CARL: Controllable Agent with Reinforcement Learning for Quadruped Locomotion

YING-SHENG LUO*, Inventec Corp., Taiwan

JONATHAN HANS SOESEN*, Inventec Corp., Taiwan

TRISTA PEI-CHUN CHEN, Inventec Corp., Taiwan

WEI-CHAO CHEN, Skywatch Innovation Inc. and Inventec Corp., Taiwan

Fig. 1. Our controllable agent produces natural movements (blue) alike depicted in the reference motion clip (red) and adapt to external perturbations.

Motion synthesis in a dynamic environment has been a long-standing problem for character animation. Methods using motion capture data tend to scale poorly in complex environments because of their larger capturing and labeling requirement. Physics-based controllers are effective in this regard, albeit less controllable. In this paper, we present CARL, a quadruped agent that can be controlled with high-level directives and react naturally to dynamic environments. Starting with an agent that can imitate individual animation clips, we use Generative Adversarial Networks to adapt high-level controls, such as speed and heading, to action distributions that correspond to the original animations. Further fine-tuning through the deep reinforcement learning enables the agent to recover from unseen external perturbations while producing smooth transitions. It then becomes straightforward to

create autonomous agents in dynamic environments by adding navigation modules over the entire process. We evaluate our approach by measuring the agent's ability to follow user control and provide a visual analysis of the generated motion to show its effectiveness.

CCS Concepts: • **Computing methodologies** → **Animation**; *Physical simulation*; *Reinforcement learning*.

Additional Key Words and Phrases: deep reinforcement learning (DRL), generative adversarial network (GAN), motion synthesis, locomotion, quadruped

ACM Reference Format:

Ying-Sheng Luo, Jonathan Hans Soeseno, Trista Pei-Chun Chen, and Wei-Chao Chen. 2020. **CARL: Controllable Agent with Reinforcement Learning for Quadruped Locomotion**. *ACM Trans. Graph.* 39, 4, Article 38 (July 2020), 10 pages. <https://doi.org/10.1145/3386569.3392433>

1 INTRODUCTION

The quality of character animation in cartoons, video games, and digital special effects have improved drastically in the past decades with new tools and techniques developed by researchers in the field. Amongst various types of characters, quadrupeds are especially challenging to animate due to their wide variations of style, cadence, and gait pattern. For real-time applications such as video games, the

*Joint first authors.

Authors' addresses: Ying-Sheng Luo, luo.ying-sheng@inventec.com, Inventec Corp., Taiwan; Jonathan Hans Soeseno, soeseno.jonathan@inventec.com, Inventec Corp., Taiwan; Trista Pei-Chun Chen, chen.trista@inventec.com, Inventec Corp., Taiwan; Wei-Chao Chen, weichao.chen@skywatch24.com, Skywatch Innovation Inc., Inventec Corp., Taiwan.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

0730-0301/2020/7-ART38

<https://doi.org/10.1145/3386569.3392433>

need to react dynamically to the environments further complicates the problem.

Traditionally, to synthesize new animations from motion capture data, one would create an interpolation structure such as a motion graph [Kovar et al. 2002; Lee et al. 2002], where the nodes represent well-defined actions from motion capture data, and the edges define the transition between the actions. Aside from the long and tedious process of labeling the graph, it is often difficult to acquire sufficient motion capture data for quadrupeds to cover different gait patterns and styles. Furthermore, the motion graph would become impractically big and complex in dynamic environments to take into account the numerous interactions between the agent and its surroundings. Despite the complexity, the motion graph would still not be useful for motion synthesis when unseen scenarios arise.

Research on the kinematic controller solves the labeling problem by reducing the need for crafting transitions between actions while allowing users to control the agent to produce the desired motions [Zhang et al. 2018]. But since a kinematic controller is designed to imitate the motion dataset, the agent would fail to respond naturally when it encounters unseen interactions between the agent and its surroundings in dynamic environments. For example, in a scenario involving a quadruped agent walking on a slippery, undulating boat, it would clearly be highly impractical to collect, or hand-engineer, enough reference motions to train the kinematic controllers. One can certainly resort to physics-based controllers to model complex phenomena effectively, as a physical simulation enables the agent to produce meaningful reactions to external perturbations without the need to collect or animate such a reaction. Although, physical constraints such as gravity, friction, and collision introduce numerous difficulties in designing a physics-based controller.

In this paper, we propose a **data-driven, physics-based, controllable** quadruped agent by adopting (1) the ability to interact physically with the dynamic environments, and (2) the natural movements learned from reference motion clips. It is a three-stage process that starts by learning the reference motions through imitation. The agent then learns to map high-level user controls such as speed and heading into joint actions with Generative Adversarial Networks (GANs). Through Deep Reinforcement Learning (DRL), the agent then gains the ability to adapt and recover from unseen scenarios, allowing us to synthesize meaningful reactions against external perturbations. One can then control the trained agent by attaching navigation modules that emulate higher-level directive controls, such as path-finding and ray-sensor. We summarize our contributions as follows:

- A GAN supervision framework to effectively map between high-level user controls and the learned natural movements,
- A physics-based controller for quadruped agents trained through DRL that can adapt to various external perturbations while producing meaningful reactions without the need for action labels, and
- A methodology to attach high-level navigation modules to the agent for tackling motion synthesis tasks involving agent-environment dynamic interactions.

2 RELATED WORKS

Real-time quadruped motion synthesis not only needs to consider the wide range of variations in gait patterns but also dynamic and unseen scenarios that happen in real-time. Works in the topic of real-time quadruped motion synthesis can be divided into two main categories: kinematic controllers and physics-based controllers. Recent advances in deep learning have also been incorporated into both categories of methods.

2.1 Kinematic Controllers

There has been an extensive amount of research works on character animation using kinematic-based methods. Classical approaches [Kovar and Gleicher 2004; Kovar et al. 2002; Lee et al. 2002; Safonova and Hodgins 2007] constructed motion graphs that included edges corresponding to each short motion clip. The resulting motions were then synthesized by graph searching. Alternatively, [Chai and Hodgins 2005; Grochow et al. 2004; Levine et al. 2012; Shin and Lee 2006] enabled smooth transitions through low-dimensional exploration. Huang et al. [2013] used an asynchronous time warping approach to handle the gait transitions of a quadruped agent. Despite their successes, motion graphs either require a large memory footprint, high order computation time, or extensive preprocessing, thus limiting the scalability to more complex movements or even larger datasets.

Recent approaches use deep learning that is known to scale better with a larger dataset. Recurrent Neural Network based approaches [Fragkiadaki et al. 2015; Lee et al. 2018; Zhou et al. 2018] predicted the future states given previous observations. Holden et al. [2016] combined a standard feed-forward neural network with an auto-encoder to produce motions with a high degree of realism. In their later work, [Holden et al. 2017] used a neural network to model the phase function of a biped agent, enabling the agent to traverse uneven terrains. However, their approach required extensive labeling of the agent's phase. Exploiting mixtures of experts model to learn quadruped locomotion, Mode-Adaptive Neural Networks [Zhang et al. 2018] suppressed the need for such a labeling process. Other recent works from [Lee et al. 2018; Starke et al. 2019] could produce less foot sliding artifact while interacting with the environment but sacrifice the control responsiveness.

In general, kinematic controllers scale with the quality of the motion clip dataset and often generate motions with higher-quality compared to physics-based controllers. However, their ability in producing dynamic movements where complex environments are presented is limited by the availability of such motions in the dataset. Capturing and hand-engineering enough motion data to accommodate a large number of possible agent behaviors are impractical.

2.2 Physics Based Controllers

Physics-based methods offer an effective way to model complex scenarios. Such types of controllers not only require to follow control commands but also need to maintain balance. Early works in locomotion controller design required human-insight [Coros et al. 2009, 2010, 2011; Yin et al. 2007]. Lee et al. [2010] designed a locomotion controller for biped agents by tracking modulated reference trajectories from motion capture data. Ye and Liu [2010] designed

an abstract dynamic model with hybrid optimization processes to achieve robust locomotion controllers. Alternatively, works in value iteration [Coros et al. 2009] developed task-based controllers for biped agents with varying body proportions. Peng et al. [2015] also used value iteration to produce a quadruped locomotion controller in 2D that could traverse terrain with gaps. However, their approach required manual selection of important features such as desired joint angles and leg forces. Although these approaches were robust, they did not scale with complex controls such as stylization or other high-level directive controls.

Trajectory-based approaches have also been explored in the past few years [Liu et al. 2015, 2010; Wampler and Popović 2009]. Liu et al. [2010] developed an open-loop controller for biped agents. These approaches did not scale well with complex and long motions. Such a problem was later solved by [Liu et al. 2015] through the reconstruction of past experiences. Recent works on biped and quadruped agents [Wampler and Popović 2009; Wampler et al. 2014] explored the gait patterns of various legged animals. Levine and Popović [2012] produced plausible variations of different locomotion skills such as walking, running, and rolling. These approaches tended to struggle with long-term planning. Recent works [Hämäläinen et al. 2015; Tassa et al. 2012, 2014] extended the offline approach to online optimization methods. Optimizing over short predictive horizons produced robust controllers [Da Silva et al. 2008; Tassa et al. 2012]. Physics-based controllers are considered an effective way to model complex phenomena, because they produce novel movements from its interactions with the environment. However, incorporating high-level controls in physics-based controllers is not a trivial task and often causes the controllers to look unnatural or is limited in its movement diversity.

2.3 Deep Reinforcement Learning (DRL)

Physics-based methods with deep neural networks, esp., Deep Reinforcement Learning (DRL), [Bansal et al. 2018; Brockman et al. 2016; Heess et al. 2017; Liu and Hodgins 2017; Merel et al. 2018; Peng et al. 2016; Schulman et al. 2015a, 2017; Sutton et al. 1998] trained controllers capable of performing diverse motions. Works from [Hwangbo et al. 2019; Won and Lee 2019] shown the success in controlling robots or agents with different morphology. Aerial locomotion was also explored through DRL algorithms [Won et al. 2017, 2018]. Works in imitation learning [Chentanez et al. 2018; Peng et al. 2018] produced high-quality motions through imitating well-defined short-clips. However, these methods' ability to scale with the size of motion capture data was limited.

Recent works from [Park et al. 2019] and [Bergamin et al. 2019] combined kinematic controllers with DRL to produce a responsive controller for biped agents. Multiplicative compositional policies (MCP) from [Peng et al. 2019] adopted a new way to blend low-level primitives to follow the high-level user intent. Despite their successes with biped controllers, these techniques could not generalize to quadruped agents. As discussed in [Zhang et al. 2018], modeling gait transitions of quadruped agents were harder compared to biped agents. This was because of the gait pattern complexity and the limited number of available motion capture data.

2.4 Adversarial Learning on DRL

The main challenge of physics-based controllers is to produce natural movements. Defining a reward function that captures a movement's naturalness is difficult. Optimizing for high-level objectives such as maintaining a certain speed or following a certain turning angle does not guarantee the resulting motion being smooth or natural. Generative adversarial imitation learning (GAIL) [Ho and Ermon 2016], incorporated a neural network that provided supervision over the agent's movements, allowing the controller to learn the manifold of natural movements. Such an adversarial approach could produce robust controllers [Fu et al. 2018], yet the generated motions were still not comparable to the kinematic controllers. The framework of generative adversarial networks (GANs) is known to be a powerful estimator of a prior known distribution [Choi et al. 2018; Engel et al. 2019; Zhu et al. 2017]. But it is also known to be unstable and sensitive to the objective function. Besides, a good objective function for natural gait transitions is difficult to define. In summary, DRL approaches struggle with defining an objective function that describes natural movements, directly adding GAN into the mix would disrupt the learning process, causing the controller to produce awkward movements.

3 PROPOSED METHOD

Our goal is to design a physics-based controller that produces natural movements and reactions under external perturbations while following high-level user controls. We consider the generated motion as natural when it resembles the movement of the reference motion. For this purpose, we train the controller in three stages (Figure 2). In the first stage, we aim to transfer the natural movements from the reference motion clips to a physics-based controller through imitation learning (Figure 2(a)). This is achieved by learning the action distribution through a policy network for the physics-based controller to follow. This policy network contains a primitive and gating network that decompose the action distribution into lower-level primitive distributions that can be multiplicatively composed using learned weightings. As a result, the policy network produces an action distribution that enables the physics-based controller to produce natural movements, successfully bridging animation and physics.

To enable the controller to adopt high-level control objectives such as target speed and heading, it would not be feasible to directly optimize for forward progression, as this would end up producing a different action distribution which may result in awkward and unnatural movements. Therefore, in the second training stage (see Figure 2(b)), we adopt a GAN Control Adapter to enable the high-level gating network to approximate the natural action distribution learned previously. However, the lack of external perturbations during the second training stage hinders the controller's ability to adapt to unseen scenarios. Therefore, we add a GAN regularized DRL fine-tuning stage to empower the controller to recover from such scenarios. We further discuss the details of each training stage in Section 3.1, Section 3.2, and Section 3.3.

Fig. 2. Our system is divided into three modules: (a) Low-Level: Imitation Learning, (b) High-Level: GAN Control Adapter, and (c) High-Level: DRL Fine-Tuning. During the fine-tuning process, we fixed the primitive network to prevent it from compensating the high-level gating network's error.

3.1 Low-Level: Imitation Learning

Our goal in this stage is to learn the action distribution for a physics-based controller. To do so, we need the action distribution to describe the articulation of the agent. This can be accomplished through imitation learning where we treat the reference motion as a form of low-level control that specifies the agent's movement at the joint level. For this purpose, we adopt the policy network of [Peng et al. 2019] which consists of two modules: gating and primitive networks (see Figure. 2(a)). The low-level gating network G_{low} takes the agent's current state s_t and the reference motion c_{low} as input and produces primitive influence w ,

$$w = G_{low}(s_t, c_{low}), \quad (1)$$

where $w \in \mathbb{R}^k$, with k being the total number of primitive motions and $c_{low} = (\hat{s}_{t+1}, \hat{s}_{t+2})$ is the joint-level control defined as the target states for the next two time steps of the reference motion. Both s_t and $\hat{s}_t$ contain the information of each joint's position, velocity, rotation, and angular velocity. All the information are represented as 3-dimensional vectors except for the rotations, which are represented as 4-dimensional quaternions. The redundant storage of both position and rotation in the state representation is a widely adopted technique for learning the locomotion skills for torque-articulated characters [Park et al. 2019; Peng et al. 2018; Todorov et al. 2012].

The primitive network P , on the other hand, takes in the agent's current state s_t and decomposes the action distribution into primitive distributions $\phi_1 \dots \phi_k$, where each primitive ϕ_i represents the lower-level action distribution that specializes in modeling a specific locomotion skill within the action space. Each ϕ_i is modeled as a Gaussian distribution with state-dependent action mean $\mu_i(s_t)$ and diagonal covariance matrix Σ_i ,

$$\phi_i = \mathcal{N}(\mu_i(s_t), \Sigma_i), i = 1, 2, \dots, k. \quad (2)$$

Unlike [Peng et al. 2019] which learns the state-dependent covariances, we use a fixed diagonal covariance matrix Σ . We found that learning Σ leads to premature convergence, similar to the findings of [Merel et al. 2019]. As discussed in [Peng et al. 2019], multiplicatively composing the Gaussian primitive distributions ϕ given its

influence w produces a composite distribution that is also a Gaussian distribution

$$\pi(a_{t+1}|s_t, c_t) = \frac{1}{Z(s_t, c_t)} \prod_{i=1}^k \phi_i^{w_i}, \quad (3)$$

where $Z(s_t, c_t)$ denotes a normalization function, and c_t denotes the current control objective $c_t = c_{low}$. The agent's next action a_{t+1} is then sampled from this action distribution. Following [Peng et al. 2018], we use their pose, velocity, and center-of-mass (COM) rewards:

$$R_p = \exp[-2(\sum_j \|\hat{q}_j \ominus q_j\|^2)], \quad (4)$$

$$R_v = \exp[-0.1(\sum_j \|\hat{q}_j - q_j\|^2)], \quad (5)$$

$$R_{com} = \exp[-10(\|\hat{p}_c - p_c\|^2)]. \quad (6)$$

The pose reward R_p encourages the controller to match the reference motion's pose by computing the quaternion difference $\ominus$ between the agent's j -th joint orientations q_j and the target $\hat{q}_j$. The velocity reward R_v computes the difference of joint velocities, where $\hat{q}_j$ and q_j represent the angular velocity for the j -th joint of the agent and the target respectively. The center-of-mass reward R_{com} discourages the agent's center-of-mass p_c to deviate from the target's $\hat{p}_c$.

We replace the end-effector reward with contact point reward R_c ,

$$R_c = \exp[-\frac{\lambda_c}{4}(\sum_e \hat{p}_e \oplus p_e)], \lambda_c = 5, \quad (7)$$

where $\oplus$ denotes the logical XOR operation, p_e denotes the Boolean contact state of the agent's end-effector, and $e \in \{left-front, right-front, left-rear, right-rear\}$. This reward function R_c is designed to discourage the agent when the gait pattern deviates from the reference motion, and to help resolve the foot-sliding artifacts. For instance, the Boolean contact state for only the *left-front* end-effector touches the ground can be denoted by $p = [1, 0, 0, 0]$. Inspired by [Peng et al. 2018], we use the exponential value where λ_c denotes a hyper-parameter that controls the slope of the exponential function, and found that $\lambda_c = 5$ yields the best outcome.

$$R = 0.65R_p + 0.1R_v + 0.1R_{\text{com}} + 0.15R_c, \quad (8)$$

is the final form of the reward function. To reduce the training complexity, we separate the reference motion of each control objective (see Section 4.2). As a result, our physics-based controller can imitate the given reference motion by learning the corresponding action distribution. The controller produces natural movements, performing different gait patterns depicted by the reference motion.

3.2 High-Level: GAN Control Adapter

Our goal for this stage is to enable high-level user control over speed and heading. This allows the user to directly control the agent without going through the tedious and laborious process of specifying reference motions. However, directly optimizing for forward progression or high-level user controls (such as target speed R_{spd} or heading R_{head} defined later in Eq. 12 and Eq. 13, respectively) causes the policy network to ignore the previously learned action distribution, and this eventually leads to awkward and unnatural movements. For the remainder of our discussion, we represent the high-level user control with two values $c_{\text{high}} = (\sigma, \Delta\theta)$, where σ denotes the agent's target speed (m/s), and $\Delta\theta$ denotes the angular difference between the agent's current heading and target heading. For instance, the control for the agent to travel at 1 m/s while rotating 90 degrees counter-clockwise is $c_{\text{high}} = (1, 0.5\pi)$. During training, we consider c_{low} and c_{high} as paired labels where c_{high} is derived from the physical states of the reference motion c_{low} .

Since replacing the reference motion with high-level user controls only affects the gating network, learning the mapping between high-level user control and the primitive influence translates to learning the agent's action distribution. Standard distance functions such as $l1$ or $l2$ only preserve the low-order statistics of a distribution and do not guarantee the samples are drawn from the correct distribution. Therefore, we use the GAN framework which, instead of estimating low-order statistics, approximates the manifold of the distribution through adversarial learning. The GAN framework consists of two modules: generator and discriminator. Given real samples of primitive influence $\mathbf{w}_{\text{real}}$ drawn from real data distribution $\mathbf{w}_{\text{real}} \sim G_{\text{low}}(s_t, c_{\text{low}})$, our high-level gating network G_{high} serves as the generator and produces fake primitive influence $\mathbf{w}_{\text{fake}}$ which is drawn from $\mathbf{w}_{\text{fake}} \sim G_{\text{high}}(s_t, c_{\text{high}})$. Our discriminator D aims to distinguish fake samples $\mathbf{w}_{\text{fake}}$ from the real samples $\mathbf{w}_{\text{real}}$ by maximizing L_{adv} defined in Eq 9.

$$\min_{G_{\text{high}}} \max_D L_{\text{adv}} = \mathbb{E}_{s_t, c_{\text{low}}} [\log D(G_{\text{low}}(s_t, c_{\text{low}}))] + \mathbb{E}_{s_t, c_{\text{high}}} [\log(1 - D(G_{\text{high}}(s_t, c_{\text{high}})))] \quad (9)$$

$$L_{\text{rec}} = \|\mathbf{w}_{\text{fake}} - \mathbf{w}_{\text{real}}\|_1 \quad (10)$$

$$L_G = \lambda_{\text{adv}} * L_{\text{adv}} + \lambda_{\text{rec}} * L_{\text{rec}} \quad (11)$$

We train G_{high} by minimizing the objective function L_G defined in Eq. 11. Reconstruction loss L_{rec} minimizes the absolute distance between samples drawn from real and fake data distributions. Through the adversarial loss L_{adv} , the discriminator D provides supervision to the generated samples by classifying it as real or fake. This guides G_{high} to learn the real data distribution, i.e., the manifold. Inspired

by [Isola et al. 2017], we weigh the importance of each loss by $\lambda_{\text{rec}} = 100$, and $\lambda_{\text{adv}} = 1$. Jointly minimizing the two losses L_{rec} and L_{adv} allows the generator to produce fake samples that are close in terms of distance and also come from the real data distribution. The end result is a control adapter that translates high-level user control over target speed and heading into natural movements.

3.3 High-Level: DRL Fine-Tuning

During the GAN training process, we only expose a small subset of possible scenarios to the controller. This means rarely sampled controls and external perturbations can cause the agent to fall, as the controller may be unable to recover from these unseen scenarios. We solve this problem by further fine-tuning the agent with physical simulations, where we expose the agent to a significantly larger number of scenarios and ask it to recover through trial and error. For example, when being bombarded with objects or when the speed suddenly changes, the controller should learn to shift the agent's body weight to prevent it from falling.

We have observed the tendency for the policy network to change the primitive network to compensate for the gating network's error. As discussed previously in Section 3.2, since high-level user control only affects the primitive influence, we freeze the parameters of the primitive network and only train the gating network to preserve the action distribution. We train the high-level gating network with the corresponding reward function R_{spd} or R_{head} depending on the control objective.

$$R_{\text{spd}} = \exp[-\lambda_{\text{spd}}(\sigma - \|\mathbf{v}\|)^2]; \lambda_{\text{spd}} = 0.8 \quad (12)$$

$$R_{\text{head}} = \left(\frac{\hat{\mathbf{u}} \cdot \mathbf{v}}{\|\hat{\mathbf{u}}\| * \|\mathbf{v}\|} + 1 \right) * 0.5 \quad (13)$$

The speed reward R_{spd} computes the $l2$ -distance between the target speed denoted by σ and the agent's current movement speed $\|\mathbf{v}\|$. We found that using $\lambda_{\text{spd}} = 0.8$ produces the best result in our case. As for the heading reward R_{head} , we compute the cosine similarity between the target heading $\hat{\mathbf{u}} = (\cos(\hat{\theta}), -\sin(\hat{\theta}))$ and agent's heading $\mathbf{v}$ projected onto the plane of motion, with $\hat{\theta}$ representing the target heading in radians. We normalize the value of cosine similarity between 0 and 1.

Lastly, to ensure the controller does not deviate too far from the learned action distribution, we impose a regularization term L_{reg} ,

$$L_{\text{reg}} = \sum_{l=1}^L \|\hat{\alpha}_l - \alpha_l\|_1, \quad (14)$$

where $\hat{\alpha}_l$ denotes the parameters of l -th fully-connected layer of the GAN trained gating network (also used as the initialization point), α_l denotes the parameters of l -th fully-connected layer of the currently trained gating network, and L denotes the total number of layers in each gating network. We apply the regularization in the parameters space, because applying it to the layer's output would penalize the real unseen scenarios. After fine-tuning with DRL, our controller gains the ability to recover from unseen scenarios, enabling smooth transitions following user high-level controls while being adaptive to external perturbations.

Fig. 3. The agent follows the user speed control from 3-4 m/s in a straight direction. (a) MCP (No Control Adapter) and (b) CARL-L1 (L1 Control Adapter) continue to perform trot at unnaturally high cadence. Whereas the (c) CARL-GAN (Ours) and (d) Reference Motion both automatically switch gait pattern from trot to canter as the target speed becomes higher.

4 IMPLEMENTATION

4.1 Data Collection

There are multiple ways to obtain the reference motion dataset, such as kinematic controllers, motion graphs, or even capturing raw motion data. For convenience, we use a readily available kinematic controller¹ provided by [Zhang et al. 2018]. Since our goal is to control both speed and heading, collecting reference motion data containing all the possible movement speeds, turning rates, and turning angles would make the reference motion undesirably complex. As discussed by [Peng et al. 2018], the complexity of the reference motion highly affects the success of the imitation learning process. We record two datasets, one for each speed and heading controls. For each control objective, we record one minute worth of reference motion clip. For speed control, the agent travels towards a fixed direction with various move speeds ranging from 0 to 4 m/s. The agent automatically changes its gait pattern depending on the speed. The agent performs pace at (0, 2) m/s, trot at [2, 3.5) m/s, and canter at [3.5, 4] m/s. As for heading control, the agent performs turning left and right in the range of 0 to 180 degrees. Here, the agent only performs pace gait-pattern as it moves at a fixed speed of 1 m/s.

4.2 Online Optimization with DRL

We train the policy network in a single PC equipped with AMD R9 3900X (12 Cores/24 Threads, clock speed at 3.8GHz). We adopt a physical simulation C++ Bullet physics library [Coumans et al. 2013] for physical simulation, which updates at 1200 frames per second. Our physical simulation modules query the policy network to obtain the action distribution at 30 frames per second. Our policy network learns with proximal policy optimization (PPO) [Schulman et al. 2017] and generalized advantage estimation (GAE) [Schulman et al. 2015b]. We follow [Peng et al. 2019]'s policy network architecture, number of primitives ($k = 8$), and action space representation. The action at time t , $a_t \in A$, is represented as PD-targets placed at each joint. However, instead of learning the action-noise variance, we use a fixed diagonal covariance matrix $\Sigma = 0.05I$, where I is the identity matrix. The value function uses multi-step returns with

¹<https://github.com/sebastianstarke/AI4Animation>

TD(λ) [Sutton et al. 1998]. The learning process is episodic, wherein each episode we incorporate early termination and reference state initialization proposed by [Peng et al. 2018].

In low-level imitation learning, we separate the learning process of speed and heading controls, each imitating the corresponding recorded reference motion clip as mentioned in Section 4.1. We then have two separate policy networks: one for speed control and the other for heading control. Learning rates of the policy networks and value function networks are $2.5e-6$ and $1e-2$ respectively. Hyperparameters for GAE(λ), TD(λ), and the discount factor are set to 0.95. The learning process takes approximately 250 hours for each policy network.

In the high-level DRL fine-tuning step (Section 3.3), the two policy networks optimize for different reward functions. The policy network for the move speed task uses R_{spd} (defined in Eq. 12) because the agent only moves towards a single direction. The policy network for the heading task only uses R_{head} (defined in Eq. 13) because the agent moves at the same speed. During training, we randomly update the control objective with an offset uniformly sampled from $[-0.25, 0.25]$ m/s and $[-0.15, 0.15]$ radian for speed and heading respectively. The updates happen at 4 frames per second. To encourage the controller to learn the rarely sampled gait transitions, we further introduce a 10% probability that the control objective being completely altered (e.g. from cantering at 4 m/s to standing at 0 m/s). The learning rates for the policy networks and value function networks are $5e-5$ and $1e-2$ for both speed and heading respectively, with 0.99 for the discount factor.

Fig. 4. The learning curves of (left) speed and (right) heading control. MCP (No Control Adapter), CARL-L1 (L1 Control Adapter), and CARL-GAN (Ours).

Table 1. The MSE of speed at different gait patterns: pace, trot, and canter.

Gait	MCP	CARL-L1	CARL-GAN
Pace	$1.7\text{e-}3 \pm 0.2\text{e-}3$	$17.7\text{e-}3 \pm 2.5\text{e-}3$	$1.6\text{e-}3 \pm 0.1\text{e-}3$
Trot	$5.7\text{e-}3 \pm 1.9\text{e-}3$	$3.2\text{e-}3 \pm 0.8\text{e-}3$	$0.5\text{e-}3 \pm 0.3\text{e-}3$
Canter	$16.4\text{e-}3 \pm 1.2\text{e-}3$	$28.1\text{e-}3 \pm 10.6\text{e-}3$	$4.1\text{e-}3 \pm 1.1\text{e-}3$
Avg.	$7.9\text{e-}3 \pm 1.1\text{e-}3$	$16.3\text{e-}3 \pm 4.6\text{e-}3$	$2.1\text{e-}3 \pm 0.5\text{e-}3$

The GAN regularization (defined in Eq. 14) is added to the DRL’s PPO clip surrogate loss as an additional term multiplied by 0.001. Our policy networks converge after approximately 25 hours for training the speed control and 40 hours for training the heading control. Alternatively, directly learning the high-level user control without GAN Control Adapter, the policy network converges after approximately 100 hours for training each of the control objectives. The learning curves are shown in Figure 4.

4.3 Offline optimization with GAN

The training process involves collecting one million data points consisting of the agent’s state, reference motions, and high-level user controls (derived from two consecutive reference motions). During training, the real samples are drawn from the real data distribution by passing the agent’s state and high-level user control to the low-level gating network.

We implement the discriminator as a series of 3-layered fully-connected neural networks with the LeakyReLU activation function in between the intermediate layers excluding the last classification layer. It optimizes for the original GAN objective where the discriminator outputs the probability of the input samples being real, as formulated by [Goodfellow et al. 2014]. Inspired by [Isola et al. 2017; Zhu et al. 2017] we added the reconstruction loss L_{rec} (defined in Eq. 14). The training process converges in approximately 1 hour (50 epochs) using Adam optimizer with a learning rate of $2\text{e-}5$.

5 EXPERIMENT

In this section, we evaluate the effectiveness of using the GAN Control Adapter to enable user’s speed and heading controls over the agent. For this purpose, we consider two additional algorithms in addition to our proposed algorithm (**CARL-GAN**). Removing the GAN Control Adapter component of our method makes it equivalent to **MCP** [Peng et al. 2019] that directly applies DRL for high-level user control. **CARL-L1** uses a standard distance function $L1$ instead of GAN as Control Adapter (i.e., $L1$ Control Adapter) to minimize the absolute distance between the output of the high-level gating network and the previously learned primitive influence.

Our evaluations focus on measuring the controller’s ability to follow the user control in addition to the quality of the motion. We use the motions generated from the kinematic controller [Zhang et al. 2018] as the ground truth. To measure the motion quality, we use an approach from robotics [Dholakiya et al. 2019; Moro et al. 2012], where the range of joint movements defines motion similarity. The visual results are shown in Figure 3.

Next, we discuss the quantitative results. A controller’s ability to follow user control is shown in Section 5.1; quality of generated motions is shown in Section 5.2; visualization and exploration of the action distribution are shown in Section 5.3; and discussion

Table 2. Deviations from the reference motions, measured in degree ($^\circ$) for angular deviation, and meter (m) for positional deviation.

Type	Control		MCP	CARL-L1	CARL-GAN
Angular (°)	180°	L	53.68 ± 6.41	17.93 ± 1.50	14.87 ± 0.75
		R	64.30 ± 8.32	26.78 ± 12.63	15.36 ± 0.65
	90°	L	46.27 ± 2.60	18.05 ± 1.23	15.42 ± 0.60
		R	44.54 ± 3.03	22.34 ± 3.03	13.04 ± 1.28
Position (m)	180°	L	0.95 ± 0.06	0.31 ± 0.03	0.25 ± 0.05
		R	1.18 ± 0.08	0.54 ± 0.21	0.19 ± 0.07
	90°	L	1.26 ± 0.07	0.31 ± 0.04	0.21 ± 0.01
		R	1.31 ± 0.09	0.38 ± 0.07	0.19 ± 0.01

about our advantages compared to the Kinematic Controller + DRL approach in Section 5.4.

5.1 Control Accuracy

To evaluate the controller’s ability in following commands, we produce 10 recordings, each with 5 seconds of duration, for each control objective. For speed control, the agent travels in a fixed direction with different gait patterns, e.g., pace, trot, and canter. Similarly for heading control, the agent travels at a constant speed (1 m/s) while performing 90- and 180-degree turns, both clockwise and counter-clockwise.

Speed Control. Table 1 shows our algorithm’s ability to follow the target speed at different gait patterns. The results show a significantly lower mean squared error (MSE) of our algorithm CARL-GAN compared to both baselines: MCP and CARL-L1. As expected, learning the target speed directly (MCP) allows the algorithm to converge better compared to L1 Control Adapter (CARL-L1). Note that the MSE of the canter gait pattern is greater than other gait patterns across all algorithms, due to its larger range of motion.

Heading Control. To investigate the controller’s responsiveness to the user’s heading control, we ask the controller to follow predefined trajectories of turning left-90 $^\circ$, right-90 $^\circ$, left-180 $^\circ$, right-180 $^\circ$, and calculate their respective angular and positional deviations from the reference motions (Table 2).

MCP’s controller ignores the previously learned gait patterns and instead overfits the target heading, resulting in relatively high deviations. CARL-L1 performs markedly better than MCP because the control adapter preserves the learned gait patterns. Results from CARL-L1 is a bit biased with better MSE in left turns than right turns, due to slight left-turn and right-turn training data imbalance. Finally, our proposed algorithm CARL-GAN performs the best in all turning angles. It handles turns well even for drastic angles while following the target speed.

5.2 Motion Quality

To measure the quality of the generated motions, we ask the controller to perform one gait pattern at a time and record its range of movements in the form of end-effector regions (Figure 5). Along the line of robotics research [Dholakiya et al. 2019; Moro et al. 2012], we measure the similarity of movements by calculating the Intersection over Union (IoU) of the end-effectors between the generated

Fig. 5. An end-effector moves in a cycle. Left-to-right: MCP (No Control Adapter), CARL-L1 (L1 Control Adapter), CARL-GAN (Ours), and Reference Motion.

Fig. 6. We apply a dimensional reduction technique (PCA) to visualize the action distribution in 2D. The visualization for each gait pattern forms a circular shape manifold corresponding to the agent's cyclic motion.

motion and the reference motion (Table 3). The results show that our controller achieves the highest IoU score consistently across gait patterns.

These quantitative results translate fairly well to the visual motion quality in Figure 3. For a reference motion that transitions from trot to canter (Figure 3(d)), the controllers in both MCP and CARL-L1 continue to perform trot at unnaturally higher cadences instead of switching to canter gait-pattern. In contrast, our controller can map the high-level user speed control to the previously learned action distribution, performing trot-to-canter transition correctly with natural movements. The high IoU score when performing canter in Table 3 further highlights this property.

5.3 Action Space Exploration

We find it useful to visualize the action distribution for assessing the quality of motions. To this end, we collect millions of data points consisting of the agent's state in different gait patterns, and apply standard principal component analysis (PCA) to visualize the data in 2D. The red dots in Figure 6 represent the action distribution obtained from the imitation learning process. MCP (orange dots) produces a different distribution, which provides a visual clue on why it fails to learn the mapping between high-level user control and the extracted action distribution. CARL-L1 (green dots) preserves the shape of the action distribution, but the manifold appears to be expanded. Ours (blue dots) successfully captures the original action distribution, which indicates the effectiveness of the GAN Control Adapter in translating high-level user controls to the action distribution.

Table 3. Intersection over Union (IoU) of an agent's end-effector while performing pace, trot, and canter.

Gait	Leg	MCP	CARL-L1	CARL-GAN
Pace	L.Front	0.03 ± 0.01	0.04 ± 0.01	0.09 ± 0.01
	R.Front	0.09 ± 0.02	0.14 ± 0.03	0.15 ± 0.02
	L.Rear	0.35 ± 0.02	0.37 ± 0.02	0.42 ± 0.01
	R.Rear	0.02 ± 0.04	0.58 ± 0.02	0.51 ± 0.01
	Avg.	0.12 ± 0.01	0.28 ± 0.02	0.29 ± 0.01
Trot	L.Front	0.34 ± 0.02	0.65 ± 0.01	0.55 ± 0.01
	R.Front	0.34 ± 0.01	0.45 ± 0.04	0.61 ± 0.05
	L.Rear	0.11 ± 0.02	0.17 ± 0.01	0.44 ± 0.03
	R.Rear	0.16 ± 0.01	0.49 ± 0.01	0.62 ± 0.01
	Avg.	0.24 ± 0.03	0.44 ± 0.02	0.56 ± 0.03
Canter	L.Front	0.37 ± 0.02	0.41 ± 0.02	0.70 ± 0.03
	R.Front	0.19 ± 0.01	0.29 ± 0.02	0.77 ± 0.03
	L.Rear	0.33 ± 0.03	0.47 ± 0.03	0.69 ± 0.04
	R.Rear	0.21 ± 0.01	0.28 ± 0.02	0.32 ± 0.03
	Avg.	0.27 ± 0.02	0.36 ± 0.02	0.62 ± 0.03

5.4 Kinematic Controller + DRL

Recent works from [Bergamin et al. 2019; Park et al. 2019] provide an end-to-end solution by manipulating the physics-based controller through a kinematic controller, i.e., Kinematic Controller + DRL. This approach can produce an accurate controller, high-quality motion, and is responsive, but requires access to action labels or a kinematic controller. In contrast, our method does not have such requirements, since we enable the controllable property by directly imposing high-level control to the physics-based controller. However, this causes an inherent problem of the distributional shift, which we address using GAN.

Given a fixed control setting, the Kinematic Controller (MANN [Zhang et al. 2018]) + DRL is equivalent to our imitation learning, which produces high quality motion as highlighted by the manifold in the last column of Figure 6. As a comparison, we ask each controller to follow a set of speed control (dotted line in Figure 7), from which the corresponding action labels are derived for the Kinematic Controller (e.g., 0 m/s for stand and 4 m/s for canter). When performing an unseen transition (4 m/s to 0 m/s), the Kinematic Controller + DRL fails to maintain balance and collapses to the ground due to the lack of such motion in the training data, as highlighted by the pink line in Figure 7. To include all of the possible transitions, the training would require more reference motion data, which leads to exponentially longer training time, as suggested by [Park et al. 2019] in their ablation study. Despite using more reference data and training for a longer time, there is no guarantee that the controller would converge. In summary, via GAN, our method does not require access to either action labels or kinematic controllers and requires less reference motion data to model the gait transitions, therefore converges faster.

Fig. 7. An example of speed control comparison given seen and unseen transitions.

Fig. 8. Attaching path-finding (left) and ray-sensor (right) navigation modules to the controller.

6 APPLICATIONS

6.1 Modular Control

Because the controller is trained to accurately follow the user's speed and heading controls, it becomes fairly straightforward to add complex navigation capability. The fact that the locomotion controller is decoupled from the navigation module, allows us to reuse pre-trained policy networks without compromising the desirable properties of the locomotion controller in complex navigation scenarios. Consider a navigation task where the controller needs to collect valuable objects scattered within a maze. In Figure 8 (left), we implement a path-finding algorithm that reads the maze as a 2D map and outputs an optimal path for the controller to follow. In another task depicted in Figure 8 (right), the controller needs to rely on its sensors and observations to locate and collect the valuable objects. For this problem, we attach a ray-sensor and implement a classical navigation module. A simple translator module then converts the path into sequences of speed and heading controls. For instance, turning 90 degrees left at 1 m/s is represented as $c_{\text{high}} = (1, 0.5\pi)$. The result is an autonomous agent that can efficiently solve the navigation task by following the sequence of controls.

6.2 Reacting to External Perturbations

Our physics-based controller can effectively model the agent's interaction with a dynamically changing environment. To demonstrate this, we expose our agent to multiple unexpected external perturbations. The first perturbation involves boxes with various volumes and density from random directions. The second perturbation is a slippery terrain with the ground's friction reduced by 90%. Lastly, we introduce a rotating and tilting terrain (in the video). Our experiment shows that the controller can react to the given perturbations, maintaining its balance while moving naturally and learn the importance of maintaining its height without being explicitly trained with the corresponding reward function (see Figure. 9).

7 CONCLUSION

In this paper, we present CARL, a quadruped agent that can respond naturally to high-level controls in dynamic physical environments. The three-stage process begins with a low-level imitation learning process to extract the natural movements perceived in the authored or captured animation clips. The GAN Control Adapter maps the high-level directive controls to action distributions that correspond to the animations. Further fine-tuning the controller with DRL enables it to recover from external perturbations while producing smooth and natural actions. The GAN Control Adapter proves instrumental in enabling the controller to respond to high-level user

Fig. 9. Our physics-based controller can produce meaningful reactions to randomly thrown boxes (left), and slippery terrain (right).

controls, all while complying with the original action distribution, as demonstrated in the experiment where the GAN Control Adapter was removed or replaced with a simple L1 adapter. We show the usefulness of our method by attaching navigation modules over the controller to enable it to operate autonomously for tasks such as traversing through mazes with goals. In addition to that, we also create an animation clip, where the agent-environment interaction changes dynamically. Equipped with natural movements, controllable, adaptive properties, our approach is a powerful tool in accomplishing motion synthesis tasks that involve dynamic environments.

In the future, we would like to investigate on imposing constraints over the manifold's morphology, such as shape and size, for controlling the movement style and transitions. Another interesting direction is to use a dynamic lambda to regularize the fine-tuning stage, as a fixed lambda may lead to a trade-off between realism and adaptiveness. We hope this research in combining the merits of GAN and DRL can inspire future research directions towards high realism controllers.

ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for the insightful comments, Yi-Chun Chen for the help in producing the supplementary video, and our colleagues for the meaningful discussions.

REFERENCES

- Trapit Bansal, Jakub Pachocki, Szymon Sidor, Ilya Sutskever, and Igor Mordatch. 2018. Emergent Complexity via Multi-Agent Competition. In *International Conference on Learning Representations*.
- Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. 2019. DReCon: Data-Driven Responsive Control of Physics-Based Characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 206.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI Gym. *arXiv preprint arXiv:1606.01540* (2016).
- Jinxiang Chai and Jessica K Hodgins. 2005. Performance Animation from Low-dimensional Control Signals. *ACM Transactions on Graphics (TOG)* 24, 3 (2005), 686–696.
- Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makovychuk, and Stefan Jeschke. 2018. Physics-Based Motion Capture Imitation with Deep Reinforcement Learning. In *Proceedings of the 11th Annual International Conference on Motion, Interaction, and Games*. ACM, 1.
- Yunjey Choi, Minje Choi, Munyoung Kim, Jung-Woo Ha, Sunghun Kim, and Jaegul Choo. 2018. StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 8789–8797.
- Stelian Coros, Philippe Beaudoin, and Michiel Van de Panne. 2009. Robust task-based control policies for physics-based characters. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 170.
- Stelian Coros, Philippe Beaudoin, and Michiel Van de Panne. 2010. Generalized biped walking control. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 130.
- Stelian Coros, Andrej Karpathy, Ben Jones, Lionel Reveret, and Michiel Van De Panne. 2011. Locomotion skills for simulated quadrupeds. In *ACM Transactions on Graphics (TOG)*, Vol. 30. ACM, 59.
- Erwin Coumans et al. 2013. Bullet physics library. *Open source: bulletphysics.org* (2013).

- Marco Da Silva, Yeui Abe, and Jovan Popović. 2008. Simulation of human motion data using short-horizon model-predictive control. In *Computer Graphics Forum*, Vol. 27. Wiley Online Library, 371–380.
- Dhaivat Dholakiya, Shounak Bhattacharya, Ajay Gunalan, Abhik Singla, Shalabh Bhatnagar, Bharadwaj Amrutur, Ashitava Ghosal, and Shishir Kolathaya. 2019. Design, development and experimental realization of a quadrupedal research platform: Stoch. In *2019 5th International Conference on Control, Automation and Robotics (ICCAR)*. IEEE, 229–234.
- Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. 2019. Gansynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710* (2019).
- Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. 2015. Recurrent Network Models for Human Dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*. 4346–4354.
- Justin Fu, Katie Luo, and Sergey Levine. 2018. Learning robust rewards with adversarial inverse reinforcement learning. In *International Conference on Learning Representations*.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative Adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- Keith Grochow, Steven L Martin, Aaron Hertzmann, and Zoran Popović. 2004. Style-based inverse kinematics. In *ACM transactions on graphics (TOG)*, Vol. 23. ACM, 522–531.
- Perttu Hämläinen, JooSe Rajamäki, and C Karen Liu. 2015. Online control of simulated humanoids using particle belief propagation. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 81.
- Nicolas Heess, Srinivasan Sriram, Jay Lemmon, Josh Merel, Greg Wayne, Yuval Tassa, Tom Erez, Ziyu Wang, SM Eslami, Martin Riedmiller, et al. 2017. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286* (2017).
- Jonathan Ho and Stefano Ermon. 2016. Generative adversarial imitation learning. In *Advances in neural information processing systems*. 4565–4573.
- Daniel Holden, Taku Komura, and Jun Saito. 2017. Phase-Functioned Neural Networks for Character Control. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 42.
- Daniel Holden, Jun Saito, and Taku Komura. 2016. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 138.
- Ting-Chieh Huang, Yi-Jheng Huang, and Wen-Chieh Lin. 2013. Real-time Horse Gait Synthesis. *Computer Animation and Virtual Worlds* 24, 2 (2013), 87–95.
- Jemin Hwangbo, Joonho Lee, Alexey Dosovitskiy, Dario Bellicoso, Vassilios Tsounis, Vladlen Koltun, and Marco Hutter. 2019. Learning agile and dynamic motor skills for legged robots. *Science Robotics* 4, 26 (2019), eaau5872.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image Translation with Conditional Adversarial Networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- Lucas Kovar and Michael Gleicher. 2004. Automated Extraction and Parameterization of Motions in Large Data Sets. In *ACM Transactions on Graphics (TOG)*, Vol. 23. ACM, 559–568.
- Lucas Kovar, Michael Gleicher, and Frédéric Pighin. 2002. Motion Graphs. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '02)*. ACM, 473–482.
- Jehee Lee, Jinxiang Chai, Paul SA Reitsma, Jessica K Hodgins, and Nancy S Pollard. 2002. Interactive control of avatars animated with human motion data. In *ACM Transactions on Graphics (TOG)*, Vol. 21. ACM, 491–500.
- Kyungho Lee, Seyoung Lee, and Jehee Lee. 2018. Interactive character animation by learning multi-objective control. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–10.
- Yoonsang Lee, Sungeun Kim, and Jehee Lee. 2010. Data-driven biped control. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 129.
- Sergey Levine and Jovan Popović. 2012. Physically Plausible Simulation for Character Animation. In *Proceedings of the 11th ACM SIGGRAPH/Eurographics conference on Computer Animation*. Eurographics Association, 221–230.
- Sergey Levine, Jack M Wang, Alexis Harauz, Zoran Popović, and Vladlen Koltun. 2012. Continuous character control with low-dimensional embeddings. *ACM Transactions on Graphics (TOG)* 31, 4 (2012), 28.
- Libin Liu and Jessica Hodgins. 2017. Learning to schedule control fragments for physics-based characters using deep q-learning. *ACM Transactions on Graphics (TOG)* 36, 3 (2017), 29.
- Libin Liu, KangKang Yin, and Baining Guo. 2015. Improving Sampling-based Motion Control. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 415–423.
- Libin Liu, KangKang Yin, Michiel van de Panne, Tianjia Shao, and Weiwei Xu. 2010. Sampling-based contact-rich motion control. In *ACM Transactions on Graphics (TOG)*, Vol. 29. ACM, 128.
- Josh Merel, Arun Ahuja, Vu Pham, Saran Tunyasuvunakool, Siqi Liu, Dhruva Tirumala, Nicolas Heess, and Greg Wayne. 2019. Hierarchical Visuomotor Control of Humanoids. In *International Conference on Learning Representations*.
- Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. 2018. Neural probabilistic motor primitives for humanoid control. *arXiv preprint arXiv:1811.11711* (2018).
- Federico Lorenzini Moro, Nikos G Tsagarakis, and Darwin G Caldwell. 2012. On the Kinematic Motion Primitives (kMPs) - Theory and Application. *Frontiers in neurorobotics* 6 (2012), 10.
- Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. 2019. Learning Predict-and-Simulate Policies From Unorganized Human Motion Data. *ACM Transactions on Graphics (TOG)* 38, 6, Article 205 (2019).
- 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), 143.
- Xue Bin Peng, Glen Berseth, and Michiel Van de Panne. 2015. Dynamic terrain traversal skills using reinforcement learning. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 80.
- Xue Bin Peng, Glen Berseth, and Michiel Van de Panne. 2016. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 81.
- Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. 2019. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. In *NeurIPS*.
- Alla Safonova and Jessica K Hodgins. 2007. Construction and optimal search of interpolated motion graphs. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 106.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015a. Trust region policy optimization. In *International conference on machine learning*. 1889–1897.
- John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. 2015b. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438* (2015).
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- Hyun Joon Shin and Jehee Lee. 2006. Motion synthesis and editing in low-dimensional spaces. *Computer Animation and Virtual Worlds* 17, 3-4 (2006), 219–227.
- Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. 2019. Neural state machine for character-scene interactions. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–14.
- Richard S Sutton, Andrew G Barto, et al. 1998. *Introduction to reinforcement learning*. Vol. 2. MIT press Cambridge.
- Yuval Tassa, Tom Erez, and Emanuel Todorov. 2012. Synthesis and stabilization of complex behaviors through online trajectory optimization. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4906–4913.
- Yuval Tassa, Nicolas Mansard, and Emo Todorov. 2014. Control-limited differential dynamic programming. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1168–1175.
- 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, 5026–5033.
- Kevin Wampler and Zoran Popović. 2009. Optimal gait and form for animal locomotion. In *ACM Transactions on Graphics (TOG)*, Vol. 28. ACM, 60.
- Kevin Wampler, Zoran Popović, and Jovan Popović. 2014. Generalizing locomotion style to new animals with inverse optimal regression. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 49.
- Jungdam Won and Jehee Lee. 2019. Learning body shape variation in physics-based characters. *ACM Transactions on Graphics (TOG)* 38, 6 (2019), 1–12.
- Jungdam Won, Jongho Park, Kwanyu Kim, and Jehee Lee. 2017. How to train your dragon: example-guided control of flapping flight. *ACM Transactions on Graphics (TOG)* 36, 6 (2017), 198.
- Jungdam Won, Jungnam Park, and Jehee Lee. 2018. Aerobatics control of flying creatures via self-regulated learning. *ACM Transactions on Graphics (TOG)* 37, 6 (2018), 1–10.
- Yuting Ye and C Karen Liu. 2010. Optimal feedback control for character animation using an abstract model. *ACM Transactions on Graphics (TOG)* 29, 4 (2010), 74.
- KangKang Yin, Kevin Loken, and Michiel Van de Panne. 2007. Simbicon: Simple biped locomotion control. In *ACM Transactions on Graphics (TOG)*, Vol. 26. ACM, 105.
- He Zhang, Sebastian Starke, Taku Komura, and Jun Saito. 2018. Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 145.
- Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. 2018. Auto-Conditioned Recurrent Networks for Extended Complex Human Motion Synthesis. (2018).
- Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. 2017. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE international conference on computer vision*. 2223–2232.