2} \mathbf{f}^T [\Sigma_t^{\text{pr}}]^{-1} \mathbf{f}\right), \quad (2)$$

where $\Sigma_t^{\text{pr}} \in \mathbb{R}^{V_t \times V_t}$, $[\Sigma_t^{\text{pr}}]_{ij} = \mathcal{K}(\mathbf{a}_{t_i}, \mathbf{a}_{t_j})$, and \mathcal{K} is a kernel. Our experiments use the squared exponential kernel. To compute the likelihood $\mathcal{P}(\mathcal{D}|\mathbf{f})$, we assume that the preferences may be corrupted by noise, such that:

$$\mathcal{P}(\mathbf{a}_{k_1} \succ \mathbf{a}_{k_2} | \mathbf{f}) = g\left(\frac{f(\mathbf{a}_{k_1}) - f(\mathbf{a}_{k_2})}{c}\right), \quad (3)$$

where $g(\cdot) \in [0, 1]$ is a monotonically-increasing link function, and $c > 0$ is a hyperparameter indicating the degree of preference noise. While previous work uses the Gaussian cumulative distribution function for g [4], [22], we empirically found that using the heavier-tailed sigmoid distribution, $g_{\text{sig}}(x) := \frac{1}{1+e^{-x}}$, as the link function improves performance. $g_{\text{sig}}(x)$ satisfies the convexity conditions for the Laplace approximation [23] and has been used to model preferences in other contexts [24]. The full likelihood expression becomes:

$$\mathcal{P}(\mathcal{D}|\mathbf{f}) = \prod_{k=1}^N g_{\text{sig}}\left(\frac{f(\mathbf{a}_{k_1}) - f(\mathbf{a}_{k_2})}{c}\right). \quad (4)$$

The posterior in (1) is estimated via the Laplace approximation as in [22], yielding a multivariate Gaussian, $\mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$.

Sampling Approach for Higher Dimensions. Existing preference-based approaches optimize over the action space \mathcal{A} by discretizing the entire space before beginning the learning process. This results in m^d combinations from m uniformly-spaced points (corresponding to actions) in each of the d dimensions of \mathcal{A} . Thus, the cardinality of this set is $A := |\mathcal{A}| = m^d$; larger m enables finer-grained search at a higher computational cost. The Bayesian preference model is updated over all A points during each iteration. This is intractable for higher d since computing the posterior over A points involves expensive matrix operations, such as inverting $\Sigma_t^{\text{pr}}, \Sigma_t \in \mathbb{R}^{A \times A}$.

Inspired by [13], LINECOSPAR overcomes this intractability by iteratively considering one-dimensional subspaces (lines), rather than the full action space. In each iteration t , LINECOSPAR selects uniformly-spaced points along a new random line \mathcal{L}_t in the action space, which is determined by a uniformly-random direction and the action \mathbf{p}_t that maximizes the posterior mean. Including \mathbf{p}_t in the subspace encourages exploration of higher-utility areas. The posterior $\mathcal{P}(\mathcal{D}|\mathbf{f})$ is

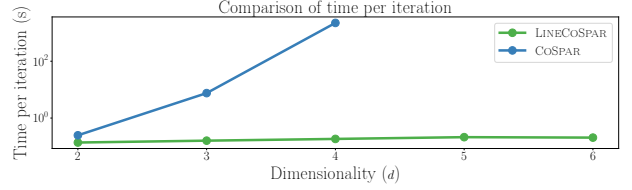


Fig. 2. **Curse of dimensionality for COSPAR.** Average time per iteration of COSPAR vs. LINECOSPAR. The y-axis is on a logarithmic scale. For LINECOSPAR, the time is roughly constant in the number of dimensions d , while the runtime of COSPAR increases exponentially. For $d = 4$, the duration of a COSPAR iteration is inconvenient in the human-in-the-loop learning setting, and for $d \geq 5$, it is intractable.

calculated over $\mathcal{V}_t := \mathcal{L}_t \cup \mathcal{W}$, where \mathcal{W} is the set of actions for which \mathcal{D} contains preference feedback. This approach reduces the model’s covariance matrices $\Sigma_t^{\text{pr}}, \Sigma_t$ from size $A \times A$ to $V_t \times V_t$. Rather than growing exponentially in d , which is impractical for online learning, LINECOSPAR’s complexity is constant in the dimension d and linear in the number of iterations T . Since queries are expensive in many human-in-the-loop robotics settings, T is typically low.

Posterior Sampling Framework. Utilities are learned using the SELFSARRING [15] approach to Thompson sampling detailed above. Specifically, in each iteration, we calculate the posterior of the utilities \mathbf{f} over the points in $\mathcal{V}_t = \mathcal{L}_t \cup \mathcal{W}$, obtaining the posterior $\mathcal{N}(\boldsymbol{\mu}_t, \Sigma_t)$ over \mathcal{V}_t . The algorithm then samples a utility function f_t from the posterior, which assigns a utility to each action in \mathcal{V}_t . Next, LINECOSPAR executes the action \mathbf{a}_t that maximizes f , $\mathbf{a}_t = \operatorname{argmax}_{\mathbf{a} \in \mathcal{V}_t} f(\mathbf{a})$. The user provides a preference (or indicates indifference, i.e. “no preference”) between \mathbf{a}_t and the preceding action \mathbf{a}_{t-1} .

In addition, for each executed action \mathbf{a}_t , the user can provide coactive feedback, specifying the dimension, direction (higher or lower), and degree in which to change \mathbf{a}_t . The user’s suggested action \mathbf{a}'_t is added to \mathcal{W} , and the feedback is added to \mathcal{D} as $\mathbf{a}'_t \succ \mathbf{a}_t$. In each iteration, preference and coactive feedback each add at most one action to \mathcal{W} . Thus, in iteration t , \mathcal{V}_t contains at most $m + 2(t - 1)$ actions, and so its size is independent of the dimensionality d . In the subsequent analysis, \mathbf{a}_{\max} is defined as the action maximizing the final posterior mean after T iterations, i.e., $\mathbf{a}_{\max} := \operatorname{argmax}_{\mathbf{a} \in \mathcal{V}_t} \mu_{T+1}(\mathbf{a})$.

III. PERFORMANCE OF LINECOSPAR

A. Simulation Results

We validate the performance of LINECOSPAR in simulation using both standard Bayesian optimization benchmarks and randomly-generated polynomials.¹ The simulations show that LINECOSPAR is sample-efficient, converges to sampling higher-valued actions, and learns a preference relation function such that actions with higher objective values have high posterior utilities.

Standard Bayesian Optimization Benchmarks. We evaluated the performance of LINECOSPAR on the standard Hart-

¹The code is at <https://github.com/myracheng/linecospar>. All experiments use the squared exponential kernel with lengthscale 0.15 in every dimension, signal variance $1e-4$, noise variance $1e-5$, and preference noise 0.005.

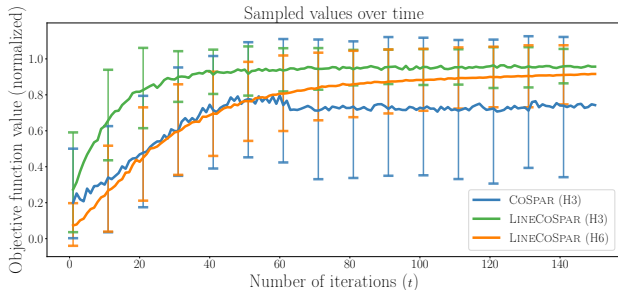


Fig. 3. **Convergence to higher values on standard benchmarks.** Mean objective value \pm SD using H3 and H6, averaged over 100 runs. The sampled actions converge to higher objective values at a faster rate with LINECOSPAR, which has an improved sampling approach and link function. It is intractable to run COSPAR on a 6-dimensional space.

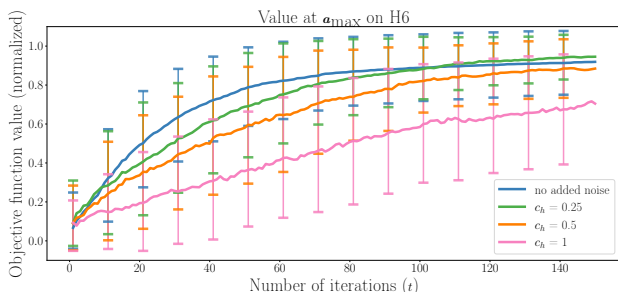


Fig. 4. **Robustness to noisy preferences.** Mean objective value \pm SD of the action \mathbf{a}_{\max} with the highest posterior utility. This is averaged over 100 runs using LINECOSPAR on H6 with varying preference noise, as quantified by c_h . Higher performance correlates with less noise (lower c_h). The algorithm is robust to noise to a certain degree ($c_h \leq 0.5$).

mann3 (H3) and Hartmann6 (H6) benchmarks (3 and 6 dimensions, respectively). We do not compare LINECOSPAR to other optimization methods because there are no other preference-based Gaussian process methods that are tractable in high dimensions. As discussed in Section II-B, we focus on Gaussian process methods because they model smooth, non-parametric utility functions. We validate LINECOSPAR with noiseless preferences and then demonstrate its robustness to noisy user preferences. Preferences are generated in simulation by comparing objective function values.

Under ideal preference feedback, $\mathbf{a}_{k_1} \succ \mathbf{a}_{k_2}$ if $f(\mathbf{a}_{k_1}) > f(\mathbf{a}_{k_2})$. The true objective values f are invisible to the algorithm, which observes only the preference dataset \mathcal{D} . Compared to COSPAR, LINECOSPAR converges to sampling actions with higher objective values at a faster rate (Fig. 3). Thus, LINECOSPAR not only enables higher-dimensional optimization, but also improves speed and accuracy of learning.

Since human preferences may be noisy, we tested the algorithm’s robustness to noisy preference feedback. In simulation, this is modeled via $\mathcal{P}(\mathbf{a}_{k_1} \succ \mathbf{a}_{k_2}) = (1 + e^{-\frac{s_k}{c_h}})^{-1}$, where $s_k = f(\mathbf{a}_{k_1}) - f(\mathbf{a}_{k_2})$ and c_h is a hyperparameter for the noise level. As $c_h \rightarrow \infty$, the preferences approach uniform randomness (i.e. become noisier). Also, actions become less distinguishable when the distance between $f(\mathbf{a}_{k_1})$ and $f(\mathbf{a}_{k_2})$ decreases. This reflects human preference generation since it is more difficult to give consistent preferences between actions with similar utilities. By simulating noisy

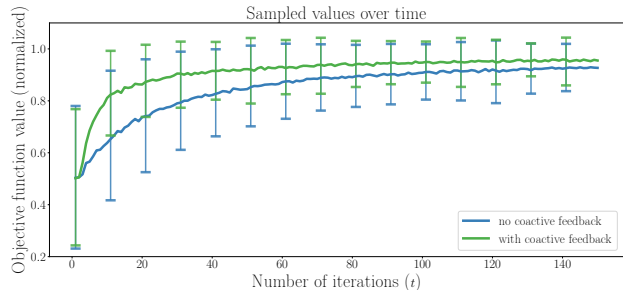


Fig. 5. **Coactive feedback improves convergence.** Mean objective value \pm SD of the sampled actions using random functions. This is averaged over 1000 runs using LINECOSPAR on 100 randomly-generated six-dimensional functions ($d = 6$). The sampled actions converge to high objective values in relatively few iterations, and coactive feedback accelerates this process.

preferences, we demonstrate that LINECOSPAR is robust to noisy feedback (see Fig. 4).

Randomly-Generated Functions. We also tested LINECOSPAR using randomly-generated d -dimensional polynomials (for $d = 6$) as objective functions: $p(\mathbf{a}) = \sum_{i=1}^d \alpha_i \sum_{j=1}^d \beta_j a_j$, where a_j denotes the j^{th} element of \mathbf{a} , and $\alpha_i, \beta_i, i \in \{1, \dots, d\}$ are sampled independently from the uniform distribution $\mathcal{U}(-1, 1)$. The dimensions’ ranges and discretizations match those in the exoskeleton experiments, so that these simulations approximate the number of human trials needed to find optimal gaits.

Coactive feedback was simulated for each sampled action \mathbf{a}_t by finding an action \mathbf{a}'_t with a higher objective value that differs from \mathbf{a}_t along only one dimension. The action \mathbf{a}'_t is determined by randomly choosing a dimension in $\{1, \dots, d\}$ and direction (positive or negative), and taking a step from \mathbf{a}_t along this vector. If the resulting action \mathbf{a}'_t has a higher objective value, it is added to the dataset \mathcal{D} as $\mathbf{a}'_t \succ \mathbf{a}_t$. This is a proxy for the human coactive feedback acquired in the exoskeleton experiments described below, in which the user can suggest a dimension and direction in which to modify an action to obtain an improved gait.

Fig. 5 displays LINECOSPAR’s performance over 100 randomly-generated polynomials (10 repetitions each) with computation time shown in Fig. 2. The results demonstrate that LINECOSPAR samples high-valued actions within relatively few iterations (≈ 20 with coactive feedback).

B. Human Subject Experiments

After the performance of LINECOSPAR was demonstrated in simulation, the algorithm was experimentally deployed on the lower-body exoskeleton Atalante (Fig. 1) to optimize six gait parameters for six able-bodied users (see Table I for results and [25] for a video).

Atalante Exoskeleton. Atalante (Fig. 1) [8], [26], [27], developed by Wandercraft, has 12 actuated joints: three at each hip, one at each knee, and two in each ankle. [9] describes the device’s mechanical components and control architecture in detail. Exoskeleton walking is achieved using pre-computed walking gaits, generated using the partial hybrid zero dynamics framework [28] and a nonlinear constrained optimization

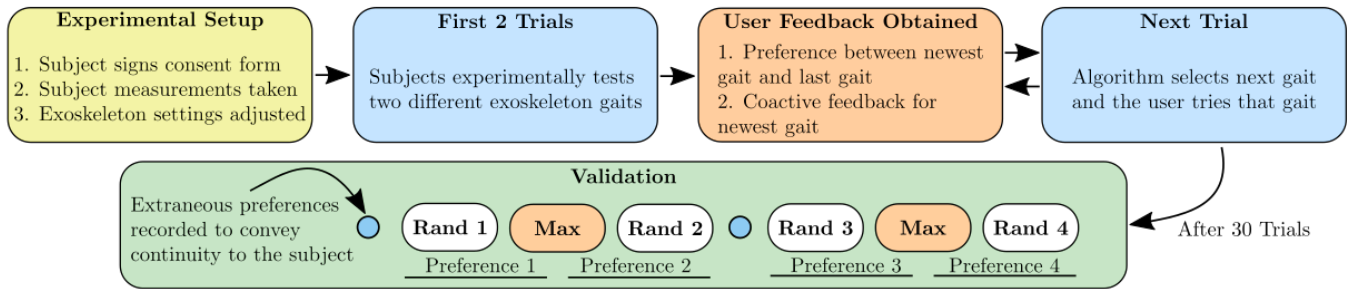


Fig. 6. **Experimental Procedure.** After setup of the subject-exoskeleton system, subjects were queried for preferences between all consecutive pairs of gaits, along with coactive feedback, in 30 gait trials (for a total of at most 29 pairwise preferences and 30 pieces of coactive feedback). After these 30 trials, the subject unknowingly entered the validation portion of the experiment, in which he/she validated the posterior-maximizing gait, \mathbf{a}_{\max} , against four randomly-selected gaits.

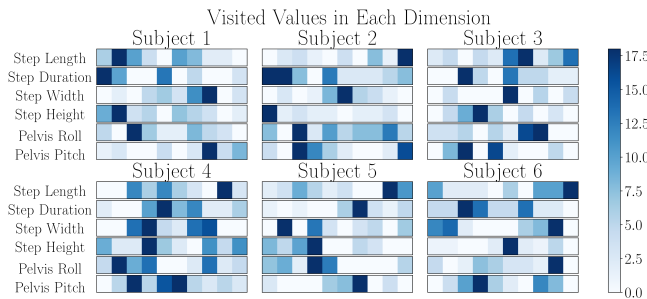


Fig. 7. **Exploration vs. exploitation in human trials.** Each row depicts the distribution of a particular gait parameter’s values across all gaits that the subject tested. Each dimension is discretized into 10 bins. Note that the algorithm explores different parts of the action space for each subject. These visitation frequencies exhibit a statistically-significant correlation with the posterior utilities across these regions (Pearson’s p -value = $1.22\text{e-}10$).

process that utilizes direct collocation. The configuration space of the human-exoskeleton system is constructed as $q = (p, \phi, q_b) \in Q \in \mathbb{R}^{18}$, where $p \in \mathbb{R}^3$ and $\phi \in \mathbb{SO}^3$ denote the position and orientation of the exoskeleton floating base frame with respect to the world frame, and $q_b \in \mathbb{R}^{12}$ denotes the relative angles of the actuated joints. The generated gaits are realized on the exoskeleton using PD control at the joint level and a high-level controller adjusting joint targets based on state feedback. The controller is executed by an embedded computer unit running a real-time operating system. Gaits are sent to the exoskeleton over a wireless connection via a custom graphical user interface.

Experimental Procedure. LINECOSPAR optimized exoskeleton gaits for six self-identified able-bodied subjects over six gait parameters (Fig. 1): step length, step duration, step width, maximum step height, pelvis roll, and pelvis pitch. These parameters were chosen from the pre-computed gait library because they are relatively intuitive for users to understand when giving coactive feedback. The parameter ranges, respectively, are: 0.08-0.18 meters, 0.85-1.15 seconds, 0.25-0.3 meters, 0.065-0.075 meters, 5.5-9.5 degrees, and 10.5-14.5 degrees. Fig. 6 illustrates the experimental procedure for testing and validating LINECOSPAR.

All subjects were volunteers without prior exoskeleton exposure. For each subject, the testing procedure lasted approximately two hours, with one hour of setup and one hour

of exoskeleton testing. The setup consisted of explaining the procedure (including how to provide preference and coactive feedback), measuring subject parameters, and adjusting the thigh and shank length of the exoskeleton to the subject. During the testing, the subjects had control over initiating and terminating each instance of exoskeleton walking and were instructed to try each walking gait until they felt comfortable giving a preference. The subjects could choose to test each gait multiple times to confirm their preference. They could also specify “no preference” between two gait trials, in which case no new information was added to the dataset \mathcal{D} .

After completing 30 trials (including trials with no preference, but not including voluntary gait repetitions), the subject began a set of “validation” trials; for consistency, the subject was not informed of the start of the validation phase. Validation consisted of six additional trials and yielded four pairwise preferences, each between the posterior-maximizing action \mathbf{a}_{\max} and a randomly-generated action. This validation step verifies that \mathbf{a}_{\max} is preferred over other parameter combinations across the search space.

Gait Optimization Results. Fig. 7 shows that the LINECOSPAR algorithm both explores across the gait parameter space and exploits regions with higher posterior utility. Over time, LINECOSPAR increasingly samples actions concentrated in regions of the search space that are preferred based on previous feedback. This results in a significant correlation between visitation frequencies and posterior utilities across these regions (Pearson’s p -value = $1.22\text{e-}10$).

For each subject, Table I lists the parameters of the predicted optimal gaits, \mathbf{a}_{\max} , identified by LINECOSPAR. Table I also illustrates the results of the validation trials for each subject. These results show that \mathbf{a}_{\max} was predominantly preferred over the randomly-selected actions during validation. For four of the six subjects, all four validation preferences matched the posterior, while the other subjects matched three and one of the four preferences, respectively.

IV. ANALYSIS OF PREFERENCE FEEDBACK AND IMPLICATIONS FOR GAIT SYNTHESIS

In addition to optimizing exoskeleton walking gaits for individual users, we aim to understand the utility functions underlying human preferences and apply this knowledge

TABLE I
GAIT PARAMETERS OPTIMIZING LINECOSPAR'S POSTERIOR MEAN (α_{MAX}) FOR EACH ABLE-BODIED SUBJECT

Subject	Height (m)	Mass (kg)	Step Length (m)	Step Duration (s)	Step Width (m)	Max Step Height (m)	Pelvis Roll (deg)	Pelvis Pitch (deg)	Validation Accuracy (%)
1	1.85	89.9	0.0835	0.943	0.278	0.0674	6.38	10.9	75
2	1.668	69.2	0.136	1.04	0.285	0.0679	6.41	12.4	100
3	1.635	51.2	0.137	0.922	0.279	0.0688	8.56	11.4	100
4	1.795	73.6	0.127	0.989	0.268	0.065	6.68	12.7	25
5	1.625	55.9	0.161	1.05	0.258	0.0689	7.32	13.2	100
6	1.66	65	0.177	1.11	0.256	0.0663	7.71	13.5	100

towards improving gait synthesis. As discussed in [4], exoskeleton gaits are generated using the partial hybrid zero dynamics framework, which is formulated by the following nonlinear optimization problem [8]:

$$\begin{aligned} \alpha^* = \underset{\alpha}{\operatorname{argmin}} \quad & \mathcal{J}(\alpha) \\ \text{s.t.} \quad & \Delta(\mathcal{S} \cap \mathcal{PZ}_\alpha) \subset \mathcal{PZ}_\alpha \\ & \mathcal{W}_i x \leq b_i \\ & \dot{\eta}_\alpha = A_{cl} \eta_\alpha, \end{aligned}$$

where α are coefficients of Bézier polynomials that yield impact-invariant periodic orbits, $\mathcal{J}(\alpha)$ is a user-determined cost, $\Delta(\mathcal{S} \cap \mathcal{PZ}_\alpha) \subset \mathcal{PZ}_\alpha$ is the impact invariance condition, $\mathcal{W}_i x \leq b_i$ are other physical constraints, and $\dot{\eta}_\alpha = A_{cl} \eta_\alpha$ is the output dynamics condition. For more details on these constraints, refer to [28].

The cost function $\mathcal{J}(\alpha)$ largely influences the behavior of the walking gaits that it generates; however, the user's cost function \mathcal{J}_{human} underlying her preferences is poorly-understood. This section aims to describe the relationship between gaits and user preferences through the underlying cost function \mathcal{J}_{human} , so that future gait synthesis can be streamlined towards user-preferred walking. Thus, we aim to identify key terms in \mathcal{J}_{human} that numerically account for the preferences captured by LINECOSPAR.

All walking gaits on the exoskeleton are flat-footed. Thus, by analyzing the center of mass (CoM) and center of pressure (CoP), we can treat the patient-exoskeleton system as a Linear Inverted Pendulum Model (LIPM). This allows us to analyze the underlying utility function \mathcal{J}_{human} using the cost structure from [29]. We first introduce Zero Moment Point (ZMP) and LIPM, and then discuss correspondences between metrics of dynamic stability and user comfort.

Zero Moment Point. The Zero Moment Point (ZMP) is a widely-used notion of stability for bipedal robots that is defined as *the point on the ground at which the net moment of the inertial forces and the gravity forces has no component along the horizontal axes* [14]. When the ZMP exists outside of the "support polygon," i.e. the convex hull of the stance foot (or stance feet in the double-support domain), the robot experiences foot roll.

Static and Dynamic Stability. For a full discussion, refer to pg. 7 of [30]. In general, static stability is the condition in which the CoM and CoP never leave the support polygon. In contrast, quasi-static stability relaxes this condition on the CoM and only requires that the CoP remains inside the

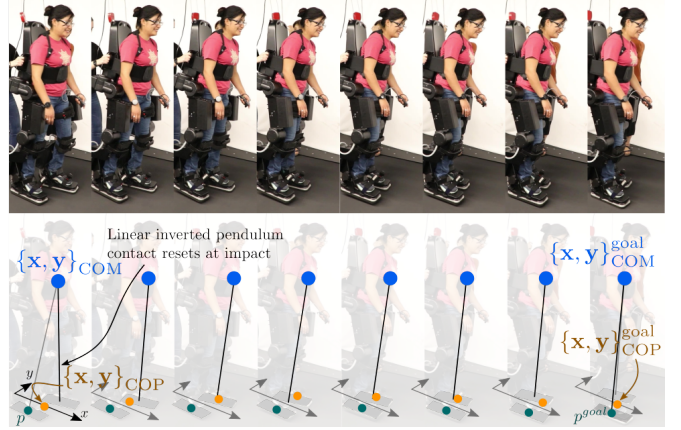


Fig. 8. Illustration of a single step with the overlaid LIPM model.

support polygon. For dynamic stability, the CoP lies on the boundary of the support polygon for a portion of the gait.

Linear Inverted Pendulum Model (LIPM). The LIPM is a low-dimensional dynamical system for reduced-order gait generation. The LIPM model assumes constant height of the center of mass, as well as zero angular momentum. The dynamics of the LIPM [31] are:

$$\begin{aligned} m\ddot{x}_{CoM} &= \frac{mg}{z_0}(x_{CoM} - x_{CoP}), \\ m\ddot{y}_{CoM} &= \frac{mg}{z_0}(y_{CoM} - y_{CoP}), \end{aligned}$$

where $\{x, y\}_{CoM}$ are the x and y positions of the CoM at constant height z_0 , and $\{x, y\}_{CoP}$ denote the x and y positions of the CoP. For planar horizontal ground walking, the ZMP is mathematically equivalent to the CoP. The CoP was experimentally obtained using the four 3-axis force sensors on the bottom of the exoskeleton's feet.

Fitting the LIPM Cost Function to User Preferences. Since flat-foot level-ground walking is well captured by the LIPM model, the cost function used in the LIPM to generate desirable walking behavior may explain the users' utility functions underlying their exoskeleton gait preferences. As defined in [29], the LIPM cost function is:

$$\begin{aligned} \mathcal{J}_{LIPM} = & w_1 \|x_{CoM}^{goal} - x_{CoM}\|^2 + w_2 \|\dot{x}_{CoM}\|^2 + \\ & w_3 \|\dot{x}_{CoP}\|^2 + w_4 \|p_x^{goal} - p_x\|^2 + \\ & w_1 \|y_{CoM}^{goal} - y_{CoM}\|^2 + w_2 \|\dot{y}_{CoM}\|^2 + \\ & w_3 \|\dot{y}_{CoP}\|^2 + w_4 \|p_y^{goal} - p_y\|^2, \end{aligned}$$

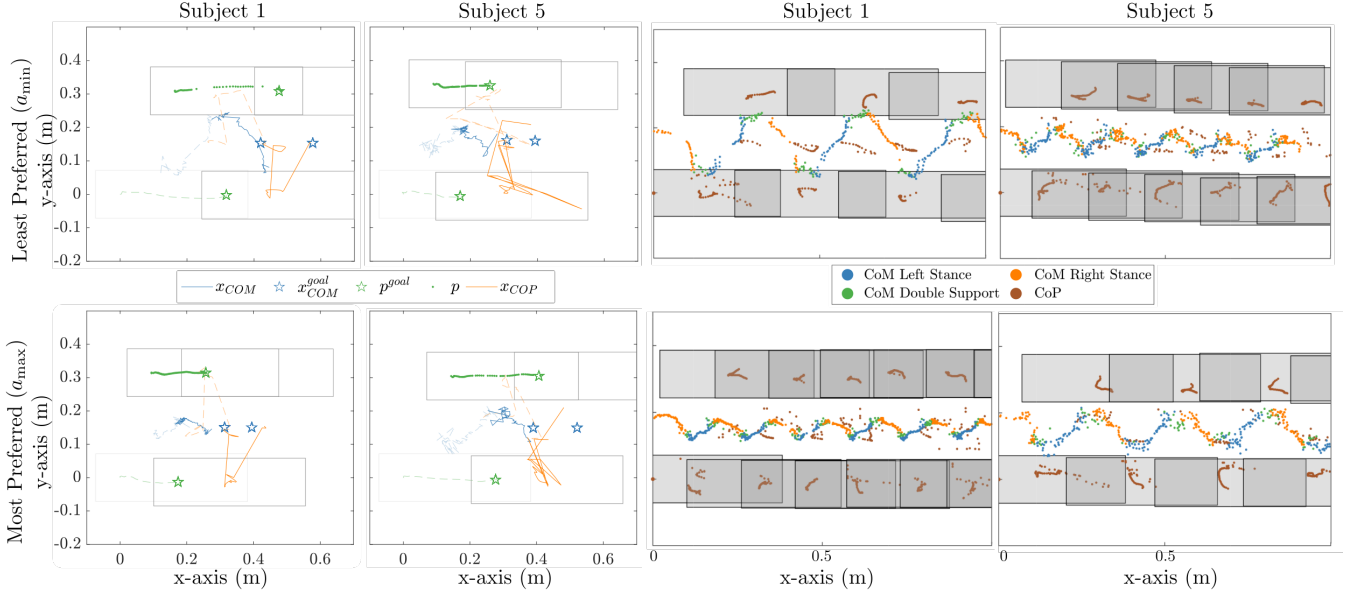


Fig. 9. **Comparison of Preferences.** This figure illustrates the trade-off between more and less dynamically-stable gaits as well as the contrasting preferences among different subjects. While all of the exoskeleton gaits are dynamically stable, both the least preferred gait (a_{\min}) of subject 1 and the most preferred gait (a_{\max}) of subject 5 exhibit behavior closer to statically-stable gaits. Subject 1 preferred dynamic gaits with a large difference between x_{CoP} and x_{CoM} ; in contrast, subject 5 preferred gaits in which x_{CoP} closely followed the center of mass. Rectangles represent the exoskeleton's feet.

where $\{x, y\}_{\text{CoM}}^{\text{goal}}$ denotes the CoM goal position in the x and the y directions, $\{\dot{x}, \dot{y}\}_{\text{CoP}}$ denotes the velocity of the CoP in the x and y directions, $\{\dot{x}, \dot{y}\}_{\text{CoM}}$ is the velocity of the CoM, $p_{\{x,y\}}^{\text{goal}}$ denotes the next stance foot position in the x and y directions, and $p_{\{x,y\}}$ denotes the x and y positions of the swing foot (Fig. 8).

We hypothesize that $\mathcal{J}_{\text{human}}(w)$ can be captured as a function of the weights $w := \{w_i\}, i \in \{1, \dots, 4\}$. Therefore, we fit the weights w of $\mathcal{J}_{\text{LIPM}}$ to the validation-stage preference data, i.e., the preferences between the most-preferred gaits (gaits with parameters \mathbf{a}_{\max}) and each of the random gaits presented during the validation phase². The weights w were optimized via the quadratic program:

$$w^* = \underset{w}{\text{argmin}} \quad \|w\|$$

$$\text{s.t.} \quad \begin{bmatrix} \delta_1^{(1)} & \delta_2^{(1)} & \delta_3^{(1)} & \delta_4^{(1)} \\ \vdots & \vdots & \vdots & \vdots \\ \delta_1^{(n)} & \delta_2^{(n)} & \delta_3^{(n)} & \delta_4^{(n)} \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \end{bmatrix} < 0,$$

where n denotes the number of pairwise preferences, and:

$$\delta_i = \left(\|x_{(i,x)}^{\text{pref}}\|^2 + \|x_{(i,y)}^{\text{pref}}\|^2 \right) - \left(\|x_{(i,x)}^{\text{not pref}}\|^2 + \|x_{(i,y)}^{\text{not pref}}\|^2 \right)$$

$$\begin{aligned} x_{(1,x)} &= x_{\text{CoM}}^{\text{goal}} - x_{\text{CoM}} & x_{(1,y)} &= y_{\text{CoM}}^{\text{goal}} - y_{\text{CoM}} \\ x_{(2,x)} &= \dot{x}_{\text{CoM}} & x_{(2,y)} &= \dot{y}_{\text{CoM}} \\ x_{(3,x)} &= \dot{x}_{\text{CoP}} & x_{(3,y)} &= \dot{y}_{\text{CoP}} \\ x_{(4,x)} &= p_x^{\text{goal}} - p_x & x_{(4,y)} &= p_y^{\text{goal}} - p_y. \end{aligned}$$

We use subject-wise holdout (leave-one-out) cross-validation across the subjects to verify the reliability of the fit. The average weights across all six holdout fits are: $w_1 = -0.1266$, $w_2 = 0.1363$, $w_3 = -0.0944$, and $w_4 = 1.0662$.

²Cost function fitting and CoP/CoM plotting code can be found at: <https://github.com/myracheng/linecospar/tree/master/gaitAnalysis>

We quantify the predictive power of each fitted cost function on the users' utility functions using the rank consistency between the cost function values and the preference data. Table II shows the predictive power of $\mathcal{J}_{\text{LIPM}}$ on the preferences, as well as the predictive power of two other cost functions, $\mathcal{J}_{\text{static}}$ and $\mathcal{J}_{\text{dynamic}}$, respectively defined as:

$$\mathcal{J}_{\text{static}} = \|\{x, y\}_{\text{CoM}} - \{x, y\}_{\text{CoP}}\|^2,$$

$$\mathcal{J}_{\text{dynamic}} = \|\{p_{\{x,y\}}^{\text{goal}} - p_{\{x,y\}}\|^2.$$

These two metrics are directly opposed: while $\mathcal{J}_{\text{dynamic}}$ is the term from $\mathcal{J}_{\text{LIPM}}$ that promotes dynamic stability, $\mathcal{J}_{\text{static}}$ penalizes dynamic stability in favor of static stability. This is because in the LIPM dynamics, the acceleration of $\{x, y\}_{\text{CoM}}$ approaches zero as $\mathcal{J}_{\text{static}}$ approaches zero. We find that $\mathcal{J}_{\text{LIPM}}$ and $\mathcal{J}_{\text{dynamic}}$ capture the preferences of five of the six subjects, while $\mathcal{J}_{\text{static}}$ completely predicts the preferences of the single outlier, subject 5.

Fig. 9 further illustrates this difference. The largest discrepancy between $\mathcal{J}_{\text{dynamic}}$ and $\mathcal{J}_{\text{static}}$ is that of subject 1 and subject 5. The preferences of subject 1 align with dynamic stability, while the preferences of subject 5 align with static stability. The diametric opposition between the cost function terms predicting these users' preferences reflects inconsistencies across users' gait utility functions. This suggests that there is most likely no single metric that entirely captures all users' underlying utilities. Thus, it is important to generate a variety of gaits that satisfy the cost functions reflecting different users' preferences.

V. CONCLUSION

This work presents two main contributions: 1) the LINECOSPAR algorithm to efficiently learn personalized, user-preferred gaits in high dimensions, and 2) an approach

TABLE II
PREDICTIVE POWER OF COST FUNCTIONS ON USER PREFERENCES

Cost Function	Correctly predicted preferences per subject (%)					
	1	2	3	4	5	6
\mathcal{J}_{LIPM} (holdout)	75	100	62.5	75	12.5	87.5
\mathcal{J}_{LIPM}	75	87.5	62.5	75	62.5	100
$\mathcal{J}_{dynamic}$	100	100	50	75	12.5	37.5
\mathcal{J}_{static}	50	75	37.5	50	100	75

for understanding the mechanisms dictating individual users' gait preferences.

LINECOSPAR identifies preferred actions in high dimensions, both in simulation and in experiments with six able-bodied subjects using the Atalante lower-body exoskeleton. We then examine the experimentally-obtained gait preferences to gain insight into the utility functions underlying users' gait preferences. We identify opposing measures of dynamicity that have predictive power for different users' preferences, implying that each user consistently prefers walking gaits that are either more dynamically or statically stable. These considerations may inform the synthesis of new exoskeleton gaits that maximize user comfort.

Future steps include conducting studies involving subjects with paraplegia, whose preferences likely differ from those of able-bodied subjects. As user preferences may change over time, creating a learning framework that accounts for these adaptations is also an important future research direction.

LINECOSPAR's high-dimensional learning capabilities provide insight into exoskeleton walking gaits that maximize user comfort, paving the way for generating new gaits beyond the gait library. This presents promising advancements for clinical trials and the broader rehabilitation community.

ACKNOWLEDGMENT

The authors would like to acknowledge the subjects who participated in exoskeleton testing, as well as the entire Wandercraft team that designed Atalante and continues to provide technical support for this project.

REFERENCES

- [1] A. Bajcsy, D. P. Losey, M. K. O'Malley, and A. D. Dragan, "Learning robot objectives from physical human interaction," in *Conf. on Robot Learning*, 2017, pp. 217–226.
- [2] S. Christen, S. Stevšić, and O. Hilliges, "Guided deep reinforcement learning of control policies for dexterous human-robot interaction," in *Int. Conf. on Robotics and Automation*. IEEE, 2019, pp. 2161–2167.
- [3] S. Cremer, S. K. Das, I. B. Wijayasinghe, D. O. Popa, and F. L. Lewis, "Model-free online neuroadaptive controller with intent estimation for physical human-robot interaction," *IEEE Trans. on Robotics*, 2019.
- [4] M. Tucker, E. Novoseller, C. Kann, Y. Sui, Y. Yue, J. Burdick, and A. D. Ames, "Preference-based learning for exoskeleton gait optimization," *arXiv preprint arXiv:1909.12316*, 2019.
- [5] J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, and S. H. Collins, "Human-in-the-loop optimization of exoskeleton assistance during walking," *Science*, vol. 356, no. 6344, pp. 1280–1284, 2017.
- [6] N. Thatte, H. Duan, and H. Geyer, "A method for online optimization of lower limb assistive devices with high dimensional parameter spaces," in *Int. Conf. on Robotics and Automation*. IEEE, 2018, pp. 1–6.
- [7] A. Agrawal, O. Harib, A. Hereid, S. Finet, M. Masselin, L. Praly, A. D. Ames, K. Sreenath, and J. W. Grizzle, "First steps towards translating hzd control of bipedal robots to decentralized control of exoskeletons," *IEEE Access*, vol. 5, pp. 9919–9934, 2017.
- [8] O. Harib, A. Hereid, A. Agrawal, T. Gurriet, S. Finet, G. Boeris, A. Duburcq, M. E. Mungai, M. Masselin, A. D. Ames, *et al.*, "Feedback control of an exoskeleton for paraplegics: Toward robustly stable, hands-free dynamic walking," *IEEE Control Systems Magazine*, vol. 38, no. 6, pp. 61–87, 2018.
- [9] T. Gurriet, S. Finet, G. Boeris, A. Duburcq, A. Hereid, O. Harib, M. Masselin, J. Grizzle, and A. D. Ames, "Towards restoring locomotion for paraplegics: Realizing dynamically stable walking on exoskeletons," in *Int. Conf. on Robotics and Automation*. IEEE, 2018, pp. 2804–2811.
- [10] C. Basu, Q. Yang, D. Hungerman, M. Sinahal, and A. D. Dragan, "Do you want your autonomous car to drive like you?" in *Int. Conf. on Human-Robot Interaction*. IEEE, 2017, pp. 417–425.
- [11] T. Joachims, L. A. Granka, B. Pan, H. Hembrooke, and G. Gay, "Accurately interpreting clickthrough data as implicit feedback," in *SIGIR*, vol. 5, 2005, pp. 154–161.
- [12] O. Chapelle, T. Joachims, F. Radlinski, and Y. Yue, "Large-scale validation and analysis of interleaved search evaluation," *ACM Transactions on Information Systems*, vol. 30, no. 1, p. 6, 2012.
- [13] J. Kirschner, M. Mutny, N. Hiller, R. Ischebeck, and A. Krause, "Adaptive and safe Bayesian optimization in high dimensions via one-dimensional subspaces," in *Int. Conf. on Machine Learning*, 2019, pp. 3429–3438.
- [14] M. Vukobratović and B. Borovac, "Zero-moment point—thirty five years of its life," *Int. journal of humanoid robotics*, vol. 1, no. 01, pp. 157–173, 2004.
- [15] Y. Sui, V. Zhuang, J. W. Burdick, and Y. Yue, "Multi-dueling bandits with dependent arms," in *Conf. on Uncertainty in Artificial Intelligence*, 2017.
- [16] Y. Sui, M. Zoghi, K. Hofmann, and Y. Yue, "Advancements in dueling bandits," in *IJCAI*, 2018, pp. 5502–5510.
- [17] Y. Yue, J. Broder, R. Kleinberg, and T. Joachims, "The k-armed dueling bandits problem," *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1538–1556, 2012.
- [18] P. Shivaswamy and T. Joachims, "Online structured prediction via coactive learning," in *Int. Conf. on Machine Learning*. Omnipress, 2012, pp. 59–66.
- [19] —, "Coactive learning," *Journal of Artificial Intelligence Research*, vol. 53, pp. 1–40, 2015.
- [20] C. Li, S. Gupta, S. Rana, V. Nguyen, S. Venkatesh, and A. Shilton, "High dimensional Bayesian optimization using dropout," *Int. Joint Conf. on Artificial Intelligence*, 2017.
- [21] Z. Wang, F. Hutter, M. Zoghi, D. Matheson, and N. D. Freitas, "Bayesian optimization in a billion dimensions via random embeddings," *J. of Artificial Intelligence Research*, vol. 55, p. 361–387, 2016.
- [22] W. Chu and Z. Ghahramani, "Preference learning with Gaussian processes," in *Int. Conf. on Machine Learning*. ACM, 2005, pp. 137–144.
- [23] U. Ezeafulukwe, M. Darus, and O. Fadipe-Joseph, "On analytic properties of a sigmoid function," *Int. Journal of Mathematics and Computer Science*, vol. 13, no. 2, pp. 171–178, 1 2018.
- [24] C. Wirth, R. Akrou, G. Neumann, and J. Fürnkranz, "A survey of preference-based reinforcement learning methods," *Journal of Machine Learning Research*, vol. 18, no. 136, pp. 1–46, 2017.
- [25] "Video of the experimental results." <https://youtu.be/c6a0kXMyML0>.
- [26] A. Duburcq, Y. Chevalere, N. Bredech, and G. Boeris, "Online trajectory planning through combined trajectory optimization and function approximation: Application to the exoskeleton Atalante," *arXiv preprint arXiv:1910.00514*, 2019.
- [27] T. Gurriet, M. Tucker, A. Duburcq, G. Boeris, and A. D. Ames, "Towards variable assistance for lower body exoskeletons," *IEEE Robotics and Automation Letters*, vol. 5, no. 1, pp. 266–273, 2019.
- [28] A. Ames, "Human-inspired control of bipedal walking robots," *IEEE Trans. on Automatic Control*, vol. 59, no. 5, pp. 1115–1130, 2014.
- [29] B. J. Stephens and C. G. Atkeson, "Push recovery by stepping for humanoid robots with force controlled joints," in *Int. Conf. on humanoid robots*. IEEE, 2010, pp. 52–59.
- [30] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. H. Choi, and B. Morris, *Feedback control of dynamic bipedal robot locomotion*. CRC press, 2018.
- [31] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," in *2003 IEEE Int. Conf. on Robotics and Automation (Cat. No. 03CH37422)*, vol. 2. IEEE, 2003, pp. 1620–1626.