Advancing Ballet Performance with Physics-Based Character Animation Using Deep Reinforcement Learning

By Christian Kim

Author Biography

Christian J. Kim is a junior at Korea International School. He leads the VEX Robotics Club and the STEM Research Journalism Club at his school. A ballet dancer and art enthusiast, he enjoys exploring the intersection of art and technology and hopes to pursue a field that bridges creativity and innovation.

Abstract

Deep reinforcement learning methods in physics-based character animation have diverse application potential for ballet performance including high-fidelity movement replication, enhanced pose estimation, and creative choreography prototyping. First, a current overview of foundational deep reinforcement learning frameworks is presented, primarily based on latest developments including DeepMimic, Residual Force Control (RFC), Adversarial Motion Priors (AMP), and Simulated Character Control for 3D Human Pose Estimation (SimPoE). DeepMimic applies reference motion data for learning control policies and can synthesize a high-fidelity of the corresponding complex motion. RFC resolves dynamics mismatches with the help of external residual forces that allow it to perform certain motions. AMP is a combination of reinforcement learning and generative adversarial networks that helps in policy learning from scarce data, while SimPoE is the integration of physics-based control with the 3D pose estimator for realism. The research focuses on exploring compatible uses of these technologies in ballet performance, from fine tuning of methods to mimic the exacting nature of artistic ballet movements to concerns of scalability, data integrity, and ethical considerations.

Keywords: Physics-Based Animation, Deep Reinforcement Learning, DeepMimic, Adversarial Motion Priors, Residual Force Control, Simulated Character Control for 3D Human Pose Estimation, Ballet Performance, Motion Control, Virtual Characters, Robotics

Introduction

The rapidly advancing field of animation and robotics has long aspired to craft lifelike simulated agents capable of moving with the agility exhibited across the natural world. Reinforcement learning formulates control as an optimization problem where agents aim to determine actions that maximize long-term cumulative rewards received from the environment (Peng et al., 2018; Heess et al., 2017). This shifts the burden of controller design in physics-based character animation techniques towards reward function specification, which offers greater flexibility in encoding desired behaviors (Sutton et al., 2018).

Deep reinforcement learning has emerged as a promising paradigm for automating the acquisition of sophisticated motor skills by learning from repeated trial-and-error interactions with high-fidelity physics simulators. In particular, the DeepMimic framework introduced the utilization of reference motion capture clips to guide the learning process (Peng et al., 2018). This is done by initializing simulations near recorded poses, providing early termination to avoid unnatural behaviors, and shaping rewards to track motion styles in addition to task objectives. The effectiveness of this approach is demonstrated through experiments that showcase learned controllers replicating challenging skills like kicking and spinning, with the resulting motions exhibiting fluidity and grace closely resembling the original human performances (Peng et al., 2018; Hess et al., 2018). To demonstrate the features of the DeepMimic framework, a humanoid agent is shown in the below Figure 1 (Peng et al., 2018) to perform accurate and sophisticated movements derived from the reference motion capture data.

Figure 1. Reference motion captures using DeepMimic highlighting a humanoid agent performing movements.

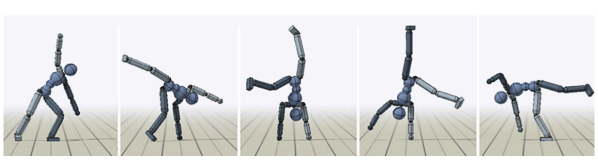

More recent studies show the integration of new frameworks such as learned locomotion controllers for real-time simulations through Residual Force Control (RFC) for characters with moving parts to maintain consistency, Dual-Policy for physically based movements, Adversarial Motion Priors (AMP) by using generative adversarial networks (GAN) for natural and expressive animations, and Simulated Character Control for 3D Human Pose Estimation (SimPoE) based on kinematic estimation and physics based dynamics for 3D human pose estimation (Peng et al., 2021; Yuan & Kitani, 2020; Yuan et al., 2021; Escontrela et al., 2022).

This literature review aims to highlight different reinforcement learning frameworks enhancing the development of using physics to characterize character animation, especially applicable to ballet performance. Key questions include: In what manner do these frameworks enhance character animation? What are the drawbacks and which methods can improve the accuracy of the motion in a virtual environment for ballet performance? As an overview, the following table summarizes the comparison of various methods and techniques within deep reinforcement learning for character animation to gain an enhanced look at what researchers have designed.

Table 1. Comparison Table of Recent Chronological Developments in Motion Synthesis

Approach	Key Features	Strengths	Limitations
DeepMimic (Peng et al., 2018)	Reference motion tracking, Proximal Policy Optimization (PPO)	High-fidelity reproduction of reference motions	Requires large amounts of reference data
RFC (Yuan et al., 2020)	Residual force incorporation	Improved agile motion imitation	May introduce unrealistic forces
AMP (Peng et al., 2021)	Adversarial training, Generative Adversarial Network (GAN)	Sample-efficient, diverse motions	May struggle with long-term planning
Advanced AMP Framework (Escontrela et al., 2022)	Reinforcement Learning + GAN for motion priors	State-of-the-art performance on dynamic skills	Potential instability in adversarial training
SimPoE (Yuan et al., 2021)	Integration of physics-based control and 3D pose estimation	Enhanced accuracy and realism of simulated motions	Computational complexity of integrating physics simulation with pose tracking

The ability of these techniques to reconstruct ballet's rigid precision and aesthetics will be examined in the later sections, along with suggestions for improvement. The purpose is to underscore the opportunities from deep reinforcement learning in improving the future of character animation and specific applications for ballet performance.

Overview of Reinforcement Learning Models for Physics-Based Character Animation

Overall, most of the Deep Reinforcement Learning frameworks learn complex strategies from the tried behaviors and their reward system, which makes them suitable for activities such as navigation and object handling (Peng et al., 2021) . Deep reinforcement learning works best in applications related to continuous control and mostly involves learning features and dependency mapping. Algorithms like Deep Deterministic Policy Gradients (DDPG), Trust Region Policy Optimization (TRPO), and Proximal Policy Optimization (PPO) methods are employed to regulate policies for walking and grasping (Peng et al., 2017; Huo, 2018; Yuan & Kitani, 2020).

Motor control tasks described as Markov Decision Processes (MDP) include the joint positions manipulating velocities, and torques (Jiang et al., 2019; Yang & Yin, 2021; Li et al., 2021). Such methods as Dynamic Time Warping (DTW) and Fast Fourier Transform (FFT) are used to warp and align reference and simulate motions. Exploration strategies like Gaussian noise and maximum entropy reduce the chances of early convergence (Plappert et al., 2017; Haarnoja et al., 2018). The use of parallel computing and physics simulation has expanded the deep reinforcement learning training to high-level motion policies for the use in matters like sports movements (Won et al., 2021; Liang et al., 2018). Frameworks such as OpenAI Gym and DeepMind Control Suite can be used for movement evaluation (Brockman et al., 2016; Carnegie Mellon University, 2021).

Virtual Character Animation through the DeepMimic Framework

Figure 2. Overview of the DeepMimic Training Process with Key Components (Peng et al., 2018)

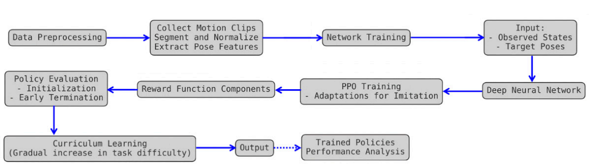

The control policies that are interposed in the Deep Mimic Framework entails the use of the deep neural network to characterize joint torques for training and involve the usage of the modified proximal policy optimization (PPO) (Peng et al., 2018). The reward function checks the reference motion, balance and contact dynamics and the states in the surrounding of the target pose are preferred to get rid of unwanted poses. Curriculum learning is the approach where the training starts from a particular position and within the first 10-30 million simulation steps the task difficulty is gradually raised. Based on the recorded result, DeepMimic enables learning of the dynamic motion skills with the average position error of joint position below 5 cm.

Currently, the DeepMimic framework has the following future development areas. Reward functions must be created manually, and it is not clear if it contains all the nuances of natural movements (Peng et al., 2018). New biomechanical principles or a concept called human pose estimation may help raise the quality of character animations. In addition, DeepMimic learns policies for individual motion clips in isolation. Developing methods to compose and transition between multiple learned skills seamlessly would enable more flexible and interactive character control. DeepMimic also requires substantial reference motion data to learn policies, which serves as a bottleneck when implementing diverse character-based animation. Despite these improvement areas, the DeepMimic framework critically enhanced deep reinforcement learning for the physics-based character animation and paved the way for the effective creation of computer animations and interactive applications (Peng et al., 2018).

Key Recent Developments in Physics-based Character Animation and Critical Analyses

Residual Force Control

Recent developments have moved towards simulating the character animations which depend on physical forces. Yuan et al. (2020) developed Residual Force Control (RFC), which employs residual forces to improve the realism of humanoid models while reducing the necessity of repeated simulations. This makes the method suitable for mimicking intricate motions like ballet and poses higher accuracy through the versatility of the learning approach. RFC can overcome discrepancies in physical characteristics and control guidelines, enabling delicate and complex movements. However, having an approach that depends on external stimulus hinders the functional capability of RFC, limiting its effects within the simulated environments excluding real world robotics (Peng et al., 2020; Cheng et al., 2024). Further adjustment of regularization terms is required to preserve the physical realism of generated frames without losing accurate representation. While RFC is superior to other methods in agile movement tasks, problems occur when the external forces are unpredictable and uncontrollable.

Adversarial Motion Priors (AMP) and Sample-Efficient Learning

Peng et al. (2021) have developed one of the recent modern algorithms of deep reinforcement learning known as the Adversarial Motion Priors (AMP), which applies both reinforcement learning and generative adversarial networks to use limited demonstration data efficiently. AMP reduces sampling noise and utilizes the discriminator network to make the distinction between the reference motion and the motion from policy generation. Simulation results reveal that even in dynamic parkour movements, AMP can achieve near-perfect results, comparable to current state-of-the-art methods with much less amount of training data. For instance, AMP can acquire skills such as the wall run and backflips from mere 5-10 mins of motion capture data (Peng et al., 2021). Nevertheless, there is an identified weakness concerning highly structured planning tasks, highlighting a challenge in sample data processing optimization. Although the AMP method uses GANs with reinforcement learning to enable learning of complex novel movements with a small amount of training data, the applications can be problematic in tasks that require a lot of future planning. Despite the challenges, AMP is a significant step forward in large-scale motion synthesis (Won et al., 2020; Peng et al., 2021; Escontrela et al., 2022).

SimPoE: Simulated Character Control for 3D Human Pose Estimation

In the same year as the introduction of AMP, researchers Yuan et al. (2021) proposed a method called SimPoE, which combines physics-based character control with the regression of 3D human pose. This approach involves the use of real pose data to navigate the physics-based simulation adding the realism of the movements in motion capture. SimPoE improves some of the shortcomings of DeepMimic by providing a more direct link between observed human motions and simulated characters, enabling better generalization to novel movements not seen in the training data. It also enhances the capability in controlling the non-periodic or complex motions which are advantageous to depict diverse characters with complex movements. Nevertheless, the integration of physics-based simulation with pose tracking may impose some difficulties when it comes to real-time processing because it requires extensive computation, limiting its applications to exclude environments with compute restraints (Yuan et al., 2021).

Potential Applications in Ballet Performance

With the help of newly developed physics-based character animation technologies, as exemplified by frameworks like DeepMimic, Residual Force Control, and Dual-policy Control (Peng et al., 2018; Yuan & Kitani, 2020; Peng et al., 2021; Yuan et al., 2021; Escontrela et al., 2022), exciting possibilities for applications in domains that require realistic and expressive human motion have opened up. One area is the potential use of these techniques in the field of ballet, which is known for its highly technical and artistic movement vocabulary. Ballets can be considered as one of the most specific and interesting fields of using physics-based animation because it involves dynamism, elegance, and the absence of apparent physical exertion.

To understand what deep reinforcement learning could be used for, in the case of ballet animation, it is pertinent to consider the nature of this art form. Most ballet movements are very controlled, and the dancers tend to aim for different positions, neat, smooth, gradual changes in different steps among others. The technique is compound as it involves the application of strength, flexibility, and coordination concurrent with complicated positions, turns, leaps, and balances (Hendry et al., 2020; Wilmerding & Krasnow, 2017). The numerous meticulous details could easily be lost when converting this character into one that is physically simulated.

One of the key advantages of approaches like DeepMimic is their ability to learn from reference motion data, which could be particularly valuable in the context of ballet. Over the years, the employment of motion capture in ballet training and analysis has become more established, offering an accumulation of information with regards to the flow of expert dancers (Wilmerding & Krasnow, 2017; Nachum et al., 2018; Hendry et al., 2020). If this information can be used then to steer the learning process, then physically simulated characters could be reproduced to mimic the fine detail of ballet.

Nevertheless, there are significant challenges to overcome in applying these techniques to ballet animation. One major challenge is the need for a large and diverse dataset of reference motions that cover the wide range of movements and variations present in ballet. While motion capture databases for general human movements are relatively common, datasets specifically focused on ballet are more limited. The creation of a comprehensive ballet motion dataset would require collaboration with skilled dancers and choreographers, as well as careful curation to ensure the quality and diversity of the captured movements.

While DeepMimic can mimic the features of human motions with a high degree of success, it might not contain the core artistic features of ballet. These approaches are aimed at minimizing the difference at the level of the quantity of the depicted movement's discrepancies (joint angles and velocities for example) in relation to the model movement, but they neglect stylistic acuteness. For example, in pirouettes where one has to swiftly turn on one leg, in grand jetés where legs are gracefully spread along with stretch, and in arabesques where one has to stand on one leg while the other has to be bent at the knee, needs precise and elegant movements that are difficult to recreate and capture in animation. The smooth, flowing transitions between positions in adagio (slow, lyrical) sequences and the quick, intricate footwork of petit allegro (small, fast jumps) also present significant hurdles for reinforcement learning algorithms (Yuan & Kitani, 2020; Yuan et al., 2021).

The combination of using Residual Force Control (RFC) and having the dual-policy control on top of the ballet animation to the DeepMimic framework provides more advantages. RFC expunges the dynamics disparity enabling the method to create infinitely more subtle and sophisticated dance moves like the pirouettes, the arabesques as well as the grand jetés. The two-policy control structure implies that if the task does not place direct control on the series of ballet sequences, then longer strings of sequences can be created, which in turn, may allow for the conduction of larger, more comprehensive movements to be made. These advancements make it possible to have more realistic, complex, and interactive ballet animations (Yuan et al., 2020; Peng et al. 2021).

Developing reward functions that can encourage the character to exhibit the desired qualities, such as grace (e.g., smooth, effortless transitions between movements), fluidity (e.g., continuous, uninterrupted flow of motion), and musicality (e.g., synchronization and responsiveness to musical elements like rhythm and phrasing), may benefit from leveraging the deep knowledge of ballet principles and conventions developed by dance experts and researchers. Furthermore, quantitatively assessing the aesthetic qualities of generated ballet motions remains an open challenge that would require input from expert dancers and choreographers to fully address. Exploring techniques for modeling and improving the realism and expressive fidelity of complex ballet movements with reinforcement learning is a rich avenue for further research.

One potential approach to address these challenges is to incorporate domain-specific knowledge and data accumulation into the learning process. Collaboration with ballet professionals, such as dancers, choreographers, and teachers, could provide valuable insights into the key features and criteria that define high-quality ballet movements. The gathered data and insights could be used to inform the design of reward functions, evaluation metrics, and training curricula that are tailored to the specific needs of ballet animation. Another approach is to look for the possibilities of applying the higher-level learning architectures and corresponding methods that attempt to mimic the compositional nature of ballet movements. Fundamentally, ballet technique entails repetitive patterns and rhythm while making the movements and poses into continuous, small units of phrases and variations. Hierarchical reinforcement learning methods, such as option learning or skill embedding (Nachum et al., 2018), could potentially enable the learning of reusable movement primitives that can be composed and adapted to create novel ballet sequences.

The possibility of applying deep reinforcement learning in ballet animation has the following advantages: From a creative point of view, it could help choreographers and artists to expand the potential of their work and the capability of robotics without limitations of a human body. Such characters can be employed for physical simulation of dance, which would enable swift designing and visualizing of any choreographic concepts. Moreover, the potential to produce natural and lifelike ballet would be useful for the development of VR and experiencing performance arts, extending the viewers' engagement possibilities (Park et al., 2019).

From a training and education point of view, physics-based ballet animation can also be of some use. For assisting students with correct technique and alignment it could supplement traditional classroom learning with an added resource for learning and development. Additionally, the application of the actors may allow for developing the individualized training activities tied to the participants' learning curve and specific features. Nevertheless, one must appreciate that it has been presented that it is only at the initial stage of using deep reinforcement learning for the purpose of animating ballet moves. It is essential to focus on the further studies for discovering how the domain knowledge can be incorporated for modifying the learning algorithms and for assessing the quality and diversification of emergent movements. Moreover, in the process of acceptance and implementation of these technologies, the cultural and artistic approaches may be required for the ballet.

Consequently, the possible opportunities in the field of ballet to use the physics-based character animation with the help of deep reinforcement learning are both promising and challenging at the same time. Although approaches employed with DeepMimic achieved state-of-the-art performance, translating the obtained results into the domain of ballet will need additional work and involvement of experts in this domain. Thus, building on the classical authority and the experience of ballet together with using the potential of deep learning a new generation of tools and applications for creative practice, education, and audience involvement can be generated. Over time the field will have to grow and find ways to effectively apply the technology and creativity it deserves when presenting ballet without losing the thorough art of it.

Current Limitations and Future Work

Self-supervised learning techniques based on deep reinforcement learning have been applied to physics-based character animation with promising results in the last years, however, there are limitations that need to be faced to fulfill the full potential of the method applications. One intriguing limitation is the considerable time for computation and large data demand when training sophisticated control policies. For example, the DeepMimic framework requires millions of simulation steps, and thousands of reference motion clips to learn the naturalistic behaviors and movements (Peng et al., 2018). This can prove costly and impractical especially in programs that need near real time or even constant processing.

Another task is the application of the learned skills to the new characters, environment, or the task. Most of the existing approaches are learned from the specific character models and the motion capture data, which is tied to the character models, which may hinder their flexibility in terms of morphology, style, or context. These problems require better and more efficient approaches to transfer learning and new aspects of domain adaptation, such as creating learned motions that are easily adaptable and customizable to different characters (Jiang et al., 2019).

Current physics-based character animation techniques are not able to capture the entire expressive and stylistic nature of complex movements especially in domains like ballet. While the focus on physical realism and dynamic consistency is important, it is also necessary to consider the artistic and aesthetic dimensions of movement. This applies to the integration of the mentioned aspects into the reward functions and metrics used in Reinforcement Learning. A potential area of opportunity in this regard is the application of adversarial learning methods that fall under the AMP framework (Peng et al. 2021; Escontrela et al., 2022), as such approaches might help to foster the creation of a greater variety of motions while ensuring that they belong to the same stylistic field.

Further limitation in deep reinforcement learning methods comes from the reliance on predefined motion capture data as reference material for learning. While this can be effective for reproducing specific movements or styles, it may not fully capture the creative potential of physics-based animation. Developing methods that can generate novel and expressive motions without explicit reference data is a key area for future research. This could involve techniques such as unsupervised learning, generative modeling, or exploration-based approaches that encourage the discovery of new movement patterns and variations. There are also challenges related to the usability and accessibility of these techniques for non-expert users. Many current frameworks require significant technical expertise in machine learning, physics simulation, and character animation. Developing more intuitive and user-friendly tools that allow choreographers, dancers, and other domain experts to easily specify and direct the desired behaviors could facilitate the adoption and creative application of these technologies.

While the application of deep reinforcement learning to physics-based character animation has shown great promise, there are still significant limitations and challenges to be addressed. Future work should focus on developing more efficient, generalizable, and expressive methods that can capture the full range of human movement and creativity. Collaboration with domain experts and consideration of the ethical and social implications will be essential to ensure that these technologies are developed and applied in a way that respects and enhances the art form of ballet. As the field continues to evolve, there is an exciting opportunity to push the boundaries of what is possible with physics-based character animation and explore new frontiers in virtual performance and training.

Ethical Considerations

The application of deep reinforcement learning for physics-based character animation in ballet performances can cause concern for ensuring that innovation does not erode the art form's artistic principles. Technology uses also have implications on human labor with new techniques in animating choreography could lead to the redundancy of the professional dancers and animators. Critical considerations should be made to explore how these technologies can work best hand in hand with existing human talent. Data privacies and consent are also crucial, as there is a need for proper rights and control of simulation data, specifically for the dancers that are captured on camera. With the ethical and social concerns present, an active discourse between the ballet community and the deep learning research community should be maintained to ensure that technology is used in a way that is responsive to the field's ethics and traditions. Desired outcomes from such discourse will be a set of principles and guidelines for the use of deep reinforcement learning technologies in real world applications for ballet, applied to researchers, dancers, choreographers, and educators alike.

Conclusion

In physics-based character animation, the employment of Deep Reinforcement Learning has proved to be highly effective in designing realistic and expressive human motions. Frameworks such as DeepMimic, Residual Force Control (RFC), Adversarial Motion Priors (AMP), and Simulated Character Control for 3D Human Pose Estimation (SimPoE) have opened new frontiers in mimicked motion's feasibility; they achieve high fidelities in impersonating intricate and dynamical motion patterns. However, few issues are still open for improvement: the computational cost is still high, generalization of the method is a problem, and the tools should be more expressive and easier to use.

Another avenue of study is to aim at improving the examination and analysis techniques for practical, nonrestrictive, and easier to apply approaches to identify all aspects of human motion and imagination. One will have to involve real life domain experts as well as consider the ethical and social issues that apply to make sure that these technologies will be used and created in such a way that benefits ballet as a human art form.

The improvements in deep reinforcement learning offer much promise for not only altering the course of how we produce and experience character movements in different forms of media and entertainment, such as film, performances, and video games, but also in changing the nature of people's interactions and relationships with computers and technology like virtual reality and robotics. Thus, we have the tools and potentialities for the development of new methods and media which allow for a better understanding and appreciation, as well as creation of movement and its various manifestations.

References

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). OpenAI gym. arXiv Preprint, arXiv:1606.01540.

Carnegie Mellon University (CMU). (2021). CMU Graphics Lab Motion Capture Database. Retrieved from http://mocap.cs.cmu.edu/

Cheng, X., Ji, Y., Chen, J., Yang, R., Yang, G., & Wang, X. (2024). Expressive whole-body control for humanoid robots.

Escontrela, A., Peng, X. B., Yu, W., Zhang, T., Iscen, A., Goldberg, K., & Abbeel, P. (2022). Adversarial motion priors make good substitutes for complex reward functions. In 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (pp. 2345-2352). IEEE.

Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In Proceedings of the International Conference on Machine Learning (pp. 1861-1870). PMLR.

Heess, N., Sriram, S., Lemmon, J., Merel, J., Wayne, G., Tassa, Y., ... & Silver, D. (2017). Emergence of locomotion behaviours in rich environments. arXiv Preprint, arXiv:1707.02286.

Hendry, D., Chai, K., Campbell, A., Hopper, L., O'Sullivan, P., & Straker, L. (2020). Development of a human activity recognition system for ballet tasks. Sports Medicine - Open, 6(1), 1-10.

Huo, J. (2018). Dynamic time warping and FFT: A data preprocessing method for electrical load forecasting. International Journal of Advanced Computer Science and Applications, 9(2), 10-17.

Jiang, Y., Van Wouwe, T., De Groote, F., & Liu, C. K. (2019). Synthesis of biologically realistic human motion using joint torque actuation. ACM Transactions on Graphics (TOG), 38(4), 1-12.

Li, C., Fussell, L., & Komura, T. (2021). Multi-agent reinforcement learning for character control. The Visual Computer, 37(12), 3115-3123.

Liang, J., Xu, H., Hou, Y., Wang, Z., Luo, P., Yuan, Y., & Zhang, L. (2018). GPU-accelerated real-time simulation and rendering of large-scale crowd and traffic. In SIGGRAPH Asia 2018 Technical Briefs (pp. 1-4). ACM.

Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., ... & Wierstra, D. (2015). Continuous control with deep reinforcement learning. arXiv Preprint, arXiv:1509.02971.

Nachum, O., Gu, S. S., Lee, H., & Levine, S. (2018). Data-efficient hierarchical reinforcement learning. In Advances in Neural Information Processing Systems (pp. 3303-3313).

Park, S. I., Ponce, S. P., Huang, J., Cao, Y., & Quek, F. (2019). Low-cost VR interactive tool for participatory urban planning. In Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems (pp. 1-6). ACM.

Peng, X. B., van de Panne, M., & Abbeel, P. (2017). DeepLoco: Dynamic locomotion skills using hierarchical deep reinforcement learning. ACM Transactions on Graphics (TOG), 36(4), 1-13.

Peng, X. B., Abbeel, P., Levine, S., & van de Panne, M. (2018). DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. ACM Transactions on Graphics (TOG), 37(4), 1-14.

Peng, X. B., Coumans, E., Zhang, T., Lee, T., Tan, J., & Levine, S. (2020). Learning agile robotic locomotion skills by imitating animals. In Robotics: Science and Systems. arXiv Preprint, arXiv:2004.00784.

Peng, X. B., Ma, Z., Abbeel, P., Levine, S., & Kanazawa, A. (2021). AMP: Adversarial motion priors for stylized physics-based character control. ACM Transactions on Graphics, 40(4), 1-20.

Plappert, M., Houthooft, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., ... & Andrychowicz, M. (2017). Parameter space noise for exploration. arXiv Preprint, arXiv:1706.01905.

Wilmerding, V., & Krasnow, D. (2017). Motor learning and teaching dance. In V. Karkou, S. Oliver, & S. Lycouris (Eds.), The Oxford Handbook of Dance and Wellbeing (pp. 421-436). Oxford University Press.

Won, J., Gopinath, D., & Hodgins, J. (2020). A scalable approach to control diverse behaviors for physically simulated characters. ACM Transactions on Graphics (TOG), 39(4), 1-14.

Won, J., Gopinath, D., & Hodgins, J. (2021). Control strategies for physically simulated characters performing two-player competitive sports. ACM Transactions on Graphics, 40(4), 1-20.

Yang, Z., & Yin, Z. (2021). Efficient hyperparameter optimization for physics-based character animation. Proceedings of the ACM on Computer Graphics and Interactive Techniques, 4(1), 1-19.

Yuan, Y., & Kitani, K. (2020). Residual force control for agile human behavior imitation and extended motion synthesis. In Advances in Neural Information Processing Systems (pp. 21763-21774).

Yuan, Y., Wei, S. E., Simon, T., Kitani, K., & Saragih, J. (2021). SimPoE: Simulated character control for 3D human pose estimation. arXiv Preprint, arXiv:2104.00683.