Check for updates

Machine Learning Approaches for Motor Learning: A Short Review

Baptiste Caramiaux[1]*, Jules Françoise[2], Wanyu Liu[1,3], Téo Sanchez[1] and Frédéric Bevilacqua[3]

[1] Université Paris-Saclay, CNRS, Inria, LRI, Gif-sur-Yvette, France, [2] Université Paris-Saclay, CNRS, LIMSI, Orsay, France, [3] STMS IRCAM-CNRS-Sorbonne Université, Paris, France

Machine learning approaches have seen a considerable number of applications in human movement *modeling* but remain limited for motor *learning*. Motor learning requires that motor variability be taken into account and poses new challenges because the algorithms need to be able to differentiate between new movements and variation in known ones. In this short review, we outline existing machine learning models for motor learning and their adaptation capabilities. We identify and describe three types of adaptation: *Parameter adaptation* in probabilistic models, *Transfer and meta-learning* in deep neural networks, and *Planning adaptation* by reinforcement learning. To conclude, we discuss challenges for applying these models in the domain of motor learning support systems.

Keywords: movement, computational modeling, machine learning, motor control, motor learning

1. INTRODUCTION

The use of augmented feedback on movements enables the development of interactive systems designed to facilitate motor learning. Such systems, which we refer to as *motor learning support systems*, require the movement data be captured and processes and that augmented feedback be returning to the users. These systems have primarily been investigated in rehabilitation [e.g., motor recovery after injury (Kitago and Krakauer, 2013)] or in other forms of motor learning-inducing contexts, such as dance pedagogy (Rivière et al., 2019) or entertainment (Anderson et al., 2013).

Motor learning support systems model human movements, taking into account the underlying learning mechanisms. While computational models have been proposed for simple forms of motor learning (Emken et al., 2007), modeling the processes at play in more complex skill learning remains challenging. Motor learning usually refers to two types of mechanisms: motor adaptation and motor skill acquisition. The former, motor adaptation, is the process by which the motor system adapts to perturbations in the environment (Wolpert et al., 2011). Adaptation tasks take place over a rather short time span (hours or days) and do not involve learning a new motor policy. The latter, motor skill acquisition, involves learning a new control policy, including novel movement patterns and shifts in speed-accuracy trade-offs (Shmuelof et al., 2012; Diedrichsen and Kornysheva, 2015). Complex skills are typically learned over months or years (Anders Ericsson, 2008; Yarrow et al., 2009).

The need for computational advances in motor learning research has recently been pointed out in the field of neurorehabilitation (Reinkensmeyer et al., 2016; Santos, 2019). We believe that data-driven strategies using machine learning represent a complementary approach to analytical models of movement learning. Recent results in machine learning have shown impressive advances in movement modeling, such as action recognition or movement prediction (Rudenko et al., 2019). However, it is still difficult to apply such approaches to motor learning support systems. In particular, computational models must meet specific adaptation requirements in order to address

the different variability mechanisms induced by motor adaptation and motor skill learning. These models have to account for both fine-grained changes in movement execution arising from motor adaptation mechanisms and more radical changes in movement execution due to skill acquisition mechanisms.

We propose a short review of the adaptation capabilities of machine learning applied to movement modeling. The objective of this review is not to be exhaustive but rather to provide an overview of recent publications on machine learning that we found significant for motor learning research. We believe that such an overview is currently missing and can offer novel research perspectives, targeting primarily researchers in the field of motor learning and behavioral sciences. In order to build the review presented in this paper, we focused on recent articles (typically <10 years old). At the time of writing (the end of 2019), we queried four online databases (Google Scholar, PubMed, Arxiv, ACM Digital Library) by combining the following keywords: "Human Movement," "Motor Model," "Modeling/Modelling," "Tracking," "Control," "Synthesis," "Movement Generation," "Movement Prediction," "On-line Learning," "Adaptation," "Gesture Recognition," "Deep Learning," and "Imitation Learning." We then compiled the papers in a spreadsheet and conducted a selection based on the type of model adaptation, the modeling technique, the field, and the input data considered. We summarize the review in **Table 1** and identify three adaptation categories in machine learning-based human modeling: (1) Parameter adaptation in probabilistic models, (2) Transfer and meta-learning in deep neural networks, and (3) Planning adaptation by reinforcement learning. We present the selected papers according to the type of adaptation and discuss their use in motor learning research.

2. PARAMETER ADAPTATION IN PROBABILISTIC MODELS

Research in movement recognition and generation has, for a long time, used parametric probabilistic approaches, such as Gaussian Mixture Models (GMM), Hidden Markov Models (HMM), or Dynamic Bayesian Networks (DBN). These models are characterized by a set of trained parameters that can be adapted during execution, either by providing new examples during the interaction or adapting the model parameters online according to the characteristics of the task.

GMMs have been used in robotics and HCI to learn movement trajectory models from a few demonstrations given by a human teacher (Calinon et al., 2007). In robotics, Calinon (2016) proposed such an approach to adapt the robot movement parameters when new target coordinates are set for the robot arm. The underlying model is a GMM trained from a few human movement demonstrations. In the context of movement-based interaction, Françoise et al. (2016) proposed a one-shot user adaptation process where the input movement associated with a sequence of sound synthesis parameters can be estimated from a single demonstration in order to retrain the underlying GMM. They showed that user-adapted feedback can support

the consistency of movement execution but that the adaptation process is efficient for limited movement variations. Sarasua et al. (2016) used GMM for soft recognition of conducting gestures that can adapt easily to user idiosyncrasies. The GMM-based mapping is learned from gesture demonstrations performed while listening to the desired musical rendering. The model is able to interpolate between demonstrations but cannot account for dramatic input variations. When tasks require that the dynamics and temporal evolution of the movement be encoded, generative sequence models, such as HMMs have been applied to gesture recognition from a few examples (Françoise and Bevilacqua, 2018) as well as movement generation (Tilmanne et al., 2012). Such adaptation techniques are often efficient when variations remain small in comparison with the overall movement dynamics.

Another approach, proposed by Caramiaux et al. (2015), consists of tracking probability distribution parameters representing input movement variations from a set of gesture templates. Tracking uses particle filtering, which updates state parameters representing movement variations (such as scale, speed, or orientation). The method can account for large, slow variations. However, the tracking method does not learn the structure of the gesture variations and forgets previously observed states.

Finally, parametric probabilistic models can be trained online to account for new movement classes. Kulić et al. (2008, 2012) proposed an HMM-based iterative training procedure for gesture recognition and generation. The method relies on unsupervised movement segmentation, from which it automatically extracts existing and new primitives (using Kullback-Leibler divergence). This strategy enables both the fine-grained adaptation of existing motor primitives and the extension of the vocabulary of motor skills. However, unsupervised segmentation remains difficult for complex gestures, and the learning remains cumulative, with an ever-growing vocabulary rather than a continuous adaptation to motor learning. Other online strategies for segmentation with adaptive behavior are described in Kulic et al. (2009).

In summary, parametric adaptation enables fine-grained adaptation to task variations and restricted input movement variations. The typical use case is learning by demonstration (in human-robot interaction), or personalization (in human-machine interaction).

3. TRANSFER AND META-LEARNING IN DEEP NEURAL NETWORKS

Transfer and meta-learning are techniques aiming to accelerate and improve the learning procedures of complex computational models, such as Deep Neural Networks (DNN). The objective is to adapt pre-trained DNN efficiently to new tasks or application domains that are unseen during training. This research is based on the literature on deep learning applied to movement modeling, which typically involves large datasets and benchmark-driven tests. The most popular approaches of this kind are Recurrent Neural Networks (RNN) (Fragkiadaki et al., 2015; Mattos et al., 2015; Alahi et al., 2016; Ghosh et al.,

TABLE 1 | Summary of the papers selected in our short survey, classified according to the type of adaptation involved in machine learning-based movement modeling.

Type of adaptation	Models	Application domains	Input data	Papers
Parameter adaptation	GMM	Gesture-based interaction	Gesture and movement data	Françoise et al., 2016
				Sarasua et al., 2016
		Human-robot interaction	Robot arm	Calinon et al., 2007
				Calinon, 2016
	One-shot HMM	Gesture-based interaction	Gesture and movement data	Françoise and Bevilacqua, 2018
	Incremental HMM	Human-robot interaction	3D motion capture	Kulić et al., 2008, 2012
	Stylistic HMM	Movement synthesis	3D motion capture	Tilmanne et al., 2012
	Particle filtering	Gesture-based interaction	Gesture and movement data	Caramiaux et al., 2015
Transfer learning	Temporal	Health, rehabilitation	Inertial sensors	Rad and Furlanello, 2016
	CNN	Interactive movement generation	Motion capture data	Holden et al., 2016
		Movement analysis	Videos and force measurements	Zecha et al., 2018
	2D CNN	Gesture-based interaction	Photo reflective data	Kikui et al., 2018
		Health & rehabilitation	EMG data	Côté-Allard et al., 2019
	RNN	Human-robot interaction		Cheng et al., 2019
		Human motion prediction	Motion capture data	Martinez et al., 2017
	Recurrent Encoder-Decoder	Human motion prediction	Motion capture data	Wang and Feng, 2019
Meta-learning	CNN-LSTM	Human-robot interaction	Robot arm trajectory	Duan et al., 2017
			Raw pixels	Finn et al., 2017b
				Yu et al., 2018
	Recurrent Encoder-Decoder	Movement generation	Motion capture data	Gui et al., 2018
Planning adaptation	IRL	Human-robot interaction	Joint dynamics	Kolter et al., 2008
				Finn et al., 2016
	GAIL	Human-robot interaction	Joint Dynamics	Ho and Ermon, 2016
				Guo et al., 2018
				Hausman et al., 2017
			Raw pixels	Zhu et al., 2018
	VAE + GAIL	Human-robot interaction	Joint dynamics	Wang Z. et al., 2017

2017; Martinez et al., 2017; Kratzer et al., 2019; Wang and Feng, 2019), and Temporal or Spatio-temporal Convolutional Neural Networks (CNN) (Gehring et al., 2017; Li et al., 2018, 2019; Zecha et al., 2018).

3.1. Transfer Learning

Transfer learning adapts a pre-trained model on a source domain to new target tasks. Several strategies exist (Scott et al., 2018). Transfer learning for movement modeling mainly relies on *embedding learning*: movement features (or embeddings) are learned from the source domain, providing well-shaped features for the target domain.

Movement embeddings are learned from large movement datasets. A first strategy involves one-dimensional convolutions over the time domain (Holden et al., 2016; Rad and Furlanello, 2016). Rad and Furlanello (2016) propose that learning be embedded using temporal convolution in order to improve diagnostic classification of autism spectrum disorder from inertial sensor data. The benefits of transfer learning are assessed on two datasets collected from the same participants three years apart. In another context, Holden et al. (2016) makes use of transfer learning to synthesize movements from high-level control parameters that are easily configurable by human-users.

Based on pre-trained movement embeddings from motion capture data, a mapping between high-level parameters and these embeddings can be efficiently learned according to the user needs.

Spatio-temporal convolutions can also be used to extract movement embeddings. Kikui et al. (2018) used this approach for inter- and intra-user adaptation of a gesture recognition system using photo reflective sensor data from a headset. They showed that transfer learning improves accuracy when the number of examples per class is low (lower than 6 ex/class). Also for classification, Côté-Allard et al. (2019) showed that embedding learning systematically improved the classification accuracy of EMG-based movement data; in particular, they found that embedding learning using CNNs on Continuous Wavelet Transform (CWT) gives the best results.

Finally, RNN can also be used to learn movement embeddings, although this is not the most common approach. In the context of human-robot interaction, Cheng et al. (2019) trained an RNN-based movement model offline and then adapted the last layer parameters through recursive least square errors. The goal is to adapt the robot control command to human behavior in real-time.

In summary, transfer learning of movement features has been proposed (1) to enable interactive movement generation or (2) to improve classification performance. Several problems remain to be addressed, especially in the context of motor learning. First, it is unclear how the model architecture and the size of the training set of the transfer task affect the approach. Second, the extent to which successive transfers would provoke dramatic forgetting of previously transferred tasks remains unexplored.

3.2. Meta-learning

Meta-learning designates the ability of a system to *learn how to learn* by being trained on a set of tasks (rather than a single task), such as learning faster (with fewer examples) on unseen tasks. Meta-learning is close to transfer learning, but, while transfer learning aims to use knowledge from a source application domain in order to improve or accelerate learning in a target application domain, meta-learning improves the learning procedure itself in order to handle various application domains.

Meta-learning of movement skills has been proposed in robotics and human-robot interaction to efficiently train robot actions from one or a few demonstrations. Duan et al. (2017) proposed a one-shot imitation learning algorithm where a regressor is trained against the output actions to perform the task, conditioned by a single demonstration sampled from a given task. This approach is close to the regression-based technique previously presented in section 2 but is formalized on a set of tasks. For example, the task could be to train the robot-arm to stack a variable number of physical blocks among a variable number of piles. The evaluation methods rely on tests on seen and unseen demonstrations during training. Their results showed that the robot performed equally well with seen and unseen demonstrations.

Adaptation through meta-learning in motor learning has also been investigated with the model-agnostic meta-learning (MAML) method (Finn et al., 2017a), which allows faster weight adaptation to new examples representing a task. Finn et al. (2017b) and Yu et al. (2018) extended the MAML approach for one-shot imitation learning by a robotic arm. Finn et al. (2017b) first demonstrated that vision-based policies can be fine-tuned from one demonstration. They conducted experiments using two types of tasks (pushing an object and placing an object on a target) on both a simulated and a real robot using video-based input data. Their results outperformed previous results (see for instance Duan et al., 2017) in terms of the number of demonstrations needed for adaptation. Yu et al. (2018) then addressed the problem of one-shot learning of motor control policies with domain shift. Their experiments on simulated and real robot actions showed good results on tasks, such as pushing, placing, and picking-and-placing objects.

The MAML method has also found applications in human motion forecasting (Gui et al., 2018), for which large amounts of annotated motion capture data are typically needed. They propose an approach based on combining MAML and model regression networks (Wang and Hebert, 2016; Wang Y.-X. et al., 2017), allowing a good generic initial model to be learned and enabling efficient adaptation to unseen tasks. They showed that the model outperforms baselines with five examples of motion capture data of walking.

4. ADAPTATION THROUGH REINFORCEMENT LEARNING

Reinforcement Learning (RL) enables robotic agents to acquire new motor skills from experience, using trial-and-error interactions with their environment (Kober et al., 2013). Contrary to the imitation learning approaches discussed in section 2, where expert demonstrations are used to train a model encoding a given behavior, RL relies on objective functions that provide feedback on the robot's performance.

Most approaches to imitation learning rely on a supervised paradigm where the model is fully specified from demonstrations without subsequent self-improvement (Billard et al., 2016). To ensure good task generalization, imitation learning requires a significant number of high-quality demonstrations that provide variability while ensuring high performance. While RL can raise impressive performance, the learning process is often very slow and can lead to unnatural behavior. A growing body of research investigates the combination of these two paradigms to improve the models' adaptation to new tasks, making the learning process more efficient and improving the generalization of the tasks from a few examples.

Demonstrations can be integrated in the RL process in various ways. One approach consists of initializing RL training with a model learned by imitation (Kober et al., 2013), typically by a human teacher. Demonstrations of such tasks are used to generate initial policies for the RL process, enabling robots to rapidly learn to perform tasks, such as reaching, ball-in-a-cup, playing pool, manipulating a box, etc. A second strategy consists of deriving cost functions from demonstrations, for instance using inverse reinforcement learning (Kolter et al., 2008; Finn et al., 2016). Finn et al. (2016) showed that using 25–30 human demonstrations (by direct manipulation) to learn the cost function was sufficient for the robot to learn how to perform dish placement and pouring tasks.

Building upon the success of Generative Adversarial Networks in other fields of machine learning, Generative Adversarial Imitation Learning (GAIL) has been proposed as an efficient method for learning movement from expert demonstrations (Ho and Ermon, 2016). In GAIL, a discriminator is trained to discriminate between expert trajectories (considered optimal) and policy trajectories generated by a generator that is trained to fool the discriminator. This approach was then extended to reinforcement learning through self-imitation, where optimal trajectories are defined by previous successful attempts (Guo et al., 2018). Several extensions of the adversarial learning framework were proposed to improve its stability or to handle unstructured demonstrations (Hausman et al., 2017; Wang Z. et al., 2017). These recent approaches have been evaluated on a standard set of tasks using simulated environments, in particular OpenAI Gym MuJoCo (Todorov et al., 2012), including continuous control tasks, such as inverted pendulums, 4-legged walk, and humanoid walk.

Recently, Zhu et al. (2018) proposed simultaneous imitation and reinforcement learning through a reward function that combines GAIL and RL. Zhu et al. (2018) evaluated their approach on several manipulation tasks (such as lifting and stacking blocks, clearing a table, and pouring) with a robot arm. Demonstrations were performed using a 3D controller, the training was done in a simulated environment, and the tasks were performed by a real robot arm. In comparison with GAIL or RL alone, the evaluation shows that the combination learns faster and achieves better performance.

5. DISCUSSION

This paper reviews three types of adaptation in machine learning applied to movement modeling. In this section, we discuss how adaptive movement models can be used to support motor learning, including both motor adaptation and motor skill acquisition.

First, motor adaptation mechanisms involve variations of an already-trained skill. Computationally, motor adaptation can be seen as an optimization process that learns and cancels external effects in order to return to baseline (Shadmehr and Mussa-Ivaldi, 1994). Accounting for these underlying variations requires rapid mechanisms and robust statistical modeling. Probabilistic model parameter adaptation (section 2) appears to be a good candidate for understanding the movement variability induced by motor adaptation processes. However, while motor adaptation has been widely studied, very little is known on the statistical structure of motor adaptation, particularly trial-to-trial motor variability (Stergiou and Decker, 2011). Transfer learning could also be used: pre-trained models (RNNs or CNNs) that capture some structure of movement parameters (i.e., low-dimensional subspaces of the parameter space), can be adapted online for fine-grained variations. Here, open questions concern how such variations can account for structural learning in motor control (Braun et al., 2009).

Second, more dramatic changes in movement patterns, as induced by learning new motor skills, might require computational adaptation that involves re-training procedures. Transfer and meta-learning (section 3) involve the adaptation of high-capacity movement models to new tasks and could be used in this context. One difficulty is to assess to what extent transferring a given model to new motor control policies would

induce the model to forget past skills. For instance, it was found that movement models relying on deep neural networks might lead to catastrophic forgetting (Kirkpatrick et al., 2017). Also, meta-learning algorithms, such as MAML (Finn et al., 2017a) are currently not suitable for adaptation to several new motor tasks. Self-imitation and reinforcement mechanisms (section 4) could help to generalize to a wider set of tasks. A current challenge is to learn suitable action selection policies. Although exploration-exploitation is known to be central in motor learning (Herzfeld and Shadmehr, 2014), it is yet unclear what process drives action selection in the brain (Carland et al., 2019; Sugiyama et al., 2020). These approaches still need to be experimentally assessed in a motor learning context.

Finally, the last challenge that we want to raise in this paper regards the continuous evolution of motor variation patterns. Motor execution may vary continuously over time due to skill acquisition and morphological changes. Accounting for such open-ended tasks may require new forms of adaptation, such as continuous online learning, as proposed by Nagabandi et al. (2018). We think that this is a promising research direction, raising the central question of computation and memory in motor learning (Herzfeld et al., 2014).

In closing, to be integrated into motor learning support systems, the aforementioned machine learning approaches should be combined with adaptation mechanisms that aim to generalize models to new movements and new tasks efficiently. We do not advocate solely for adaptive machine learning explaining motor learning processes. We propose adaptation procedures that can account for variation patterns observed in behavioral data, leading to performance improvements in motor learning support systems.

AUTHOR CONTRIBUTIONS

BC made the first draft. All authors contributed to the manuscript, adding content, and revising it critically for important intellectual content.

FUNDING

REFERENCES

Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., and Savarese, S. (2016). "Social LSTM: human trajectory prediction in crowded spaces," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (Las Vegas, NV), 961–971. doi: 10.1109/CVPR.2016.110

Anders Ericsson, K. (2008). Deliberate practice and acquisition of expert performance: a general overview. Acad. Emerg. Med. 15, 988–994. doi: 10.1111/j.1553-2712.2008.00227.x

Anderson, F., Grossman, T., Matejka, J., and Fitzmaurice, G. (2013). "Youmove: enhancing movement training with an augmented reality mirror," in Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (St. Andrews: ACM), 311–320. doi: 10.1145/2501988.2502045

Billard, A. G., Calinon, S., and Dillmann, R. (2016). "Learning from humans," in Springer Handbook of Robotics (Springer), 1995–2014. doi: 10.1007/978-3-319-32552-1_74

Braun, D. A., Aertsen, A., Wolpert, D. M., and Mehring, C. (2009). Motor task variation induces structural learning. Curr. Biol. 19, 352–357. doi: 10.1016/j.cub.2009.01.036

Côté-Allard, U., Fall, C. L., Drouin, A., Campeau-Lecours, A., Gosselin, C., Glette, K., et al. (2019). Deep learning for electromyographic hand gesture signal classification using transfer learning. IEEE Trans. Neural Syst. Rehabil. Eng. 27, 760–771. doi: 10.1109/TNSRE.2019.2896269

Calinon, S. (2016). A tutorial on task-parameterized movement learning and retrieval. Intell. Serv. Robot. 9, 1–29. doi: 10.1007/s11370-015-0187-9

Calinon, S., Guenter, F., and Billard, A. (2007). On learning, representing, and generalizing a task in a humanoid robot. *IEEE Trans. Syst. Man Cybern. B Cybern.* 37, 286–298. doi: 10.1109/TSMCB.2006.886952

Caramiaux, B., Montecchio, N., Tanaka, A., and Bevilacqua, F. (2015). Adaptive gesture recognition with variation estimation for interactive systems. *ACM Trans. Interact. Intell. Syst.* 4:18. doi: 10.1145/2643204

Carland, M. A., Thura, D., and Cisek, P. (2019). The urge to decide and act: implications for brain function and dysfunction. *Neuroscientist* 25, 491–511. doi: 10.1177/1073858419841553

Cheng, Y., Zhao, W., Liu, C., and Tomizuka, M. (2019). "Human motion prediction using semi-adaptable neural networks," in *2019 American Control Conference (ACC)* (Philadelphia, PA: IEEE), 4884–4890. doi: 10.23919/ACC.2019.8814980

Diedrichsen, J., and Kornysheva, K. (2015). Motor skill learning between selection and execution. *Trends Cogn. Sci.* 19, 227–233. doi: 10.1016/j.tics.2015.02.003

Duan, Y., Andrychowicz, M., Stadie, B., Ho, O. J., Schneider, J., Sutskever, I., et al. (2017). "One-shot imitation learning," in *Advances in Neural Information Processing Systems* (Los Angeles, CA), 1087–1098.

Emken, J. L., Benitez, R., Sideris, A., Bobrow, J. E., and Reinkensmeyer, D. J. (2007). Motor adaptation as a greedy optimization of error and effort. *J. Neurophysiol.* 97, 3997–4006. doi: 10.1152/jn.01095.2006

Finn, C., Abbeel, P., and Levine, S. (2017a). "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70 (Sydney, NSW: JMLR.org), 1126–1135.

Finn, C., Levine, S., and Abbeel, P. (2016). "Guided cost learning: deep inverse optimal control via policy optimization," in *International Conference on Machine Learning* (New York, NY), 49–58.

Finn, C., Yu, T., Zhang, T., Abbeel, P., and Levine, S. (2017b). One-shot visual imitation learning via meta-learning. *arXiv* 1709.04905.

Fragkiadaki, K., Levine, S., Felsen, P., and Malik, J. (2015). "Recurrent network models for human dynamics," in *Proceedings of the IEEE International Conference on Computer Vision* (Santiago), 4346–4354. doi: 10.1109/ICCV.2015.494

Françoise, J., and Bevilacqua, F. (2018). Motion-sound mapping through interaction: an approach to user-centered design of auditory feedback using machine learning. *ACM Trans. Interact. Intell. Syst.* 8, 1–30. doi: 10.1145/3211826

Françoise, J., Chapuis, O., Hanneton, S., and Bevilacqua, F. (2016). "Soundguides: adapting continuous auditory feedback to users," in *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems* (San Josï£¡, CA: ACM), 2829–2836.

Gehring, J., Auli, M., Grangier, D., Yarats, D., and Dauphin, Y. N. (2017). "Convolutional sequence to sequence learning," in *Proceedings of the 34th International Conference on Machine Learning*, Vol. 70 (Sydney, NSW: JMLR.org), 1243–1252.

Ghosh, P., Song, J., Aksan, E., and Hilliges, O. (2017). "Learning human motion models for long-term predictions," in *2017 International Conference on 3D Vision (3DV)* (Qingdao: IEEE), 458–466. doi: 10.1109/3DV.2017.00059

Gui, L.-Y., Wang, Y.-X., Ramanan, D., and Moura, J. M. (2018). "Few-shot human motion prediction via meta-learning," in *Proceedings of the European Conference on Computer Vision (ECCV)* (Munich), 432–450. doi: 10.1007/978-3-030-01237-3_27

Guo, Y., Oh, J., Singh, S., and Lee, H. (2018). Generative adversarial self-imitation learning. *arXiv* 1812.00950.

Hausman, K., Chebotar, Y., Schaal, S., Sukhatme, G., and Lim, J. J. (2017). "Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets," in *Advances in Neural Information Processing Systems* (Los Angeles, CA), 1235–1245.

Herzfeld, D. J., and Shadmehr, R. (2014). Motor variability is not noise, but grist for the learning mill. *Nat. Neurosci.* 17, 149–150. doi: 10.1038/nn.3633

Herzfeld, D. J., Vaswani, P. A., Marko, M. K., and Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science* 345, 1349–1353. doi: 10.1126/science.1253138

Ho, J., and Ermon, S. (2016). "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems* (Barcelona), 4565–4573.

Holden, D., Saito, J., and Komura, T. (2016). A deep learning framework for character motion synthesis and editing. *ACM Trans. Graph.* 35:138. doi: 10.1145/2897824.2925975

Kikui, K., Itoh, Y., Yamada, M., Sugiura, Y., and Sugimoto, M. (2018). "Intra-/inter-user adaptation framework for wearable gesture sensing device," in *Proceedings of the 2018 ACM International Symposium on Wearable Computers, ISWC '18* (New York, NY: ACM), 21–24. doi: 10.1145/3267242.3267256

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proc. Natl. Acad. Sci. U.S.A.* 114, 3521–3526. doi: 10.1073/pnas.1611835114

Kitago, T., and Krakauer, J. W. (2013). Motor learning principles for neurorehabilitation. *Handb. Clin. Neurol.* 110, 93–103. doi: 10.1016/B978-0-444-52901-5.00008-3

Kober, J., Bagnell, J. A., and Peters, J. (2013). Reinforcement learning in robotics: a survey. *Int. J. Robot. Res.* 32, 1238–1274. doi: 10.1177/0278364913495721

Kolter, J. Z., Abbeel, P., and Ng, A. Y. (2008). "Hierarchical apprenticeship learning with application to quadruped locomotion," in *Advances in Neural Information Processing Systems* (Vancouver, BC), 769–776.

Kratzer, P., Toussaint, M., and Mainprice, J. (2019). Motion prediction with recurrent neural network dynamical models and trajectory optimization. *arXiv* 1906.12279.

Kulić, D., Ott, C., Lee, D., Ishikawa, J., and Nakamura, Y. (2012). Incremental learning of full body motion primitives and their sequencing through human motion observation. *Int. J. Robot. Res.* 31, 330–345. doi: 10.1177/0278364911426178

Kulić, D., Takano, W., and Nakamura, Y. (2008). Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains. *Int. J. Robot. Res.* 27, 761–784. doi: 10.1177/0278364908091153

Kulic, D., Takano, W., and Nakamura, Y. (2009). Online segmentation and clustering from continuous observation of whole body motions. *IEEE Trans. Robot.* 25, 1158–1166. doi: 10.1109/TRO.2009.2026508

Li, C., Zhang, Z., Sun Lee, W., and Hee Lee, G. (2018). "Convolutional sequence to sequence model for human dynamics," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Salt Lake City, UT), 5226–5234. doi: 10.1109/CVPR.2018.00548

Li, Y., Wang, Z., Yang, X., Wang, M., Poiana, S. I., Chaudhry, E., et al. (2019). Efficient convolutional hierarchical autoencoder for human motion prediction. *Vis. Comput.* 35, 1143–1156. doi: 10.1007/s00371-019-01692-9

Martinez, J., Black, M. J., and Romero, J. (2017). "On human motion prediction using recurrent neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Honolulu, HI), 2891–2900. doi: 10.1109/CVPR.2017.497

Mattos, C. L. C., Dai, Z., Damianou, A., Forth, J., Barreto, G. A., and Lawrence, N. D. (2015). Recurrent gaussian processes. *arXiv* 1511.06644.

Nagabandi, A., Finn, C., and Levine, S. (2018). Deep online learning via meta-learning: continual adaptation for model-based rl. *arXiv* 1812.07671.

Rad, N., and Furlanello, C. (2016). "Applying deep learning to stereotypical motor movement detection in autism spectrum disorders," in *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)* (Barcelona). doi: 10.1109/ICDMW.2016.0178

Reinkensmeyer, D. J., Burdet, E., Casadio, M., Krakauer, J. W., Kwakkel, G., Lang, C. E., et al. (2016). Computational neurorehabilitation: modeling plasticity and learning to predict recovery. *J. Neuroeng. Rehabil.* 13:42. doi: 10.1186/s12984-016-0148-3

Riviére, J.-P., Alaoui, S. F., Caramiaux, B., and Mackay, W. E. (2019). Capturing movement decomposition to support learning and teaching in contemporary dance. *Proc. ACM Hum. Comput. Interact.* 3:86. doi: 10.1145/3359188

Rudenko, A., Palmieri, L., Herman, M., Kitani, K. M., Gavrila, D. M., and Arras, K. O. (2019). Human motion trajectory prediction: a survey. *arXiv* 1905.06113.

Santos, O. C. (2019). Artificial intelligence in psychomotor learning: modeling human motion from inertial sensor data. *Int. J. Artif. Intell. Tools* 28:1940006. doi: 10.1142/S0218213019400062

Sarasua, A., Caramiaux, B., and Tanaka, A. (2016). "Machine learning of personal gesture variation in music conducting," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San José, CA: ACM), 3428–3432. doi: 10.1145/2858036.2858328

Scott, T. R., Ridgeway, K., and Mozer, M. C. (2018). "Adapted deep embeddings: a synthesis of methods for k-shot inductive transfer learning," in *Advances in Neural Information Processing Systems* (Montreal, QC), Vol. 2018, 76–85.

Shadmehr, R., and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14, 3208–3224. doi: 10.1523/JNEUROSCI.14-05-03208.1994

Shmuelof, L., Krakauer, J. W., and Mazzoni, P. (2012). How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J. Neurophysiol.* 108, 578–594. doi: 10.1152/jn.00856.2011

Stergiou, N., and Decker, L. M. (2011). Human movement variability, nonlinear dynamics, and pathology: is there a connection? *Hum. Mov. Sci.* 30, 869–888. doi: 10.1016/j.humov.2011.06.002

Sugiyama, T., Schweighofer, N., and Izawa, J. (2020). Reinforcement meta-learning optimizes visuomotor learning. *bioRxiv*. doi: 10.1101/2020.01.19.912048

Tilmanne, J., Moinet, A., and Dutoit, T. (2012). Stylistic gait synthesis based on hidden markov models. *EURASIP J. Adv. Signal Process.* 2012:72. doi: 10.1186/1687-6180-2012-72

Todorov, E., Erez, T., and Tassa, Y. (2012). "Mujoco: a physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Vilamoura: IEEE), 5026–5033. doi: 10.1109/IROS.2012.6386109

Wang, H., and Feng, J. (2019). Vred: A position-velocity recurrent encoder-decoder for human motion prediction. *arXiv* 1906.06514.

Wang, Y.-X., and Hebert, M. (2016). "Learning to learn: model regression networks for easy small sample learning," in *European Conference on Computer Vision* (Amsterdam: Springer), 616–634. doi: 10.1007/978-3-319-46466-4_37

Wang, Y.-X., Ramanan, D., and Hebert, M. (2017). "Learning to model the tail," in *Advances in Neural Information Processing Systems* (Los Angeles, CA), 7029–7039.

Wang, Z., Merel, J. S., Reed, S. E., de Freitas, N., Wayne, G., and Heess, N. (2017). "Robust imitation of diverse behaviors," in *Advances in Neural Information Processing Systems* (Los Angeles, CA), 5320–5329.

Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nat. Rev. Neurosci.* 12:739. doi: 10.1038/nrn3112

Yarrow, K., Brown, P., and Krakauer, J. W. (2009). Inside the brain of an elite athlete: the neural processes that support high achievement in sports. *Nat. Rev. Neurosci.* 10, 585–596. doi: 10.1038/nrn2672

Yu, T., Finn, C., Xie, A., Dasari, S., Zhang, T., Abbeel, P., et al. (2018). One-shot imitation from observing humans via domain-adaptive meta-learning. *arXiv* 1802.01557. doi: 10.15607/RSS.2018.XIV.002

Zecha, D., Eggert, C., Einfalt, M., Brehm, S., and Lienhart, R. (2018). "A convolutional sequence to sequence model for multimodal dynamics prediction in ski jumps," in *Proceedings of the 1st International Workshop on Multimedia Content Analysis in Sports* (Seoul: ACM), 11–19. doi: 10.1145/3265845.3265855

Zhu, Y., Wang, Z., Merel, J., Rusu, A., Erez, T., Cabi, S., et al. (2018). Reinforcement and imitation learning for diverse visuomotor skills. *arXiv* 1802.09564. doi: 10.15607/RSS.2018.XIV.009