



Demystifying Reward Design in Reinforcement Learning for Upper Extremity Interaction: Practical Guidelines for Biomechanical Simulations in HCI

Hannah Selder

Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, Leipzig University
Leipzig, Germany
hannah.selder@uni-leipzig.de

Florian Fischer

Department of Engineering
University of Cambridge
Cambridge, United Kingdom
fjf33@cam.ac.uk

Per Ola Kristensson

Department of Engineering
University of Cambridge
Cambridge, United Kingdom
pok21@cam.ac.uk

Arthur Fleig

Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI) Dresden/Leipzig, Leipzig University
Leipzig, Germany
arthur.fleig@uni-leipzig.de




Figure 1: Training musculoskeletal user models via Reinforcement Learning can be a daunting task. We analyze reward function design to overcome typical failure modes, such as erratic movement without pressing the correct button in a choice reaction task (A) or moving towards the target button without pressing it (B). We distill guidelines on how to arrive at successful policies in the choice reaction task (C), bi-manual typing (D), and other interaction tasks (pointing, tracking, remote control).

Abstract

Designing effective reward functions is critical for reinforcement learning-based biomechanical simulations, yet HCI researchers and practitioners often waste (computation) time with unintuitive trial-and-error tuning. This paper demystifies reward function design by systematically analyzing the impact of effort minimization, task completion bonuses, and target proximity incentives on typical HCI tasks such as pointing, tracking, and choice reaction. We show that proximity incentives are essential for guiding movement, while completion bonuses ensure task success. Effort terms, though optional, help refine motion regularity when appropriately scaled. We perform an extensive analysis of how sensitive task success and completion time depend on the weights of these three reward components. From these results we derive practical guidelines to

create plausible biomechanical simulations without the need for reinforcement learning expertise, which we then validate on remote control and keyboard typing tasks. This paper advances simulation-based interaction design and evaluation in HCI by improving the efficiency and applicability of biomechanical user modeling for real-world interface development.

CCS Concepts

• **Human-centered computing** → **User models.**

Keywords

Reward design, biomechanical models, deep reinforcement learning

ACM Reference Format:

Hannah Selder, Florian Fischer, Per Ola Kristensson, and Arthur Fleig. 2025. Demystifying Reward Design in Reinforcement Learning for Upper Extremity Interaction: Practical Guidelines for Biomechanical Simulations in HCI. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*, September 28–October 01, 2025, Busan, Republic of Korea. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3746059.3747779>



This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International License.

UIST '25, Busan, Republic of Korea

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2037-6/25/09

<https://doi.org/10.1145/3746059.3747779>

1 Introduction

Reinforcement Learning (RL) has emerged as a transformative force across a wide range of domains, including robotics [22, 29], gaming [68], and autonomous systems [75]. In Human-Computer Interaction (HCI), RL is increasingly used to drive biomechanical user simulations that model plausible human behavior during interaction [17, 24, 28], opening the door to a dynamic, embodied understanding of movement-based tasks and virtual prototyping. Like moths to a flame, researchers are drawn to these powerful tools – only to discover that the flame might burn them too: polished results often come at the cost of extensive, unintuitive trial-and-error.

At the heart of RL lies the *reward function*, a mathematical description of the agent’s goals, such as typing a word on a keyboard (Figure 1). This function encodes what it means to act successfully and plausibly: For biomechanical simulations, did the agent, i.e., the biomechanical user model, reach the correct target? Was the motion biomechanically plausible? While prior works have shown that RL-trained agents can simulate plausible user behavior in HCI [17, 28], they rarely reveal how much time and effort went into tuning these reward functions. For state-of-the-art musculoskeletal models, each design tweak can take 12–72 hours on modern workstations to train – potentially resulting in subtle, non-intuitive differences. A lack of general guidance on reward function design, limited understanding of how different components affect outcomes, and high computational costs of unstructured trial-and-error tuning all constitute a major entry barrier for HCI researchers without RL expertise.

Extending preliminary results [64], this paper aims to demystify reward function design and provide practical guidance that lowers barriers for newcomers while refining practice for experts. While individual reward terms may seem reasonable, combining them unsystematically can cause agents to exploit loopholes and fail the task. Band-aid fixes, like adding penalty terms for specific unwanted behaviors, tend to result in overly complex and “overfitted” rewards. This problem is especially pronounced in high-dimensional, over-actuated musculoskeletal models and black-box dynamics. To tackle this issue, we systematically analyze how reward components – target proximity, task completion bonuses, and effort minimization – impact learning outcomes across movement-based HCI tasks such as pointing, tracking, and choice reaction. We look at how sensitive these outcomes are to changes in the weights of the individual reward components. In total, we trained and tested over 500 policies – so other researchers don’t have to. From this, we derive practical, sequential guidelines to accelerate and simplify biomechanical simulation design. For instance, finding that effort terms are not always necessary to generate successful movement trajectories, we suggest to omit them unless movement instabilities such as trembling arise. We validate these guidelines on two more advanced HCI tasks that do not only involve aimed arm movements, but also require learning to control a physical input device (joystick and keyboard).

While our goal is to support more principled and accessible reward design, some limitations are inherent to our scope. Our analysis focuses on basic visuomotor interaction tasks that do not require much creative or strategic thinking. Tasks that involve highly ambiguous goal structures or leave room for a diversity of

creative strategies may require complementary methods such as hierarchical RL and more sophisticated cognitive modeling. Similarly, our evaluation is currently grounded in arm-based interaction.

In summary, we address the following **research questions**:

- RQ.1 (Plausibility)** What combinations and relative weightings of reward function components (i.e., proximity, effort, and task completion bonuses) produce plausible human movement trajectories across a range of goal-directed visuomotor tasks (pointing, tracking, choice reaction)?
- RQ.2 (Sensitivity)** What are the sensitivities of interaction outcomes (e.g., completion time, success rate) to variations in individual reward components in these tasks?
- RQ.3 (Generalizability)** To what extent do the observed findings on the effects of reward components generalize across a set of goal-directed visuomotor tasks?

With **RQ.1**, we aim to support the reward function design process by identifying combinations that lead to plausible simulated behavior. With **RQ.2**, we investigate how sensitive key performance metrics are to individual reward components, helping researchers anticipate the effects of design changes. With **RQ.3**, we seek to uncover recurring effects and trade-offs across multiple tasks, enabling the formulation of general reward function design guidelines.

For these guidelines to be useful, we believe they should exhibit the following key qualities: they should be (i) **succinct** to reduce cognitive load; (ii) **clear** in what each point contributes; (iii) **sequential** to offer a logical flow that eases step-by-step implementation; (iv) **actionable** with concrete steps or heuristics; and (v) **justified** by empirical evidence or clear reasoning. These principles informed both our analysis and the way we present our findings.

We **contribute** to the computational modeling of movement-based interaction by:

- presenting the first systematic and empirical analysis of reward function components (proximity, effort, and task bonuses) across three canonical goal-directed HCI tasks (pointing, tracking, choice reaction);
- revealing key sensitivities and trade-offs that emerge during reward function tuning, thus deepening the understanding of RL behavior in HCI contexts;
- providing clear, actionable, and empirically grounded guidelines for reward function design that lower the entry barrier for newcomers and support experts.

To support reproducibility and facilitate further research, we release a dataset comprising 517 policies trained for the reward functions discussed in this paper [63]. These policies are intended for use within the User-in-the-Box framework [28]. The dataset can be found at <https://doi.org/10.5281/zenodo.15845404>.

2 Related Work

Our goal in this section is twofold: (i) to contextualize our contribution within the growing body of work on RL-based biomechanical simulations, and (ii) to offer an entry point for HCI researchers, especially those without prior RL expertise, who are interested in leveraging these methods for interaction design. To this end, we survey how biomechanical models have been applied in HCI to simulate, prototype, and evaluate movement-based interactions in Section 2.1. We then discuss the role of reward function design in

RL and its implications for plausible agent behavior in Section 2.2, before summarizing the identified research gaps in Section 2.3. While extensive literature has developed guidelines and workflows supporting different aspects of developing and testing user simulation, including computational models of human behavior and rationality [9, 77], empirical design decisions [56], and integrating behavioral data into model formulation [72], this work addresses practical aspects in RL-based user modeling, namely the development and iterative refinement of reward functions.

2.1 Biomechanical Simulations in HCI

Simulation-based approaches are gaining increasing attention in HCI research, particularly for modeling user behavior during interaction. Recent work highlights the relevance of RL-based user modeling, including simulations of touchscreen typing [65] and adaptive interfaces supported by learned user models [35]. In parallel, biomechanical simulations have evolved as a beneficial tool for developing and validating HCI technologies [3, 45]. They are fundamentally based on kinematic and dynamic models, which, depending on their complexity, provide information on the skeletal structure, inertial properties, and the neuromuscular system of the human body [36]. While early models were limited to calculating mechanical loads in static postures [2, 20, 73], advancements have led to physiologically more accurate musculoskeletal models [14, 15, 25, 36, 41], which can be used to predict ergonomics, fatigue, and user strategies in addition to traditional performance metrics [19]. Biomechanical simulations therefore provide valuable insights into predictability, safety, and accessibility, making them a promising approach for system evaluation prior to and as a complement to user testing [45].

Traditionally, biomechanical models integrated into physics engines such as OpenSim [15] have been used for *inverse simulation*. Inverse methods aim to infer dynamic properties and motor patterns from reference motion data, enabling in-depth analysis of the physiological state and the development of assistive tools, for example for rehabilitation [32, 58]. In the context of virtual reality, Hwang et al. apply simulations with electrical muscle stimulation to create kinesthetic force feedback for the lower body [27].

A second and more recent branch of research uses *forward simulation* to predict human movements and ergonomic states from muscle control signals [1, 39]. These methods require defining a controller that actively selects neural muscle signals based on the sensed body state and the given interaction task. The framework of optimal feedback control has emerged as the de facto standard to model how these muscle signals are selected in aimed movements [18, 70]. Optimal feedback control is based on the assumption that information on the controlled quantity (e.g., the arm, hand, and controller in VR interaction) is continuously fed back to the controller, resulting in a closed interaction loop between the user and their environment. Optimal feedback control methods such as LQG [18], MPC [31], and intermittent control [42] have been successfully applied to predict movement trajectories in standard interaction tasks such as mouse pointing and mid-air pointing. However, they impose severe restrictions on the complexity of the biomechanical models and tasks considered. For example, simplified activation models based on torque actuators directly attached

to the joints have been introduced [24, 31] to overcome the complexities arising from the highly over-actuated, noisy, and delayed neuromuscular system [37].

In recent years, deep RL has emerged as the go-to method for forward simulations of movement-based interaction [8, 66]. Building on the theory of optimal feedback control, deep RL methods approximate the optimal control policy through efficient exploration and learning from interaction with an unknown environment [67]. In contrast to traditional methods, RL methods do not require a closed formulation of the system to be controlled, which makes them particularly useful for non-linear musculoskeletal systems and black-box physics engines such as OpenSim [15] or MuJoCo [69]. For example, Fischer et al. use SAC, an entropy-based RL method, to learn controlling the muscles of a state-of-the-art shoulder model for a mid-air pointing task. Hetzel et al. apply deep RL to simulate arm movement in mid-air typing.

The combination of motor and sensory control models, as discussed in [48], has been particularly relevant for modeling human-computer interaction. Ikkala et al. present *User-in-the-Box*, an RL-based simulation framework to generate task-specific movement trajectories based on the user's visual and proprioceptive perception of the interaction environment [28]. Moon et al. use movement data from biomechanical simulations to infer aimed target positions via neural density estimation [44]. Fischer et al. develop a system for running biomechanical models directly in VR applications, ensuring that the simulated user "perceives and controls" exactly the same environment as a real user [19]. With recent improvements to training efficiency, the scope of RL-based simulations has been extended to dexterous manipulation and grasping tasks [4, 7, 13, 62].

2.2 Reward Function Design for RL-Based Simulations

In RL-based biomechanical simulations, the design of the reward function is identified as a key factor in the effectiveness of the learning process [34]. This is because most RL methods (with a few noteworthy exceptions [57, 78]) do not learn from human reference data but solely from a predefined reward function. While different formulations of effective reward functions have been examined [23, 28, 34, 55], most of them have been handcrafted for a specific task [7, 8, 28], which limits their generalizability across tasks and contexts. In this process, even if a "working" reward function is found, it is often unclear to what extent it accurately encodes the preferences of a human stakeholder [46]. In addition, there is a lack of benchmark tasks and environments that would enable a fair comparison between proposed reward functions. Instead, the common practice is to identify and validate a reward function based on trial-and-error within a unique, task-specific setting that contains a large number of potentially confounding variables (and is often difficult to replicate). These confounding factors may include the biomechanical model used, different types of control and exploration noise [10], as well as factors modulating the training process such as early termination constraints [28, 33], predefined or adaptive learning curricula [17, 19, 24], and even the RL algorithm used [6].

In addition, movement-based interaction usually involves a trade-off between two or more opposing objectives, e.g., between accuracy and stability [38] or speed and accuracy [47]. While for many practical tasks it is natural to only define *sparse rewards* (e.g., providing a single positive reward if and only if the desired target position or velocity has been reached [10, 17], or when the game score increases [19, 43]), this sparsity of information leads to very slow training and often hinders convergence. As a consequence, RL literature has explored adding continuous *dense* reward terms that provide information about the utility of a state-action pair much earlier during an episode (e.g., moving closer towards vs. moving away from the desired target position). In particular, the theory of *reward shaping* has been developed [53], and bi-level approaches have been proposed to increase the robustness of RL training while adhering to the goals of the original sparse reward formulation [16, 26]. However, these approaches are rarely used in practice, probably because of their complexity and the lack of toolkits implementing them. Instead, ad-hoc approaches based on trial-and-error are predominate [6, 55]. Since 'failed' reward functions explored along the way are rarely reported, little insight is gained into which reward patterns are critical for successful RL training, essentially leaving all RL designers with the same difficulties. In this work, we propose an intermediate approach that systematically adds task-relevant 'guiding' terms in order to shape the reward function without introducing systematic bias, although the latter can only be experimentally validated.

The complexity of designing an appropriate reward function is further exacerbated in biomechanical simulations, since humans can perform tasks with an infinite number of different admissible joint trajectories and muscle activations [5] ("movement redundancy" [51]). To limit the possibilities to "sensible" movements, at least one *effort* term is typically added for regularization, e.g., to restrict the use of rapid and abrupt arm movements by penalizing large muscle activations, thus ensuring that available resources are used efficiently. Several effort cost models, including jerk [20], energy consumption [52], and commanded torque change [49], have been proposed and investigated from a motor control perspective [5, 10, 11, 71]. The role of an effort model in motor adaptation has been explored in empirical studies in [59, 76]. While there is consensus in the literature that effort reduction is likely to play a key role in movement planning, it is unclear how this incentive can be reconciled with other task- or performance-related goals [37].

2.3 Identified Research Gaps

To the best of our knowledge, there exist no guidelines on how to design and balance the individual reward components in practical HCI tasks, especially in combination with complex musculoskeletal systems. From an engineering point of view, it is unclear to what extent the different effort models can help guide the RL training towards biomechanically plausible movements and improve convergence in RL-based simulations. We therefore anticipate a strong need to explore the design of reward functions for realistic use cases of biomechanical models. In this work, we take a decisive step towards this goal by starting with a number of goal-directed visuomotor tasks that are highly relevant for HCI and analyzing

the individual and combined effects of different reward function components on RL-based learning of interactive body movements.

3 Methodology

In this work, we systematically analyze how different reward components (which we introduce in Section 3.1) influence emergent user strategies in RL-based biomechanical simulations. Our focus lies on visuomotor interaction tasks involving direct input via the index finger, such as pointing, tracking, and typing. Section 3.2 outlines the full range of tasks we consider, while Section 3.3 details our experimental design and rationale.

3.1 Reward Components

In designing the reward function, we focus on three components:

- The *task bonus* is a term that incentivizes reaching states that are beneficial for successfully completing the given task. Often an episode terminates immediately after such a state is reached. We integrate the many possibilities in the function $f_{\text{bonus}}(\cdot)$, where (\cdot) is a placeholder for all relevant function arguments. Defining such a bonus only requires splitting the state space into a set of desirable "goal states" to be awarded (e.g., when a desired button is hit), and a set of neutral states, where no information regarding their desirability is provided.
- The *distance* term provides a proximity incentive by rewarding the agent more the closer they get to the target, expressed by the function $f_{\text{distance}}(\cdot)$.
- The *effort* model aims to incentivize controls that require low physiological effort. Depending on which measure of effort is used, this model can be quite versatile: Designers can choose, e.g., to penalize jerky movements, or reward movements that require lower energy. We encompass the possibilities in the function $f_{\text{effort}}(\cdot)$.

To evaluate the intricacies of how the individual components work independently and in conjunction, for each component we introduce respective weights $w_{\text{distance}}, w_{\text{effort}}, w_{\text{bonus}} \geq 0$. In total, the most generic reward function amounts to

$$r_t = w_{\text{bonus}} \cdot f_{\text{bonus}}(\cdot) - w_{\text{distance}} \cdot f_{\text{distance}}(\cdot) - w_{\text{effort}} \cdot f_{\text{effort}}(\cdot). \quad (1)$$

If we set $w_{\text{effort}} = 0$ or $w_{\text{distance}} = 0$, we speak of *zero effort* or *zero distance*, respectively.

The **task bonus** is formulated as follows:

$$f_{\text{bonus}}(\cdot) = \begin{cases} 1, & \text{if model reaches a goal state,} \\ 0 & \text{else.} \end{cases} \quad (\text{Bonus})$$

We investigate three different **distance** reward functions, each based on the distance between the index finger and the task-specific target position (either a point or a surface), *dist*, as measured by a MuJoCo distance sensor [69]. In every timestep, we can measure:

- (1) The (absolute) value of the MuJoCo distance sensor:

$$f_{\text{distance}}(\text{dist}) = \|\text{dist}\|_2 \quad (D_{\text{absolute}})$$

- (2) The squared distance, which has been successfully used in RL tasks [30, 54]:

$$f_{\text{distance}}(\text{dist}) = \|\text{dist}\|_2^2 \quad (D_{\text{squared}})$$

- (3) An exponential transformation of the distance, as used in [28]:

$$f_{\text{distance}}(\text{dist}) = \frac{1 - e^{-10 \cdot \|\text{dist}\|_2}}{10} \quad (D_{\text{exponential}})$$

We also compare different **effort** models. The first one, denoted as *EJK* in the following, was first presented in [10] to simulate realistic arm movements. It consists of three components, motivated by the observation that combining multiple effort terms can improve the plausibility of generated movements [5, 74]. The components of *EJK* penalize the mean value of the muscle stimulation commands (r_{energy}), the jerk, i.e., the change in joint acceleration (r_{jerk}), and the total work done by the shoulder and elbow (r_{work}) in terms of angular velocities and torques. These components are normalized and weighted by coefficients c_1 , c_2 , and c_3 , respectively, resulting in the following effort model:

$$f_{\text{effort}}(r_{\text{energy}}, r_{\text{jerk}}, r_{\text{work}}) = \frac{c_1 r_{\text{energy}} + c_2 r_{\text{jerk}} + c_3 r_{\text{work}}}{c_1 + c_2 + c_3} \quad (\text{EJK})$$

In addition, we consider the three effort models from [31] (*DC*, *CTC*, and *JAC*), where their suitability to predict mid-air pointing movements using a non-RL optimization method (MPC) was examined. All three models include a penalty for large muscle stimulation commands u , motivated by the fact that humans seek to minimize their control effort during movement [70]. In the following models, these muscle controls are penalized in the squared Euclidean norm (the r_{energy} component of the EJK model instead penalizes muscle controls using the l^1 -norm [10]). The *DC* effort model only consists of this penalty term, weighted by a coefficient c_1 :

$$f_{\text{effort}}(u) = c_1 \|u\|_2^2 \quad (\text{DC})$$

The *CTC* model adds a penalty on large changes in commanded torque τ , which is the torque at the joints that directly results from the controlled muscle activations. This term is motivated by a study from Nakano et al. [49], where the minimum commanded torque *change* criterion, i.e., the derivative of τ , was found to be the most effective in explaining the temporal characteristics of actual hand trajectories. The *CTC* model is formulated as follows:

$$f_{\text{effort}}(u, \dot{\tau}) = c_1 \|u\|_2^2 + c_2 \|\dot{\tau}\|_2^2 \quad (\text{CTC})$$

Similarly, the *JAC* model adds a penalty in the kinematic space, namely on the joint accelerations x_{qacc} , thus avoiding "jerky" movements. This effort model was introduced in [49] and later found to provide the most comprehensive explanation of mid-air pointing movements [31]. In contrast to the EJK effort model, this model penalizes the acceleration values themselves instead of their changes (i.e., the jerk). The resulting *JAC* model is defined as follows:

$$f_{\text{effort}}(u, x_{\text{qacc}}) = c_1 \|u\|_2^2 + c_2 \|x_{\text{qacc}}\|_2^2 \quad (\text{JAC})$$

3.2 Tasks

To develop practical guidelines, which we will propose in Section 4.3, we trained policies for the choice reaction, pointing, and tracking task implemented in the User-in-the-Box (UitB) framework¹. In each task, we use the default *MoBL Arms Model* [60] with 5 DoFs (three independent shoulder joints, elbow, wrist) and 26 muscles enabled, and provide visual, proprioceptive, and tactile

information as input to the agent. Each episode starts with the arm hanging down (see Figure 3 (A)).

To validate our guidelines, we define a separate set of movement-based interaction tasks, consisting of a remote control [28] and a keyboard typing task [24]. In the following, we describe each of the five tasks in more detail. Table 1 provides an overview of the HCI relevance and objectives for each task, as well as the associated challenges for our RL-based user simulations.

Choice Reaction. In the choice reaction task, the agent faces four colored buttons and a display presenting a stimulus, i.e., one of the four colors. The objective is to press the physical button that corresponds to the displayed color as quickly as possible, with a time limit of four seconds per trial. The agent reaches a goal state when it presses the correct button with sufficient force. This triggers a change of the display color to indicate the start of the next trial. Each episode incorporates 10 trials. Training was conducted for 35 million steps, as further training beyond this threshold did not yield significant improvements.

Pointing. In the pointing task, the agent must move its index finger to a virtual target sphere of varying radius, positioned in front of the user. A goal state is reached when the agent maintains its fingertip within the target for 500 milliseconds, thereby completing a trial. After either a successful trial or a timeout of four seconds without completion, a new target location is sampled. Each episode contains 10 trials. Each policy has been trained for 50 million steps, after initial tests suggested no further improvements past this.

Tracking. In the tracking task, the agent is required to continuously follow a virtual moving target sphere with its fingertip as precisely as possible. The aim is to point at the target and then keep the fingertip inside the moving sphere. At each timestep, the model is considered to be in a goal state if the fingertip remains inside the target. Each episode lasts ten seconds, regardless of the agent's success. Training was conducted for 50 million steps.

Remote Control. In the remote control driving task, the agent controls acceleration and braking of a toy car by tilting the left joystick of a physical gamepad along a single dimension (forward or backward). The objective is to park the car inside the marked area. An obvious goal state is reached when the car is positioned inside this area. Due to the sequential nature of the task, making contact with the joystick can be considered an intermediate goal state. Each episode runs for a fixed duration of ten seconds. Being a more involved task, we heuristically trained each policy for 200 million steps.

Keyboard Typing. For the keyboard typing task, the agent interacts with a physical keyboard, attempting to hit specified keys. Replicating the setup from [24], the agent does not use visual perception, but instead obtains the position of the current target key as part of the observation. A goal state is reached when the correct key is successfully pressed, which also terminates the episode. If the agent fails to press the key within three seconds, the episode ends without success. Training was conducted using a single-hand model, after which a mirrored second hand was introduced to enable bimanual typing. Each policy was trained for two million steps.

¹<https://github.com/User-in-the-Box/user-in-the-box>

Task	HCI relevance	Challenges	Objective
Mid-air pointing	Natural direct control method; fundamental component of VR/AR interaction; required for both selection and manipulation tasks	Learn muscle control of shoulder and arm from multi-sensory perception (visuomotor task); keep finger inside target (precise aimed reaching); handle a range of target positions and sizes (multi-goal task)	Finger inside target (dwell time)
Object tracking	Dynamic environments requiring fast visuomotor adaptation	Perform continuous error corrections based on perceptual input (closed loop task)	Finger close to target
Choice reaction	Prototype of a physical selection task; sequential button clicking	Visuomotor coordination; abstraction; handling contact forces	Press correct button (force constraints)
Remote control	Sequential task; indirect manipulation via joystick control	Complex control task of 6th order; requires end-to-end learning of multiple skills (mid-air pointing, joystick control, visuomotor coordination to steer the car position)	Park car inside target
Keyboard typing	Standard bimanual interaction task; sequential task involving physical selection	High precision; handling contact forces; large sequences of aimed movements; generalizability across buttons (multi-goal task)	Press correct buttons (touch)

Table 1: HCI relevance, challenges, and objectives of the five HCI tasks considered in this work.

3.3 Experimental Design

In total, we trained over 500 RL policies across the five tasks introduced in Section 3.2, with more than 450 policies dedicated to analyzing the effects of reward function components in the core tasks of pointing, tracking, and choice reaction. Throughout, our experimental design sought to balance depth and coverage with the high computational cost of training biomechanical agents. We recall that each training may require 12-72 hours to converge. To ensure robustness and allow for confidence intervals, we initially trained each reward configuration multiple times with seven different initial network parameters (weights and biases) for the choice reaction task. As preliminary results showed low variance in most settings, we subsequently reduced the number of repetitions, in particular for cases where we observed low performance: three repetitions for policies without task bonus, where we averaged results, and one repetition for policies without a distance term. For other tasks, we performed three repetitions except for policies without a distance term, where we limited ourselves to one repetition.

To investigate how individual reward components influence learning outcomes (RQ.1), we adopted an incremental strategy: we first trained agents using only the task bonus or distance term, then explored their combinations, and finally introduced the effort model to assess its interaction with the other components. Our choice for the respective effort coefficients c_i (see Section 3.1) was informed by related literature and refined through initial pilot training runs. For the sensitivity analysis (RQ.2), we varied the weight of one reward component while keeping the others fixed, across multiple configurations and tasks. As we observed little variance across different formulations of the distance term, our sensitivity analysis only included a single distance term (exponential distance) to avoid redundant computation. Based on early-stage convergence patterns and similarity in component behavior, we focus on the choice reaction task in our sensitivity analysis, and add observations from the tracking task where differences occurred. Finally, to evaluate how our findings generalize beyond the initial task set (RQ.3), we applied our reward design guidelines to the remaining two tasks: keyboard typing and remote control.

To evaluate the trained policies, we ran five episodes per policy. Each episode consisted of 10 trials for the choice reaction and pointing tasks, and 10 seconds of continuous motion for the tracking task. In each episode, the agent starts with the arm hanging down, and targets are randomly sampled, using the same seed for all policies to ensure comparability. We report the mean performance across episodes.

For the choice reaction and pointing tasks, we measured success rate, defined as the percentage of successful trials across the five episodes, and task completion time, which reflects the average time required to complete a trial. If a trial is not completed within four seconds, this timeout was recorded as the completion time. For the tracking task, we evaluated the average distance to target, calculated as the mean Euclidean distance between the fingertip and the surface of the moving target sphere, and the percentage of time the fingertip remained inside the target (measured at a frequency of 20Hz).

4 Results

This section presents the results of training policies with various reward function configurations, in the order of the research questions. We provide both qualitative and quantitative evaluations; the former is based on behavioral observations from evaluation videos, while the latter uses task-specific metrics as described in Section 3.3. Parameter values for all experimental conditions (referred to as IDs in the following) are listed in the Supplementary Material, and key outcomes are illustrated in the accompanying video figure.

4.1 Plausibility (RQ.1)

We present the results for the choice reaction, pointing, and tracking tasks, in that order.

4.1.1 Choice Reaction. Figure 2 shows the success rates and completion times (IDs: 1–35). With **task bonus** only (IDs: 31, 32), agents fail to learn generalized behavior – only the green button is successfully pressed most of the time (orange circle for the "Zero" effort

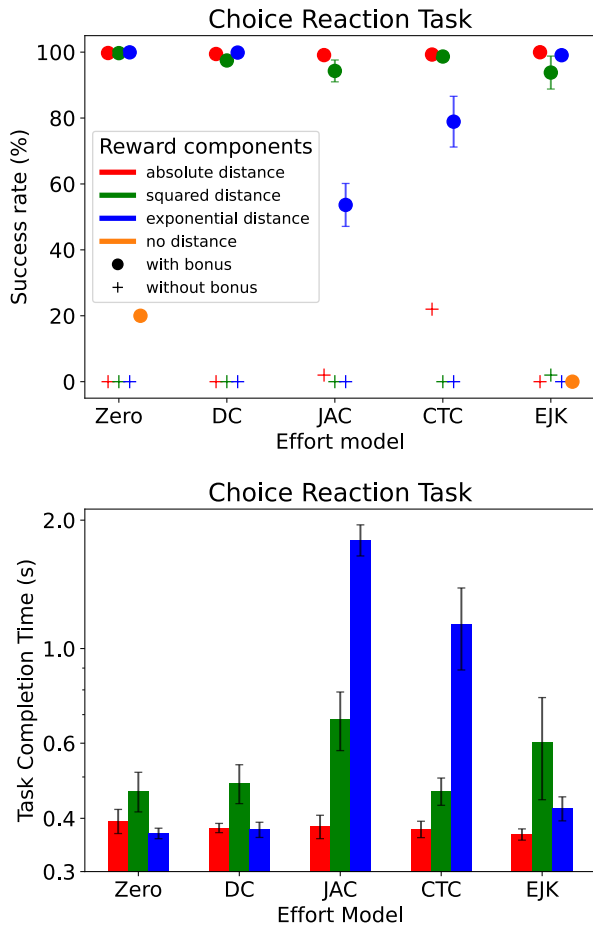


Figure 2: Success rates (top) and task completion times (bottom) for policies trained with various reward functions in the choice reaction task. Error bars indicate the standard error of the mean. Policies trained without either the task bonus (pluses) or distance term (orange circles) achieve success rates below 25%. Including both terms yields near-perfect success, except for JAC and CTC effort models. Among high-performing agents, those trained with a squared distance term complete the task considerably slower than with other distance terms.

model in Figure 2 (top)). Movements for the remaining three buttons are noisy and uncoordinated, even when increasing the bonus weight (ID: 32). With **distance term** only (pluses for the "Zero" effort model in Figure 2 (top)), while the agents learn to identify the correct button and move towards it, they fail to press buttons properly and resort to side contacts (IDs: 4–6, 10–12, 16–18, 22–24) as depicted in Figure 1(B). **Effort-only** rewards would lead to no movement, so were excluded.

Combining distance and task bonus results in successful movements (circles for the "Zero" effort model in Figure 2 (top)) as depicted in Figure 1(C), with one caveat: the choice of the distance function influences the qualitative behavior for the red button.

The squared and exponential distances (IDs: 25, 27) lead to the agent pressing the red button using the index finger's proximal phalanx, which often requires multiple attempts. This behavior is not observed with the absolute distance (ID: 26).

Combining distance and effort (pluses in Figure 2 (top)), agents mostly fail, either touching (but not pressing) buttons from the side or not reaching them (IDs: 4–6, 10–12, 16–18, 22–24). Agents trained with the JAC and CTC effort models only get close to the button with the absolute distance, whereas the exponential distance "works" for the DC and EJK models in that they touch the buttons from the side (IDs: 4, 22). With increased effort weights, the distinctions between the effort models become more apparent (IDs: 45, 53, 49). The penalization of joint accelerations in the JAC model causes the arm to remain extended, whereas the DC model bends the arm and lifts it towards the buttons, but stops below them (Figure 3 (C)).

A reward function **combining only effort and task bonus** (only tested for the EJK model; orange circle for EJK in Figure 2 (top)) fails to initiate movement, further emphasizing the need for a complementary, task-specific "guidance" term in the reward function (IDs: 33–35). In combination with the task-bonus-only results, we refrained from testing other effort models.

Finally, **combining all three components** (circles in Figure 2 (top)), we observe qualitative and quantitative differences between the distance terms. Across all effort models, absolute distance leads to the highest success rates ($\geq 96\%$). For JAC and CTC effort models, both absolute and squared distances outperform the exponential distance, with longer completion times for the squared distance (Figure 2 (bottom)). However, among the configurations with $\geq 90\%$ success rate, the exponential and absolute distance terms show comparable task completion times. Notably, we did not observe a faster task completion for agents trained without an effort model. Qualitatively, with the CTC effort, the agent is unable to press the upper two buttons and instead keeps the fingertip at one of the lower buttons (IDs: 13–15). With squared or exponential distance, we observe similar behavior for the DC effort model as without effort, where in both cases the red button is pressed with the back of the hand (IDs: 19, 21) – an effect that vanishes with increasing effort weight.

4.1.2 Pointing. Figure 4 presents success rates and completion times for agents trained on the **pointing** task (IDs: 117–150). The results are generally similar to those of the choice reaction task.

With **task bonus** only (orange circle for the "Zero" effort model in Figure 4 (top)), agents successfully point at targets close to the arm's initial position, but fail to reach other targets (IDs: 147, 148). Agents trained with **distance-only** (pluses for the "Zero" effort model in Figure 4 (top)) struggle to stay inside the sphere for the required dwell time, often hovering near the boundary (IDs: 144–146).

Combining distance and task bonus (circles for the "Zero" effort model in Figure 4 (top)) leads to stable positioning within the sphere ($\geq 90\%$ success rate), with reduced trembling (IDs: 141–143). In contrast to the choice reaction task, we do not observe any qualitative differences between the distance terms.

Combining distance and effort (pluses in Figure 4 (top)) yields agents that exhibit similar behavior as for distance-only. Squared distance combined with any effort model (IDs: 122, 128, 134, 141) and exponential distance combined with the JAC or CTC effort




Figure 3: Comparison of (failed) movement patterns of different effort models with (too) high weighting. (A) The CTC model does not move at all. (B) The JAC model raises the arm but does not bend it. (C) The DC model bends the arm but struggles to reach the button. (D) The EJK model reaches only the lower buttons.

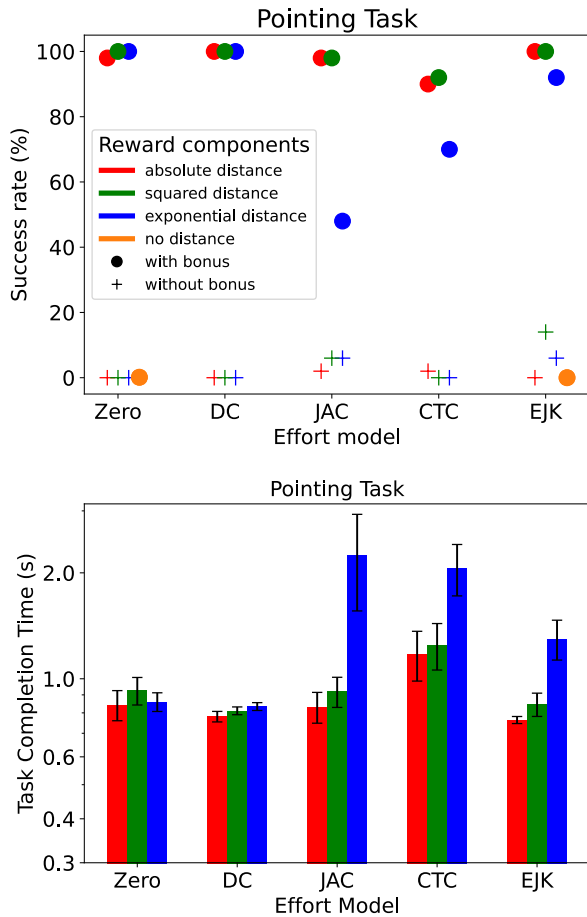


Figure 4: Success rates (top) and task completion time (bottom) for policies trained with various reward functions for the pointing task. Error bars indicate the standard error of the mean. Policies trained without either the task bonus (pluses) or distance term (orange circles) achieve success rates below 25%. Including both terms yields near-perfect success, except for JAC or CTC effort models.

model (IDs: 126, 132) result in agents that tend to stop short of the target. Other combinations hover over the boundary, not staying inside long enough (IDs: 120, 121, 127, 133, 140, 141). **Combining effort and task bonus** (orange circle for the EJK effort in Figure 4 (top)) results in no movement (IDs: 149, 150).

Finally, **combining all three components** leads to successful task performance, except for the combination of exponential distance with either EJK, JAC, or CTC effort models, where the agent sometimes keeps the arm outstretched, unable to reach inside the sphere (IDs: 117, 123, 129). Agents trained without an effort model are again not necessarily faster than those trained with effort model.

4.1.3 Tracking. Figure 5 illustrates the average target distance (lower is better; note the logarithmic scale) and the percentage of time spent within the target sphere (higher is better) for the tracking task (IDs: 151–187).

In contrast to the previous tasks, the agent trained with just the **task bonus** (orange circle for the "Zero" effort model in Figure 5) follows the target with high accuracy. A reward function with only the **distance term** (pluses for the "Zero" effort model in Figure 5) results in agents that stay close to the target but show a lot of trembling (IDs: 181, 182).

Agents trained with **distance and task bonus** (non-orange circles for the "Zero" effort model in Figure 5) remain in close proximity to the target sphere (IDs: 175–177). However, these agents occasionally lack the precision necessary to maintain their position within the target sphere. Notably, **combining distance and effort models** (pluses in Figure 5) results in agents mostly maintaining a closer distance to the target than those with an additional bonus (non-orange circles in Figure 5). The performance for **combining effort and task bonus** (orange circles in Figure 5) differs between effort models. While agents trained with the JAC effort model do not move at all (ID: 187), those trained with the DC, CTC or EJK effort models stay very close to the target sphere (ID: 184–186).

When **combining all three components**, we observed qualitative differences between the distance terms. Agents trained with the absolute distance tend to remain to the left of the sphere (IDs: 152, 158, 164, 170, 176), while those trained with the squared distance tend to remain to the right (IDs: 153, 159, 165, 171, 177). Agents trained with the exponential distance term exhibit a mixture of these behaviors (IDs: 151, 157, 163, 169, 175). In combination with

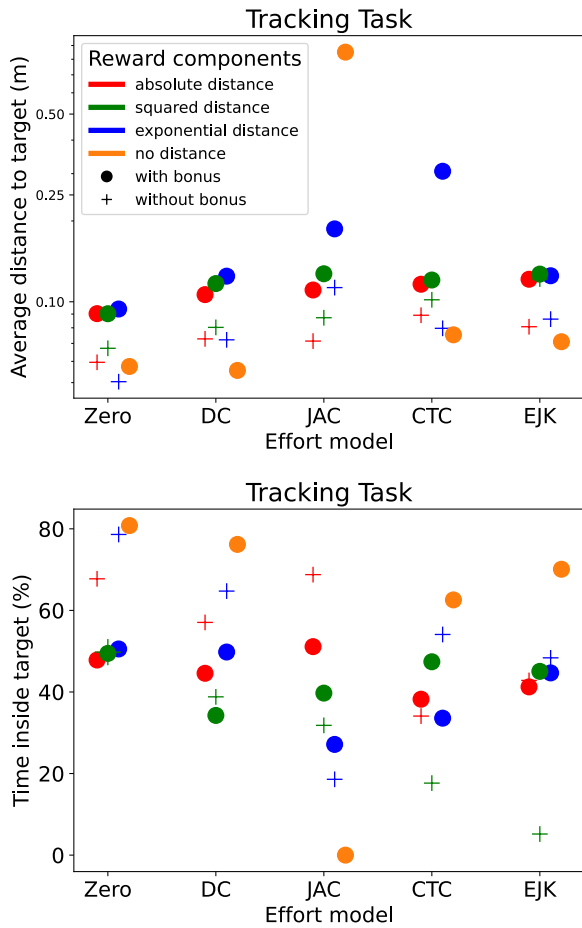


Figure 5: Performance comparison of policies trained with various reward functions in the tracking task. The top plot shows the average distance to the target (lower is better), while the bottom plot presents the time spent inside the target area (higher is better). The policy with task bonus and the JAC effort model fails to point inside the target at all, similar to the policy combining a task bonus, the CTC effort model, and exponential distance.

task bonus and either the JAC or the CTC effort model, the exponential distance leads to poorer tracking performance than the other distance terms.

4.2 Sensitivity (RQ.2)

We analyze the sensitivity of each reward component by varying one weight while keeping the others fixed. Based on the results from Section 4.1, we decided to focus on exponential distance as it showed the greatest variability between effort models, but otherwise performed similarly to absolute and squared distance. For the choice reaction task, we use $w_{\text{distance}} = 1$ (exponential), $w_{\text{effort}} = 1$, and $w_{\text{bonus}} = 8$ as our baseline. For the tracking task, we select $w_{\text{bonus}} = 0.5$ as baseline, as it constitutes a promising candidate for high success rates based on initial findings.

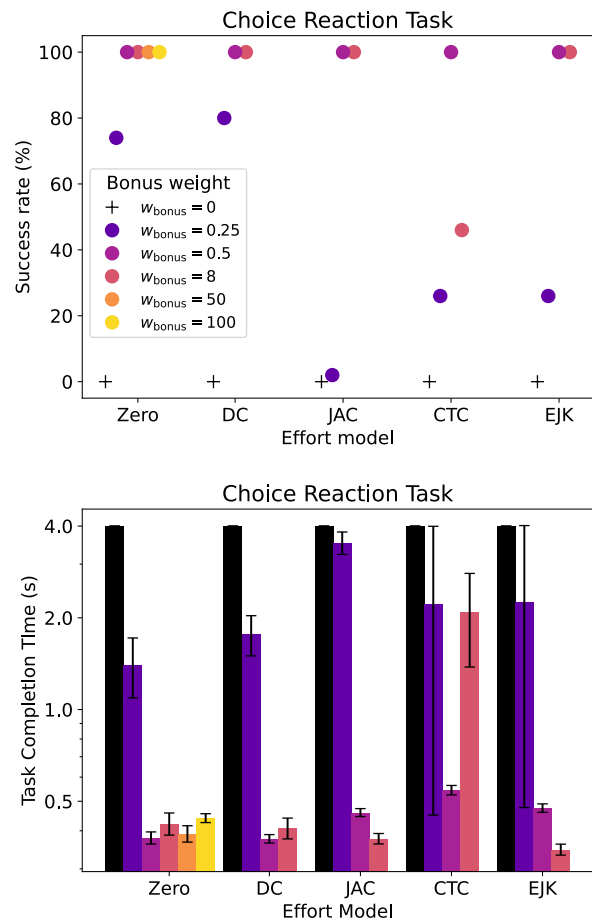


Figure 6: Effect of varying bonus weight on success rate (top) and task completion time (bottom) in the choice reaction task. The task completion time is presented for policies including task bonus and distance. Error bars correspond to the standard error of the mean. Policies trained with a bonus weight greater than or equal to 0.5 achieve success rates of 100%, except when using the CTC effort model. Lower bonus weights result in lower success rates.

4.2.1 Bonus Weight. Figure 6 illustrates the success rates and completion times for the **choice reaction** task (IDs: 1-30, 66-84), for varying bonus weights $w_{\text{bonus}} \in \{0, 0.25, 0.5, 8, 50, 100\}$. Agents trained with a bonus weight of 0 or 0.25 exhibit low success rates, typically failing to press buttons despite approaching them (IDs: 71-76, 78, 80, 82, 84). In contrast, a bonus weight of 0.5 or 8 (IDs: 1-3, 7-9, 14-15, 19-21, 25-28, 68-70, 77, 79, 81, 83) yields high success (and low task completion time) across all effort models except CTC: for 0.5, the agent first touches the buttons from the side before succeeding (ID: 77); for a value of 8, the success rate diminishes. While no negative effect of a large bonus on success rate was observed for policies trained without effort term, too large bonus weights can slow down learning. For instance, weights of

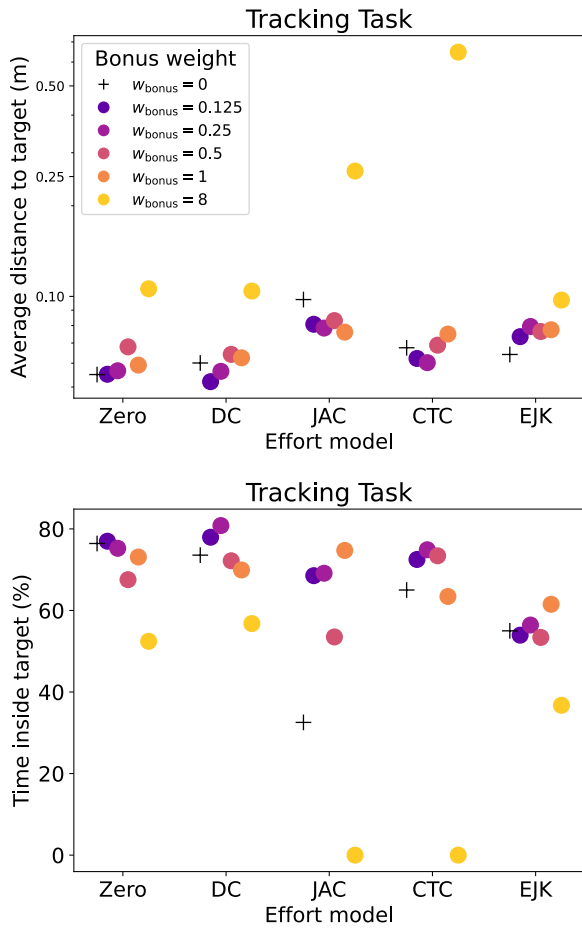


Figure 7: Effect of varying bonus weight on the average distance to the target (top; logarithmic scale) and time spent inside target (bottom) in the tracking task. Small bonus weights (≤ 1) generally result in smaller distances and more time spent inside the target.

8, 50, and 100 lead to convergence in 12M, 17M, and 24M steps, respectively (IDs: 31, 66, 67).

Figure 7 depicts the success rates and completion times for the **tracking** task (IDs: 151-180, 188-207), for varying bonus weights $w_{\text{bonus}} \in \{0, 0.125, 0.25, 0.5, 1, 8\}$. In contrast to the choice reaction task, a "large" bonus weight (in this case 8) does degrade performance compared to lower weights, including 0 (note the logarithmic scale for the average distance). The JAC and CTC models seem to profit most from a carefully chosen bonus weight.

4.2.2 Distance Weight. Figure 8 shows the effects of varying the distance weight ($w_{\text{distance}} \in \{0.1, 0.2, 1, 5, 10\}$) in the **choice reaction** task (IDs: 36–65). We recall that agents are trained with the exponential distance term and a fixed bonus weight of $w_{\text{bonus}} = 8$. A distance weight larger than 1 has limited impact on policies (IDs: 42–48, 50, 51, 54, 55, 58, 59, 62, 63). Lower distance weights lead to lower success for some effort models (outstretched arm for JAC,

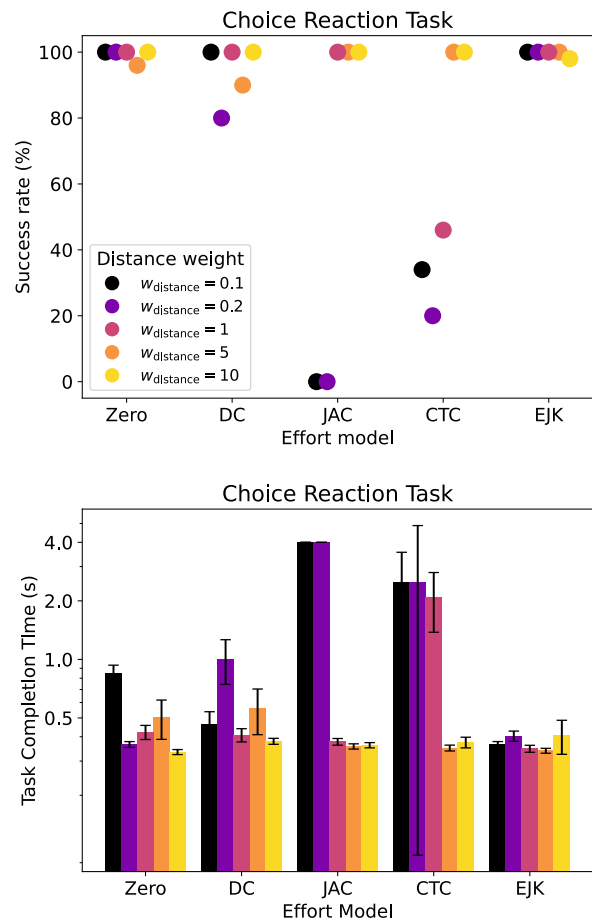


Figure 8: Effect of varying distance weights on success rate (top) and task completion time (bottom) in the choice reaction task. Error bars represent the standard error of the mean. Combined with the CTC or JAC effort models, small distance weights (≤ 1 and ≤ 0.2 , respectively) may lead to failure.

bent arm for CTC), and may lead to larger completion times (zero effort), where agents tap the button with the back of the hand, requiring several attempts to press a button (IDs: 45–47). The EJK effort model is the most robust to changes in the distance weight, while the CTC effort model combined with a distance weight of 0.2 exhibits a notably high standard error of the mean.

For the **tracking** task (Figure 9, IDs: 208–227), results indicate that increasing the distance weight (slightly) improves accuracy (note the logarithmic scale) and time inside target for agents trained with any non-zero effort models, particularly for the JAC, CTC, and EJK effort models. For the lowest distance weights, the JAC model tries to reach the target with an outstretched arm, while the CTC model follows the target shakily.

4.2.3 Effort Weight. Figure 10 shows the impact of varying effort weights ($w_{\text{effort}} \in \{0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10, 20\}$) on success rate and task completion time in the **choice reaction** task (IDs: 85–116).

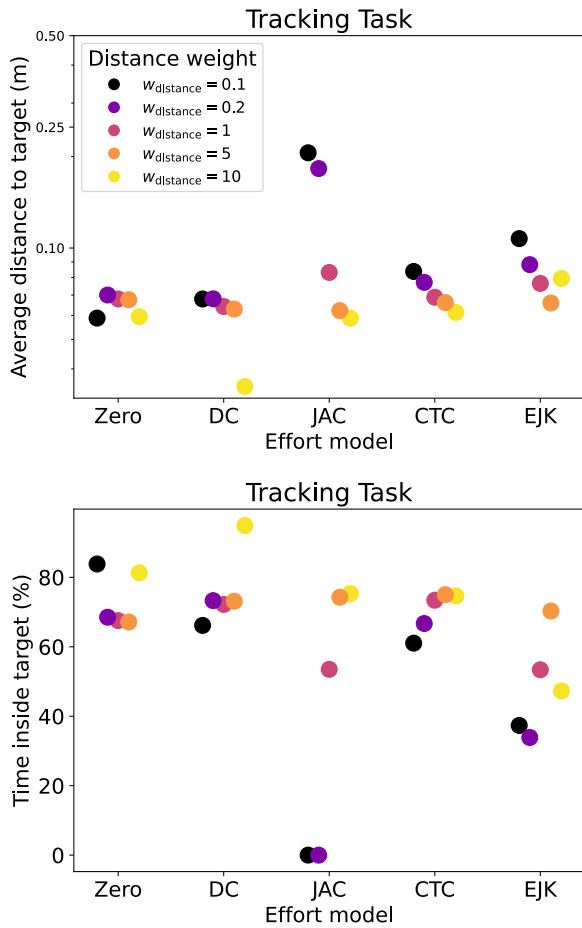


Figure 9: Effect of varying distance weights on the average distance to the target (top; logarithmic scale) and time spent inside target (bottom) in the tracking task. Combined with the JAC effort model, small distance weights (≤ 0.2) lead to agents not being able to reach the moving target. Overall, small distance weights lead to a less time inside the target.

We recall that $w_{\text{distance}} = 1$ and $w_{\text{bonus}} = 8$. Too large effort weights (≥ 2 , depending on the effort model) tend to reduce success rates and negatively affect completion time. Effort weights greater than 10 prevent the agent with the JAC, CTC, and EJK models from completing the task (IDs: 85–108), or allow them to press only the button closest to their initial position (ID: 36). Lowering effort weights does not necessarily lead to lower completion times, and weights below 0.2 may introduce a tendency to stick to sub-optimal movements and behaviors, e.g. with one agent occasionally struggling to correctly press the red button (ID: 43) The DC models seems most robust to effort weight changes.

Similarly, in the **tracking** task (Figure 11, IDs: 228–259), large effort weights lead to increased failure to point inside the target (again note the logarithmic scale). Effort weights greater than 10 show similar behavior to distance weights that are too low; in the EJK model, too large effort weights cause the agent to not move




Figure 10: Effect of varying effort weight on success rate (top) and task completion time (bottom) in the choice reaction task. Error bars represent the standard error of the mean. Larger effort weights lead to task failure for the JAC, CTC, and EJK effort models.

at all. With small effort weights, agents tend to stay closer to the target. As with the choice reaction task, the DC model appears to be least sensitive to changes in the effort weight.

4.3 Guidelines

Based on the above results, we distill guidelines, meant to be followed sequentially, to assist in the design of a reward function. The guidelines are presented in the order in which we suggest they should be followed.

First, we find that a guiding term, such as a distance-based reward, is crucial for training agents to perform goal-directed HCI tasks. Without this component, models in the choice reaction and pointing tasks fail to approach the target and instead perform random uncontrolled movements. Incorporating a distance term allows agents to effectively move towards the target. We therefore propose the following guideline:

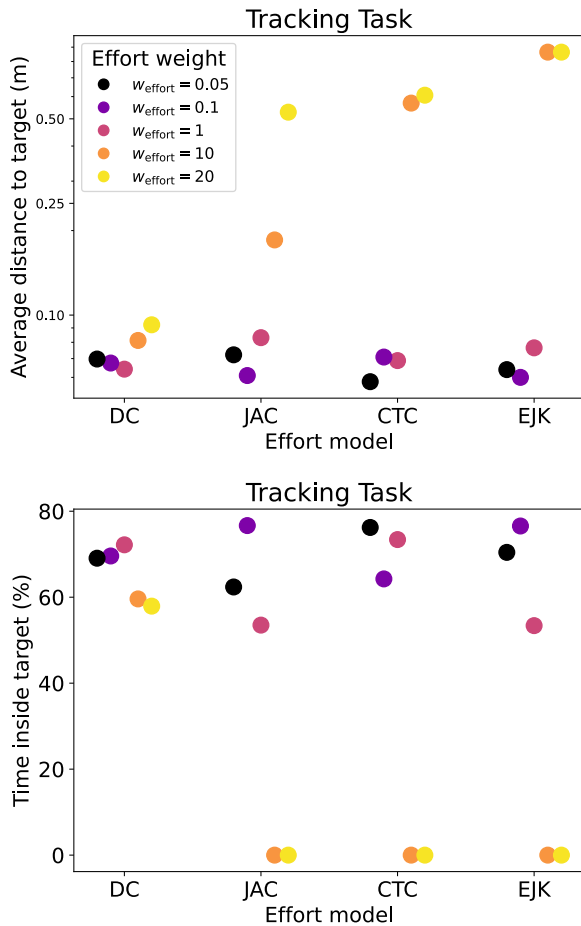


Figure 11: Effect of varying effort weight on the average distance to the target (top; logarithmic scale) and time spent inside target (bottom) in the tracking task. Large effort weights for the JAC, CTC, and EJK effort models lead to agents that fail to point inside the target at all.

G1 Integrate guidance components, such as distance-based rewards, to effectively guide the agent toward task completion.

Second, our findings highlight the need for a task completion bonus to achieve high success rates. Without this bonus, agents often adopt suboptimal strategies, such as touching buttons from the side (choice reaction), or exhibit excessive trembling, which prevents them from keeping the fingertip inside the target sphere for the required dwell time (pointing). We thus recommend:

G2 Include a task bonus to ensure task success.

The value of the bonus weights is also crucial. Setting the bonus weights in the choice reaction task too small results in agents that move the finger close the button but fail to successfully press them. While increasing the bonus weight eventually enables proper task execution, too large bonus weights can slow down the training. When the task bonus can be awarded multiple times per trial, as it was the case for tracking, we observed a preference toward smaller

bonus weights. In this case, a too large bonus weight can lead to inaccuracies.

G3 Adapt the bonus weight:

- If the trial terminates when the bonus is given: increase the bonus weight if the agent reaches the target but fails to complete the trial successfully. Reduce the bonus weight if the agent requires excessive training steps.
- If the bonus is awarded multiple times per trial, in addition to a guidance component: keep the bonus weight low and adjust for accuracy.

The results also suggest that effort models are not always necessary to generate successful movement trajectories. While we observed some irregularities in the choice reaction task (trying to repeatedly press the button with the back of the hand), none occurred in the pointing and tracking tasks. Hence, we suggest:

G4 Try without effort models first. If instabilities arise (e.g., excessive trembling), add an effort model.

Effort-based penalties offer a mechanism to encourage smoother or more efficient movement. However, they must be balanced carefully with distance-based rewards. A sufficiently high effort weight can suppress undesirable behavior like erratic motion (e.g., pressing with the back of the hand), but overly strong penalties may discourage movement altogether. Striking the right balance is task dependent and typically requires iterative weight tuning based on first results. We therefore recommend:

G5 Dynamically balance distance and effort weights: reduce the effort weight or increase the distance weight if the agent fails to approach the target; do the opposite if movement instability is observed.

4.4 Case Studies: Keyboard Typing and Remote Control Driving

We evaluate our proposed reward design guidelines on two HCI tasks: keyboard typing [24] and remote control [28], as shown in Figure 12. Unlike prior work that relies on curriculum learning and multi-stage training processes [24], our goal is to demonstrate that an appropriately designed reward function, utilizing our guidelines, can lead to successful policy learning with a straightforward training setup. Notably, the keyboard typing task is implemented and evaluated outside the user-in-the-box framework, demonstrating the broader applicability of our approach. We train agents for two million steps on the keyboard typing task and 200 million steps on the remote control task.

Keyboard Typing. For the **keyboard typing** task, performance is measured using the character error rate over 500 stimulus phrases (each fewer than 40 characters long), which are sampled from the same dataset as used in [24]; results are shown in Figure 13.

Following our guidelines, the initial reward function combined a distance term penalizing the distance between the fingertip and target key (G1) with a task bonus for correctly pressing the key (G2). Since the bonus immediately terminates the episode, we chose a large bonus weight $w_{\text{bonus}} = 30$ (G3). The resulting agent successfully learns the task and achieves character error rates $\leq 10\%$. Importantly, prior work required a complex two-stage curriculum approach and 5 million steps to reach a character error rate below




Figure 12: Illustrations of the keyboard typing task: (A) An agent trained with a distance-only reward attempting to press the keys from below, (B) an agent successfully pressing the keys, and (C) the remote control task.

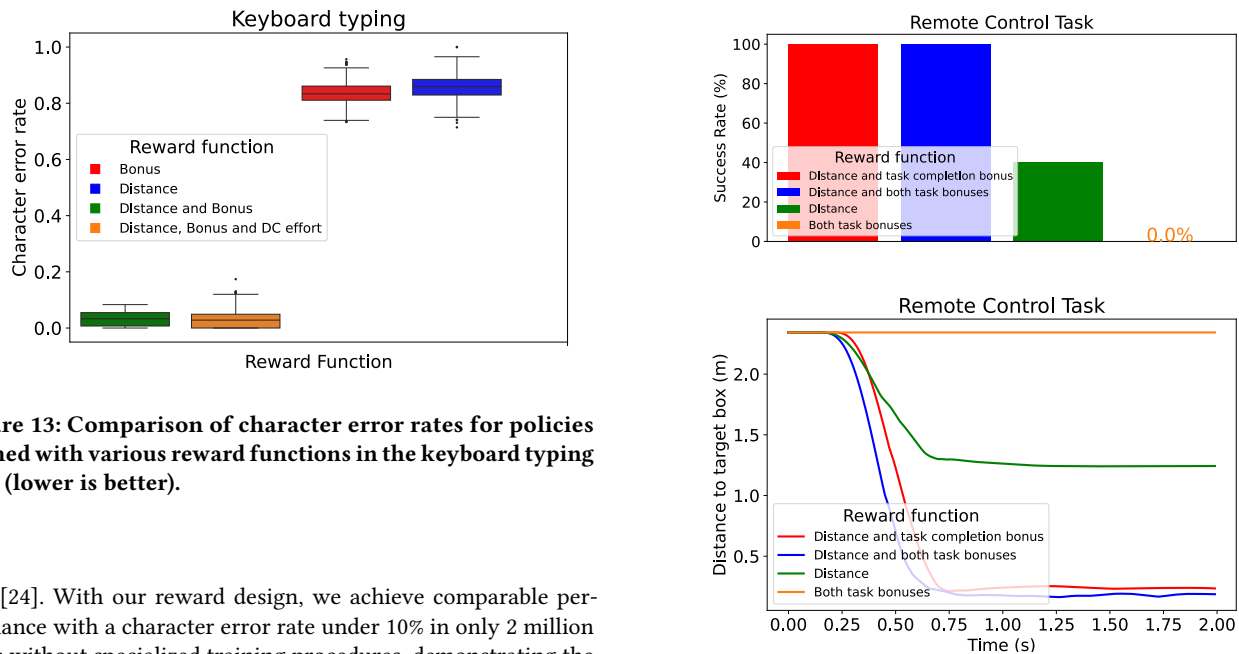


Figure 13: Comparison of character error rates for policies trained with various reward functions in the keyboard typing task (lower is better).

10% [24]. With our reward design, we achieve comparable performance with a character error rate under 10% in only 2 million steps without specialized training procedures, demonstrating the effectiveness of our guidelines.

Recognizing that effort models can be beneficial for shaping specific behavior, we also evaluate agents including the DC effort in the reward function, previously shown to perform robustly with respect to effort weights in related tasks. Using the same effort weight as in our earlier experiments ($w_{\text{effort}} = 1$), the resulting agents again succeed in the typing task, achieving character error rates below 20%, confirming that our guidelines generalize well to this setting.

To further validate the necessity of the distance and bonus reward components, we conduct more studies. The agent trained with only the task bonus (no distance guidance) fails to learn the right behavior, producing arbitrary movements and a character error rate above 75%. Conversely, the agent trained with only the distance term approaches the target keys but fails to press them, sometimes trying to press them from below, see Figure 12 (A). This policy results in a character error rate of over 75%. These findings provide empirical support for the necessity of both guidance and task completion incentives, as outlined in guidelines G1 and G2, with a bonus weight not too large to slow down the learning process (G3).

Figure 14: Comparison of success rates (top) and distances to the target (bottom) for policies trained with various reward functions in the remote control task. Combining distance and bonus terms is essential to ensure success. Without a distance term, the car does not move.

Remote Control. We additionally apply our framework to the **remote control** task and present the results in Figure 14. Each agent is tested over 10 episodes. Following our guidelines, we again start with a reward function consisting of a distance term and a task bonus (G1, G2). The task requires a sequence of actions: moving the fingertip to a joystick and using it to steer a car into a target box. Accordingly, we include two distance terms in the reward function: one penalizing the distance between the fingertip and the joystick, and one penalizing the distance between the car and the target box (G1). For both distance terms, we use the exponential distance. Since the task only requires the car to reach the target box, but allows for further corrections ("parking") after the time limit is

reached (i.e., no early termination is applied) [28], a success bonus for reaching the target box can be given for multiple timesteps. Therefore, we choose a small bonus weight $w_{\text{bonus}} = 1$, following G3. The resulting agent achieves a 100% success rate, consistently steering the car into the target.

Since this task is sequential, one might also consider joystick contact as an intermediate goal state required for correct task execution (G2). An agent trained with this additional, second bonus term also achieves 100% success and exhibits slightly faster car movement toward the target. Further results reconfirm the necessity of both distance guidance and a task bonus (G1, G2). The agent trained using only distance costs succeeds in only 40% of the episodes. Failed episodes result from the agent’s inability to reliably control the joystick to steer the car into the target box. The agent trained with task bonuses only fails completely, resulting in a 0% success rate. It either performs arbitrary movements or does not move at all, demonstrating the importance of guidance in the reward function. Due to the high computational cost of this task and the success of our baseline reward design, we do not train policies with additional effort models for this case (G4).

5 Discussion

We here summarize and contextualize our findings with respect to the research questions (Sections 5.1-5.3) and identify limitations and avenues for future work (Section 5.4).

5.1 RQ.1: Plausibility

Our analyses identify a dense distance term as an essential reward component to produce plausible, aimed movements across the five considered tasks. This distance term ‘guides’ the simulated user towards the desired target, e.g., the correct button in a keyboard typing task or the controller in an instrumented interaction task.

Many HCI tasks impose additional constraints beyond simple distance minimization, such as dwell time, force control, object grasping, or discrete gesture execution. This leads to a division of the task into at least two phases: (1) a reaching/aiming phase, where distance-based rewards are usually sufficient; and (2) a completion phase. For (2), additional reward terms are required. We found that providing a bonus term once upon task completion to be sufficient for the considered tasks, e.g., to learn keeping the fingertip inside the target sphere for a fixed dwell time (pointing), or applying enough force for a successful button pressing (choice reaction). Depending on the time and visuomotor complexity of the completion phase, we expect continuous reward terms, which, e.g., incentivize increasing the button pressing force, to provide additional guidance and thus simplify training. However, further analysis of such constraint-specific guidance terms is left for future work. In the tracking task, which simply requires keeping the fingertip close to the moving target, the bonus term proved mostly unnecessary when combined with a guiding distance. Rather, the bonus available in every time step plays a “guiding” role similar to the distance term, which reinforces the need to balance the two.

Our observations suggest that an effort model is not necessary for successful task completion. While biomechanical simulations traditionally use an effort model to constrain the overactuated action space and regularize the optimization landscape [5, 10, 21],

our observations do confirm earlier findings from [17]. In this work, a biomechanical arm model trained without an effort model was shown to generate mid-air movements that follow well-established characteristics such as Fitts’ Law [40] and the 2/3 Power Law [61].

Based on these results, we propose a set of reward design guidelines that are meant to be applied sequentially. While trade-offs, such as between distance and continuous bonus terms, may occur, they can be effectively balanced by following Guideline G3.

5.2 RQ.2: Sensitivity

While we found that the weighting of the reward function components affected the quality of the trained agents, the performance metrics and the observed qualitative behavior were quite robust to weight changes. In the choice reaction task, any bonus weight between 0.5 and 100 resulted in 100% accuracy when trained without an effort model, albeit with slower training times for larger bonus weight. Similarly, hardly any differences were observed between the tested distance weights, which spanned three orders of magnitude, when effort models were excluded. However, it is important to strike a good balance between effort and distance weights. Effort weights should not be chosen too large to avoid task failure, while the range of ‘working’ weights depends on the specific effort model. The DC effort model that directly penalizes squared muscle controls proved most robust against changes in the weighting. In the tracking task, the bonus term is not given only once per trial, but can be obtained multiple times, resulting in a larger accumulated sum of these reward terms over a trial. We hypothesize that this is the reason for the observed advantage of smaller bonus weights. However, any bonus weight between 0.125 and 1 resulted in good performance and plausible movement. Effects of distance and effort weights were similar to the choice reaction task.

Our findings suggest that designing reward functions does not require finding a very specific, unique set of cost weights, but rather allows for a fairly wide range of combinations that HCI researchers and designers might come up with, as long as some basic principles are followed. Adhering to the proposed guidelines promotes a more principled design process, which replaces prevailing unstructured trial-and-error approaches. While a few iterations of adjusting weights within Guidelines G3 and G4 might still be required, these guidelines provide clear instructions how to systematically proceed.

5.3 RQ.3: Generalizability

Based on the insights gained from over 450 policies trained in the choice reaction, pointing, and tracking task, we were able to distill a set of concise and clear guidelines for reward design. The analysis focused on three reward components chosen for their relevance and demonstrated effectiveness in goal-directed visuomotor tasks: (1) task completion bonuses, which are standard in RL; (2) proximity rewards, which help guide agents when task success is sparse; and (3) effort terms, which support human-plausible behavior as shown in prior HCI and motor control research [5, 10]. The resulting guidelines propose a sequence of actionable steps that address HCI researchers and designers with little or none experience, and may help them developing their own task-specific biomechanical simulations. We have validated our guidelines on two independent standard interaction tasks (keyboard typing and

remote control), demonstrating their applicability across diverse goal-directed visuomotor tasks.

While the five tasks considered cover a wide range of motor requirements, success conditions and contexts, there are many more interesting use cases that could benefit from biomechanical testing. We believe that the proposed guidelines can inspire the community to create simulated users for their own interaction techniques, task requirements, and user models. Extending the scope of RL-based biomechanical simulations is essential to consolidate and deepen our understanding of the role of composite reward functions. Such future work can ultimately form the basis for ongoing evaluation and refinement of the proposed guidelines, which we hope will catalyze this research activity. In addition, it would be interesting to see how advanced RL training methods, such as adaptive curricula [12, 50] and muscle-specific exploration strategies [4, 13, 62] complement reward design. These approaches could improve learning efficiency and robustness, but they also increase the complexity of the training process.

5.4 Limitations and Future Work

While our reward function design provides a structured approach for designing reward functions for RL agents in HCI tasks, several limitations remain.

In our analysis, we have only considered task bonus terms that were awarded for *successful* task performance. In environments that enable unsuccessful early termination (e.g., when a walking agent falls down), penalties for undesirable or failure states may help shape agent behavior. Integrating such penalties is straightforward within our framework, but requires further investigations.

Similarly, this work focused on distance terms as the main guiding concept. Extensions to higher-order reference signals, such as maintaining a fixed target speed, matching orientations in 6 DoF docking tasks, or tracking a reference velocity for virtual world navigation are conceivable. Moreover, additional reward components incentivizing aspects beyond success criteria such as ergonomics (e.g. reducing discomfort or pain), novelty or multi-agent collaboration could be integrated. The interplay between these additional factors and the three core reward components considered in this paper warrants further investigation.

Future work should also test our guidelines for a wider range of biomechanical models, e.g., including head or eye movements or actuable fingers, and develop more quantitative metrics to assess the motion plausibility of predicted movements and reduce reliance on manual inspection of movement videos prone to subjectivity. In addition, potential biases introduced by shaped reward terms, such as favoring certain targets or body postures, require further investigation to ensure that learned behaviors remain natural and unbiased. This is particularly relevant for modeling more strategic or creative interaction tasks.

Finally, modular simulation designs that separate the perceptual, cognitive, and motor control tasks that are currently intertwined in our end-to-end visuomotor learning approach will allow for better interpretability of individual reward design choices and more targeted diagnosis of learned behaviour.

6 Conclusion

Reward function design is critical for RL-based biomechanical simulations. We systematically investigated the influence of effort minimization, task bonuses, and target proximity incentives across representative HCI tasks: pointing, tracking, and choice reaction. We then performed an extensive analysis of how sensitive task success and completion time depend on the weights of these three reward components. Our analysis revealed that a guidance reward component is essential for task completion, while task bonuses are necessary for tasks with additional constraints beyond simple distance minimization. Notably, agents can succeed without an effort model, and the addition of effort models requires careful tuning to avoid performance degradation. Building on these findings, we proposed practical guidelines for reward function design that enable plausible, biomechanical simulations without requiring much RL expertise. We validated these guidelines on two additional tasks: remote control and keyboard typing. This work highlights the need for a deeper understanding of the subtleties involved in training RL-based user models, and contributes to making biomechanical user simulations a viable, low-barrier tool for HCI research and design.

Acknowledgments

This work was supported by EPSRC grant EP/W02456X/1. Hannah Selder and Arthur Fleig acknowledge the financial support by the Federal Ministry of Education and Research of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research „Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig“, project identification number: ScaDS.AI.

References

- [1] Marko Ackermann and Antonie J Van den Bogert. 2010. Optimality principles for model-based prediction of human gait. *Journal of biomechanics* 43, 6 (2010), 1055–1060.
- [2] M. M. Ayoub, Arun Walvekar, and Mike Petruno. 1974. A Biomechanical Model for the Upper Extremity Using Optimization Techniques. *SAE Transactions* 83 (1974), 1139–1148. <http://www.jstor.org/stable/44734426>
- [3] Myroslav Bachynskyi, Gregorio Palmas, Antti Oulasvirta, Jürgen Steimle, and Tino Weinkauff. 2015. Performance and ergonomics of touch surfaces: A comparative study using biomechanical simulation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, 1817–1826.
- [4] Cameron Berg, Vittorio Caggiano, and Vikash Kumar. 2024. Sar: Generalization of physiological agility and dexterity via synergistic action representation. *Autonomous Robots* 48, 8 (2024), 28.
- [5] Bastien Berret, Enrico Chiovetto, Francesco Nori, and Thierry Pozzo. 2011. Evidence for Composite Cost Functions in Arm Movement Planning: An Inverse Optimal Control Approach. *PLoS Computational Biology* 7, 10 (Oct. 2011), e1002183. doi:10.1371/journal.pcbi.1002183
- [6] Serena Booth, W. Bradley Knox, Julie Shah, Scott Niekum, Peter Stone, and Alessandro Allievi. 2023. The Perils of Trial-and-Error Reward Design: Misdesign through Overfitting and Invalid Task Specifications. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 5 (June 2023), 5920–5929. doi:10.1609/aaai.v37i5.25733
- [7] Vittorio Caggiano, Sudeep Dasari, and Vikash Kumar. 2023. MyoDex: a generalizable prior for dexterous manipulation. In *Proceedings of the 40th International Conference on Machine Learning (Honolulu, Hawaii, USA) (ICML '23)*. JMLR.org, New York, NY, USA, 20 pages.
- [8] Vittorio Caggiano, Huawei Wang, Guillaume Durandau, Massimo Sartori, and Vikash Kumar. 2022. MyoSuite: A Contact-rich Simulation Suite for Musculoskeletal Motor Control. In *Proceedings of The 4th Annual Learning for Dynamics and Control Conference (Proceedings of Machine Learning Research, Vol. 168)*, Roya Firoozi, Negar Mehr, Esen Yel, Rika Antonova, Jeannette Bohg, Mac

- Schwager, and Mykel Kochenderfer (Eds.). PMLR, New York, NY, USA, 492–507. <https://proceedings.mlr.press/v168/caggiano22a.html>
- [9] Suyog Chandramouli, Danqing Shi, Aini Putkonen, Sebastiaan De Peuter, Shanshan Zhang, Jussi Jokinen, Andrew Howes, and Antti Oulasvirta. 2024. A Workflow for Building Computationally Rational Models of Human Behavior. *Computational Brain and Behavior* 7, 3 (Sept. 2024), 399–419. doi:10.1007/s42113-024-00208-6 Publisher Copyright: © The Author(s) 2024..
- [10] Jhon P.F. Charaja, Isabell Wochner, Pierre Schumacher, Winfried Ilg, Martin Giese, Christophe Maufroy, Andreas Bulling, Syn Schmitt, Georg Martius, and Daniel F.B. Haeufle. 2024. Generating Realistic Arm Movements in Reinforcement Learning: A Quantitative Comparison of Reward Terms and Task Requirements. In *2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, Piscataway, NJ, USA, 562–568. doi:10.1109/biorob60516.2024.10719719
- [11] Noshaba Cheema, Laura A. Frey-Law, Kourosh Naderi, Jaakko Lehtinen, Philipp Slusallek, and Perttu Hämmäläinen. 2020. Predicting Mid-Air Interaction Movements and Fatigue Using Deep Reinforcement Learning. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–13. doi:10.1145/3313831.3376701
- [12] Alberto Silvio Chiappa, Pablo Tano, Nisheet Patel, Abigail Ingster, Alexandre Pouget, and Alexander Mathis. 2024. Acquiring musculoskeletal skills with curriculum-based reinforcement learning. *Neuron* 112, 23 (2024), 3969–3983.
- [13] Alberto Silvio Chiappa, Alessandro Marin Vargas, Ann Zixiang Huang, and Alexander Mathis. 2023. Latent exploration for reinforcement learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems (New Orleans, LA, USA) (NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 2466, 23 pages.
- [14] Michael Damsgaard, John Rasmussen, Søren Tørholm Christensen, Egidijus Surma, and Mark De Zee. 2006. Analysis of musculoskeletal systems in the AnyBody Modeling System. *Simulation Modelling Practice and Theory* 14, 8 (Nov. 2006), 1100–1111. doi:10.1016/j.simpat.2006.09.001 Publisher: Elsevier BV.
- [15] Scott L. Delp, Frank C. Anderson, Allison S. Arnold, Peter Loan, Ayman Habib, Chand T. John, Eran Guendelman, and Darryl G. Thelen. 2007. OpenSim: Open-Source Software to Create and Analyze Dynamic Simulations of Movement. *IEEE Transactions on Biomedical Engineering* 54, 11 (2007), 1940–1950. doi:10.1109/TBME.2007.901024
- [16] Rati Devidze, Goran Radanovic, Parameswaran Kamalaruban, and Adish Singla. 2024. Explicable reward design for reinforcement learning agents. In *Proceedings of the 35th International Conference on Neural Information Processing Systems (NIPS '21)*. Curran Associates Inc., Red Hook, NY, USA, Article 1539, 14 pages.
- [17] Florian Fischer, Miroslav Bachinski, Markus Klar, Arthur Fleig, and Jörg Müller. 2021. Reinforcement learning control of a biomechanical model of the upper extremity. *Scientific Reports* 11, 1 (2021), 14445.
- [18] Florian Fischer, Arthur Fleig, Markus Klar, and Jörg Müller. 2022. Optimal Feedback Control for Modeling Human–Computer Interaction. *ACM Transactions on Computer-Human Interaction* 29, 6 (Dec. 2022), 1–70. doi:10.1145/3524122
- [19] Florian Fischer, Aleks Ikkala, Markus Klar, Arthur Fleig, Miroslav Bachinski, Roderick Murray-Smith, Perttu Hämmäläinen, Antti Oulasvirta, and Jörg Müller. 2024. SIM2VR: Towards Automated Biomechanical Testing in VR. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. ACM, Pittsburgh PA USA, 1–15. doi:10.1145/3654777.3676452
- [20] T Flash and N Hogan. 1985. The coordination of arm movements: an experimentally confirmed mathematical model. *The Journal of Neuroscience* 5, 7 (July 1985), 1688–1703. doi:10.1523/JNEUROSCI.05-07-01688.1985
- [21] Emmanuel Guigon, Pierre Baraduc, and Michel Desmurget. 2007. Computational Motor Control: Redundancy and Invariance. *Journal of Neurophysiology* 97, 1 (Jan. 2007), 331–347. doi:10.1152/jn.00290.2006
- [22] Tuomas Haarnoja, Ben Moran, Guy Lever, Sandy H. Huang, Dhruva Tirumala, Jan Humplik, Markus Wulfmeier, Saran Tunyasuvunakool, Noah Y. Siegel, Roland Hafner, Michael Bloesch, Kristian Hartikainen, Arunkumar Byravan, Leonard Hasenclever, Yuval Tassa, Fereshteh Sadeghi, Nathan Batchelor, Federico Casarini, Stefano Saliceti, Charles Game, Neil Sreendran, Kushal Patel, Marlon Gwira, Andrea Huber, Nicole Hurley, Francesco Nori, Raia Hadsell, and Nicolas Heess. 2024. Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *Science Robotics* 9, 89 (2024), eadi8022. doi:10.1126/scirobotics.adi8022 arXiv:https://www.science.org/doi/pdf/10.1126/scirobotics.adi8022
- [23] Yixu He, Yang Liu, Lan Yang, and Xiaobo Qu. 2024. Exploring the design of reward functions in deep reinforcement learning-based vehicle velocity control algorithms. *Transportation Letters* 16, 10 (Nov. 2024), 1338–1352. doi:10.1080/19427867.2024.2305018
- [24] Lorenz Hetzel, John Dudley, Anna Maria Feit, and Per Ola Kristensson. 2021. Complex Interaction as Emergent Behaviour: Simulating Mid-Air Virtual Keyboard Typing using Reinforcement Learning. *IEEE Transactions on Visualization and Computer Graphics* 27, 11 (2021), 4140–4149. doi:10.1109/TVCG.2021.3106494
- [25] Katherine RS Holzbaur, Wendy M Murray, and Scott L Delp. 2005. A model of the upper extremity for simulating musculoskeletal surgery and analyzing neuromuscular control. *Annals of biomedical engineering* 33 (2005), 829–840.
- [26] Yujing Hu, Weixun Wang, Hangtian Jia, Yixiang Wang, Yingfeng Chen, Jianye Hao, Feng Wu, and Changjie Fan. 2020. Learning to utilize shaping rewards: A new approach of reward shaping. *Advances in Neural Information Processing Systems* 33 (2020), 15931–15941.
- [27] Seokhyun Hwang, Jeongseok Oh, Seongjun Kang, Minwoo Seong, Ahmed Ibrahim Ahmed Mohamed Elsharkawy, and Seungjun Kim. 2024. ErgoPulse: Electrifying Your Lower Body With Biomechanical Simulation-based Electrical Muscle Stimulation Haptic System in Virtual Reality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–21. doi:10.1145/3613904.3642008
- [28] Aleks Ikkala, Florian Fischer, Markus Klar, Miroslav Bachinski, Arthur Fleig, Andrew Howes, Perttu Hämmäläinen, Jörg Müller, Roderick Murray-Smith, and Antti Oulasvirta. 2022. Breathing Life Into Biomechanical User Models. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. ACM, Bend OR USA, 1–14. doi:10.1145/3526113.3545689
- [29] Figure AI Inc. 2025. Natural Humanoid Walk Using Reinforcement Learning. Retrieved March 31, 2025 from <https://www.figure.ai/news/reinforcement-learning-walking>
- [30] Janghyeon Kim, Ho-Jin Jung, Dae Han Sim, Ji-Hyeon Yoo, Song Woo Kim, and Han Ul Yoon. 2023. An Approach to Design a Biomechanically-Inspired Reward Function to Solve a Patience Cube Under Reinforcement Learning Framework. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Piscataway, NJ, USA, 5653–5660. doi:10.1109/IROS55552.2023.10341831
- [31] Markus Klar, Florian Fischer, Arthur Fleig, Miroslav Bachinski, and Jörg Müller. 2023. Simulating Interaction Movements via Model Predictive Control. *ACM Transactions on Computer-Human Interaction* 30, 3 (June 2023), 1–50. doi:10.1145/3577016
- [32] Bart Koopman, Henk J Grootenboer, and Henk J De Jongh. 1995. An inverse dynamics model for the analysis, reconstruction and prediction of bipedal walking. *Journal of biomechanics* 28, 11 (1995), 1369–1376.
- [33] Visak Kumar. 2021. Modelling Human Kinetics and Kinematics during Walking using Reinforcement Learning. arXiv:2103.08125 [cs.RO] <https://arxiv.org/abs/2103.08125> arXiv preprint arXiv:2103.08125.
- [34] Ariel Kwiatkowski, Vicky Kalogeiton, Julien Pettré, and Marie-Paule Cani. 2023. Reward Function Design for Crowd Simulation via Reinforcement Learning. In *Proceedings of the 16th ACM SIGGRAPH Conference on Motion, Interaction and Games (Rennes, France) (MIG '23)*. Association for Computing Machinery, New York, NY, USA, Article 4, 7 pages. doi:10.1145/3623264.3624452
- [35] Thomas Langerak, Sammy Christen, Mert Albaba, Christoph Gebhardt, Christian Holz, and Otmar Hilliges. 2024. MARLUI: Multi-Agent Reinforcement Learning for Adaptive Point-and-Click UIs. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (June 2024), 1–27. doi:10.1145/3661147
- [36] Sung-Hee Lee, Eftychios Sifakis, and Demetri Terzopoulos. 2009. Comprehensive biomechanical modeling and simulation of the upper body. *ACM Transactions on Graphics (TOG)* 28, 4 (2009), 1–17.
- [37] Raz Leib, Ian S. Howard, Matthew Millard, and David W. Franklin. 2024. Behavioral Motor Performance. *Comprehensive Physiology* 14, 1 (2024), 5179–5224. doi:10.1002/j.2040-4603.2024.tb00286.x
- [38] Dan Liu and Emanuel Todorov. 2007. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *Journal of Neuroscience* 27, 35 (2007), 9354–9368.
- [39] Ramona Maas and Sigrid Leyendecker. 2013. Biomechanical optimal control of human arm motion. *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics* 227, 4 (2013), 375–389.
- [40] I Scott MacKenzie. 1992. Fitts' law as a research and design tool in human-computer interaction. *Human-computer interaction* 7, 1 (1992), 91–139.
- [41] Pascal Madeleine, Afshin Samani, Mark de Zee, and Uwe Kersting. 2011. Biomechanics of Human Movement. In *IFMBE Proceedings*. Springer Berlin Heidelberg, Berlin, Heidelberg, 237–240. doi:10.1007/978-3-642-21683-1_60 ISSN: 1680-0737, 1433-9277.
- [42] J. Alberto Álvarez Martín, Henrik Gollee, Jörg Müller, and Roderick Murray-Smith. 2021. Intermittent Control as a Model of Mouse Movements. *ACM Transactions on Computer-Human Interaction* 28, 5 (Oct. 2021), 1–46. doi:10.1145/3461836 Publisher: Association for Computing Machinery (ACM).
- [43] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning.
- [44] Hee-Seung Moon, Yi-Chi Liao, Chenyu Li, Byungjoo Lee, and Antti Oulasvirta. 2024. Real-time 3D Target Inference via Biomechanical Simulation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–18. doi:10.1145/3613904.3642131
- [45] Roderick Murray-Smith, Antti Oulasvirta, Andrew Howes, Jörg Müller, Aleks Ikkala, Miroslav Bachinski, Arthur Fleig, Florian Fischer, and Markus Klar. 2022. What simulation can do for HCI research. *Interactions* 29, 6 (Nov. 2022), 48–53. doi:10.1145/3564038
- [46] Calarina Muslimani, Kerrick Johnstonbaugh, Suyog Chandramouli, Serena Booth, W. Bradley Knox, and Matthew E. Taylor. 2025. Towards Improving Reward Design in RL: A Reward Alignment Metric for RL Practitioners. doi:10.48550/

- arXiv:2503.05996 arXiv:2503.05996 [cs].
- [47] Arne J Nagengast, Daniel A Braun, and Daniel M Wolpert. 2011. Risk sensitivity in a motor task with speed-accuracy trade-off. *Journal of neurophysiology* 105, 6 (2011), 2668–2674.
- [48] Masaki Nakada, Tao Zhou, Honglin Chen, Tomer Weiss, and Demetri Terzopoulos. 2018. Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Transactions on Graphics* 37, 4 (Aug. 2018), 1–15. doi:10.1145/3197517.3201305 Publisher: Association for Computing Machinery (ACM).
- [49] Eri Nakano, Hiroshi Imamizu, Rieko Osu, Yoji Uno, Hiroaki Gomi, Toshinori Yoshioka, and Mitsuo Kawato. 1999. Quantitative examinations of internal representations for arm trajectory planning: minimum commanded torque change model. *Journal of Neurophysiology* 81, 5 (1999), 2140–2155.
- [50] Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. 2020. Curriculum learning for reinforcement learning domains: A framework and survey. *Journal of Machine Learning Research* 21, 181 (2020), 1–50.
- [51] Peter D Neilson. 1993. The problem of redundancy in movement control: the adaptive model theory approach. *Psychological Research* 55, 2 (1993), 99–106.
- [52] Winston L Nelson. 1983. Physical principles for economies of skilled movements. *Biological cybernetics* 46 (1983), 135–147.
- [53] Andrew Y. Ng, Daishi Harada, and Stuart J. Russell. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *Proceedings of the Sixteenth International Conference on Machine Learning (ICML '99)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 278–287.
- [54] Jorge Nocedal and Stephen J. Wright. 2006. *Numerical optimization* (2. ed. ed.). Springer, New York, NY.
- [55] Katharine Nowakowski, Philippe Carvalho, Jean-Baptiste Six, Yann Maillet, Anh Tu Nguyen, Ismail Seghiri, Loick M’Pemba, Theo Marcille, Sy Toan Ngo, and Tien-Tuan Dao. 2021. Human locomotion with reinforcement learning using bioinspired reward reshaping strategies. *Medical & Biological Engineering & Computing* 59, 1 (Jan. 2021), 243–256. doi:10.1007/s11517-020-02309-3
- [56] Andrew Patterson, Samuel Neumann, Martha White, and Adam White. 2024. Empirical Design in Reinforcement Learning. doi:10.48550/arXiv.2304.01315 arXiv:2304.01315 [cs].
- [57] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. 2018. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)* 37, 4 (2018), 1–14.
- [58] Claudio Pizzolato, Monica Reggiani, Luca Modenese, and David G Lloyd. 2017. Real-time inverse kinematics and inverse dynamics for lower limb applications using OpenSim. *Computer methods in biomechanics and biomedical engineering* 20, 4 (2017), 436–445.
- [59] Tommaso Proietti, Emmanuel Guigon, Agnès Roby-Brami, and Nathanaël Jarrassé. 2017. Modifying upper-limb inter-joint coordination in healthy subjects by training with a robotic exoskeleton. *Journal of NeuroEngineering and Rehabilitation* 14, 1 (Dec. 2017), 55. doi:10.1186/s12984-017-0254-x
- [60] Katherine R Saul, Xiao Hu, Craig M Goehler, Meghan E Vidt, Melissa Daly, Anca Velisar, and Wendy M Murray. 2015. Benchmarking of dynamic simulation predictions in two software platforms using an upper limb musculoskeletal model. *Computer methods in biomechanics and biomedical engineering* 18, 13 (2015), 1445–1458.
- [61] Stefan Schaal and Dagmar Sternad. 2001. Origins and violations of the 2/3 power law in rhythmic three-dimensional arm movements. *Experimental brain research* 136 (2001), 60–72.
- [62] Pierre Schumacher, Daniel F.B. Haeufle, Dieter Büchler, Syn Schmitt, and Georg Martius. 2023. DEP-RL: Embodied Exploration for Reinforcement Learning in Overactuated and Musculoskeletal Systems. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*. ICLR, Appleton, WI, USA, 1–29. https://openreview.net/forum?id=C-xa_D3oTj6
- [63] Hannah Selder, Florian Fischer, Per Ola Kristensson, and Arthur Fleig. 2025. Supporting dataset for Demystifying Reward Design in Reinforcement Learning for Upper Extremity Interaction: Practical Guidelines for Biomechanical Simulations in HCI. Zenodo. doi:10.5281/zenodo.15845404 [Data set].
- [64] Hannah Selder, Florian Fischer, Per Ola Kristensson, and Arthur Fleig. 2025. What Makes a Model Breathe? Understanding Reinforcement Learning Reward Function Design in Biomechanical User Simulation. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 592, 10 pages. doi:10.1145/3706599.3719699
- [65] Danqing Shi, Yujun Zhu, Jussi P. P. Jokinen, Aditya Acharya, Aini Putkonen, Shumin Zhai, and Antti Oulasvirta. 2024. CRTypist: Simulating Touchscreen Typing Behavior via Computational Rationality. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. ACM, Honolulu HI USA, 1–17. doi:10.1145/3613904.3642918
- [66] Seungmoon Song, Łukasz Kidziński, Xue Bin Peng, Carmichael Ong, Jennifer Hicks, Sergey Levine, Christopher G Atkeson, and Scott L Delp. 2021. Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation. *Journal of neuroengineering and rehabilitation* 18 (2021), 1–17.
- [67] Richard S Sutton, Andrew G Barto, et al. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge, London.
- [68] SIMA Team, Maria Abi Raad, Arun Ahuja, Catarina Barros, Frederic Besse, Andrew Bolt, Adrian Bolton, Bethanie Brownfield, Gavin Buttimore, Max Cant, Sarah Chakera, Stephanie C. Y. Chan, Jeff Clune, Adrian Collister, Vikki Copeman, Alex Cullum, Ishita Dasgupta, Dario de Cesare, Julia Di Trapani, Yani Donchev, Emma Dunleavy, Martin Engelcke, Ryan Faulkner, Frankie Garcia, Charles Gbadamosi, Zhitao Gong, Lucy Gonzales, Kshitij Gupta, Karol Gregor, Arne Olav Hallingstad, Tim Harley, Sam Haves, Felix Hill, Ed Hirst, Drew A. Hudson, Jony Hudson, Steph Hughes-Fitt, Danilo J. Rezende, Mimi Jasarevic, Laura Kampis, Rosemary Ke, Thomas Keck, Junkyung Kim, Oscar Knagg, Kavya Kopparapu, Rory Lawton, Andrew Lampinen, Shane Legg, Alexander Lerchner, Marjorie Limont, Yulan Liu, Maria Loks-Thompson, Joseph Marino, Kathryn Martin Cussons, Loic Matthey, Siobhan McLaughlin, Piermaria Mendolicchio, Hamza Merzic, Anna Mitenkova, Alexandre Moufarek, Valeria Oliveira, Yanko Oliveira, Hannah Openshaw, Renke Pan, Aneesh Pappu, Alex Platonov, Ollie Purkiss, David Reichert, John Reid, Pierre Harvey Richemond, Tyson Roberts, Giles Ruscoe, Jaume Sanchez Elias, Tasha Sanders, Daniel P. Sawyer, Tim Scholtes, Guy Simmons, Daniel Slater, Hubert Soyer, Heiko Strathmann, Peter Stys, Allison C. Tam, Denis Teplyashin, Tayfun Terzi, Davide Verelli, Bojan Vujatovic, Marcus Wainwright, Jane X. Wang, Zhengdong Wang, Daan Wierstra, Duncan Williams, Nathaniel Wong, Sarah York, and Nick Young. 2024. Scaling Instructable Agents Across Many Simulated Worlds. arXiv:2404.10179 [cs.RO] <https://arxiv.org/abs/2404.10179>
- [69] Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Piscataway NJ USA, 5026–5033. doi:10.1109/IROS.2012.6386109
- [70] Emanuel Todorov and Michael I. Jordan. 2002. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5, 11 (Nov. 2002), 1226–1235. doi:10.1038/nn963 Publisher: Springer Science and Business Media LLC.
- [71] Rubin Wang and Fanji Gu (Eds.). 2011. *Advances in Cognitive Neurodynamics (II): Proceedings of the Second International Conference on Cognitive Neurodynamics - 2009*. Springer Netherlands, Dordrecht. doi:10.1007/978-90-481-9695-1
- [72] Robert C Wilson and Anne Ge Collins. 2019. Ten simple rules for the computational modeling of behavioral data. *eLife* 8 (Nov. 2019), e49547. doi:10.7554/eLife.49547
- [73] DA Winter. 1984. Biomechanics of human movement with applications to the study of human locomotion. *Critical reviews in biomedical engineering* 9, 4 (1984), 287–314. <http://europepmc.org/abstract/MED/6368126>
- [74] Isabell Wochner, Danny Driess, Heiko Zimmermann, Daniel F. B. Haeufle, Marc Toussaint, and Syn Schmitt. 2020. Optimality Principles in Human Point-to-Manifold Reaching Accounting for Muscle Dynamics. *Frontiers in Computational Neuroscience* 14 (May 2020), 38. doi:10.3389/fncom.2020.00038
- [75] Jingda Wu, Chao Huang, Hailong Huang, Chen Lv, Yuntong Wang, and Fei-Yue Wang. 2024. Recent advances in reinforcement learning-based autonomous driving behavior planning: A survey. *Transportation Research Part C: Emerging Technologies* 164 (2024), 104654. doi:10.1016/j.trc.2024.104654
- [76] Yangmengfei Xu, Vincent Crocher, Justin Fong, Ying Tan, and Denny Oetomo. 2021. Inducing Human Motor Adaptation Without Explicit Error Feedback: A Motor Cost Approach. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 29 (2021), 1403–1412. doi:10.1109/TNSRE.2021.3096516
- [77] Lei Zhang, Lukas Lengersdorff, Nace Mikus, Jan Gläscher, and Claus Lamm. 2020. Using reinforcement learning models in social neuroscience: frameworks, pitfalls and suggestions of best practices. *Social Cognitive and Affective Neuroscience* 15, 6 (July 2020), 695–707. doi:10.1093/scan/nsaa089
- [78] Henry Zhu, Abhishek Gupta, Aravind Rajeswaran, Sergey Levine, and Vikash Kumar. 2019. Dexterous Manipulation with Deep Reinforcement Learning: Efficient, General, and Low-Cost. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, Montreal, QC, Canada, 3651–3657. doi:10.1109/ICRA.2019.8794102