



HAL
open science

Machine Learning Approaches For Motor Learning: A Short Review

Baptiste Caramiaux, Jules Françoise, Wanyu Liu, Téo Sanchez, Frédéric Bevilacqua

► **To cite this version:**

Baptiste Caramiaux, Jules Françoise, Wanyu Liu, Téo Sanchez, Frédéric Bevilacqua. Machine Learning Approaches For Motor Learning: A Short Review. *Frontiers in Computer Science*, inPress, 10.3389/fcomp.2020.00016 . hal-02558779

HAL Id: hal-02558779

<https://hal.science/hal-02558779>

Submitted on 29 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MACHINE LEARNING APPROACHES FOR MOTOR LEARNING: A SHORT REVIEW

Baptiste Caramiaux¹, Jules François², Wanyu Liu³, Téo Sanchez¹, and Frédéric Bevilacqua³

¹ Université Paris-Saclay, CNRS, Inria, LRI, France

² Université Paris-Saclay, CNRS, LIMSI, 91400, Orsay, France

³ STMS IRCAM - CNRS - Sorbonne Université, France

ABSTRACT

Machine learning approaches have seen considerable applications in human movement *modeling*, but remain limited for motor *learning*. Motor learning requires accounting for motor variability, and poses new challenges as the algorithms need to be able to differentiate between new movements and variation of 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 Modelling, 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, that we refer as *motor learning support systems*, require capturing and processing movement data and returning augmented feedback to the users. These systems have primarily been investigated in rehabilitation (e.g. motor recovery after injury [1]), or in other forms of motor learning induced contexts, such as dance pedagogy [2] or entertainment [3].

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 [4], 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 [5]. Adaptation tasks take place over a rather short time span (hours or days) and does 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 [6, 7]. Complex skills are typically learned over months or years [8, 9].

The need for computational advances in motor learning research has recently been pointed out in the field of neurorehabilitation [10, 11]. 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 [12]. 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 on 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 behavioural

Type of adaptation	Models	Application domains	Input data	Papers
Parameter adaptation	GMM	Gesture-based Interaction	Gesture and movement data	[13]
		Human-robot interaction	Robot arm	[14] [15] [16]
	One-shot HMM	Gesture-based Interaction	Gesture and movement data	[17]
	Incremental HMM	Human-robot interaction	3D Motion Capture	[18, 19]
	Stylistic HMM	Movement synthesis	3D Motion Capture	[20]
Particle filtering	Gesture-based Interaction	Gesture and movement data	[21]	
Transfer learning	Temporal CNN	Health, rehabilitation	Inertial sensors	[22]
		Interactive movement generation	Motion capture data	[23]
	Movement analysis	Videos and force measurements	[24]	
	2D CNN	Gesture-based interaction	Photo reflective data	[25]
RNN	Health & rehabilitation	EMG data	[26]	
	Human-robot interaction	Motion capture data	[27]	
Recurrent Encoder-Decoder	Human motion prediction	Motion capture data	[28]	
	Human motion prediction	Motion capture data	[29]	
Meta-learning	CNN-LSTM	Human-robot interaction	Robot arm trajectory	[30]
			Raw pixels	[31] [32]
	Recurrent Encoder-Decoder	Movement generation	Motion capture data	[33]
Planning adaptation	IRL	Human-robot interaction	Joint dynamics	[34]
				[35]
	GAIL	Human-robot interaction	Joint Dynamics	[36]
				[37] [38] [39]
VAE + GAIL	Human-robot interaction	Joint Dynamics	[40]	

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

sciences. In order to build the review presented in this paper, we focused on recent articles (typically less than 10 years). At the time of writing (end of 2019), we queried four online databases (Google Scholar, PubMed, Arxiv, ACM Digital Library) 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”, “Imitation Learning”. We then compiled the papers in a spreadsheet and conducted a selection based on the type of model adaptations, 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 along 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 few demonstrations given by a human teacher [15]. In robotics, [16] 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 few human movement demonstrations. In the context of movement-based interaction, [13] 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. [14] 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 to encode the dynamics and temporal evolution of the movement, generative sequence models such as HMMs have been applied to gesture recognition from few examples [17] as well as movement generation [20]. Such adaptation techniques are often efficient when variations remain small in comparison with the overall movement dynamics.

Another approach, proposed by [21] 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. [18, 19] proposed a 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 behaviour are described in [41].

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 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, unseen during training. This research is based on the literature in 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) [42, 43, 44, 45, 28, 46, 29], and Temporal or Spatio-temporal Convolutional Neural Networks (CNN) [47, 48, 49, 24].

3.1 Transfer learning

Transfer learning adapts a pre-trained model on a source domain to new target tasks. Several strategies exist [50]. 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 [22, 23]. [22] propose embedding learning using temporal convolution in order to improve diagnosis classification of autism spectrum disorder from inertial sensor data. The benefit of transfer learning is assessed on two datasets collected from the same participants, three years apart. In another context, [23] makes use of transfer learning to synthesize movements from high-level control parameters 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. [25] uses 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, [26] 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, [27] proposed to train offline a RNN-based movement model and adapts the last layer parameters through recursive least square errors. The goal is to adapt the robot control command to human behaviour 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 affects the approach. Second, it remains unexplored the extent to which successive transfers would provoke dramatic forgetting of previously transferred tasks.

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 was proposed in robotics and human-robot interaction, to efficiently train robot actions from one or few demonstrations. [30] 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 previous regression-based technique presented in Section 2, but formalised on a set of tasks. For example, tasks can be training the robot-arm of stacking a variable number of physical blocks among a variable number of piles. The evaluation methods rely on tests on seen and unseen demonstration during training. Their results showed that the robot performed equally well with seen and unseen demonstrations.

Adaptation process through meta-learning in motor learning has also been investigated with the model-agnostic meta-learning (MAML) method [51], allowing faster weight adaptation to new examples representing a task. [31, 32] extended the MAML approach for one-shot imitation learning by a robotic arm. [31] first demonstrate that vision-based policies can be fine-tuned from one demonstration. They conducted experiments using two types of tasks (pushing object and placing object in a recipient) on both a simulated and a real robot using video-based input data. Their results outperformed previous results (see for instance [30]) in terms of the number of demonstrations needed for adaptation. Then, [32] 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 push, place, pick-and-place objects.

The MAML method has also found applications in human motion forecasting [33], for which large annotated motion capture are typically needed. They propose an approach based combining MAML and model regression networks [52, 53], allowing for learning a good generic initial model and for adaptation efficiently to unseen tasks. They showed that the model outperforms baselines with 5 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 its environment [54]. 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 [55]. To ensure a 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 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 [54], 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 tasks, 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 [34, 35]. [35] 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 [36]. 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 [37]. Several extensions of the adversarial learning framework were proposed to improve its stability or to handle unstructured demonstrations [38, 40]. These recent approaches have been evaluated on a standard set of tasks using simulated environments, in particular OpenAI Gym MuJoCo [56], including continuous control tasks such as inverted pendulums, 4-legged walk, humanoid walk.

Recently, [39] proposed simultaneous imitation and reinforcement learning through a reward function that combines GAIL and RL. [39] evaluated their approach on several manipulation tasks (such as block lifting and stacking, clearing a table, pouring content) 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 reaches 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 [57]. Accounting for these underlying variations require rapid mechanisms and robust statistical modeling. Probabilistic model parameter adaptation (Section 2) appears to be a good candidate to understand movement variability induced by motor adaptation processes. However, while motor adaptation has been widely studied, very few is known on the statistical structure of motor adaptation, in particular trial-to-trial motor variability [58]. 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 parameters space), can be adapted online for fine-grained variations. Here, open questions concern how such variations can account for structural learning in motor control [59].

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) describe 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 [60]. Also, meta-learning algorithms such as MAML [51] are currently not suitable to adapt 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 [61], it is yet unclear what process drives action selection in the brain [62, 63]. These approaches still need to be experimentally assessed in a motor learning context.

Finally, a last challenge that we want to raise in this paper regards the continuous evolution of motor variation patterns. Motor execution may continuously vary over time, due to skill acquisition and morphological changes. Accounting for such open-ended task may require new form of adaptation such as continuous online learning, as proposed by [64]. We think that this is a promising research direction, raising the central question of computation and memory in motor learning [65].

In closing, to be integrated in 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 behavioural data, leading to performance improvements in motor learning support systems.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

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

Funding

This research was supported by the ELEMENT project (ANR-18-CE33-0002) and the ARCOL project (ANR-19-CE33-0001) from the French National Research Agency.

References

- [1] TOMOKO Kitago and John W Krakauer. Motor learning principles for neurorehabilitation. In *Handbook of clinical neurology*, volume 110, pages 93–103. Elsevier, 2013.

- [2] Jean-Philippe Rivière, Sarah Fdili Alaoui, Baptiste Caramiaux, and Wendy E Mackay. Capturing movement decomposition to support learning and teaching in contemporary dance. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):86, 2019.
- [3] Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. Youmove: enhancing movement training with an augmented reality mirror. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pages 311–320. ACM, 2013.
- [4] Jeremy L Emken, Raul Benitez, Athanasios Sideris, James E Bobrow, and David J Reinkensmeyer. Motor adaptation as a greedy optimization of error and effort. *Journal of neurophysiology*, 97(6):3997–4006, 2007.
- [5] Daniel M Wolpert, Jörn Diedrichsen, and J Randall Flanagan. Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12):739, 2011.
- [6] Jörn Diedrichsen and Katja Kornysheva. Motor skill learning between selection and execution. *Trends in cognitive sciences*, 19(4):227–233, 2015.
- [7] Lior Shmuelof, John W Krakauer, and Pietro Mazzoni. How is a motor skill learned? change and invariance at the levels of task success and trajectory control. *Journal of neurophysiology*, 108(2):578–594, 2012.
- [8] Kielan Yarrow, Peter Brown, and John W. Krakauer. Inside the brain of an elite athlete: the neural processes that support high achievement in sports. *Nature Reviews Neuroscience*, 10:585–596, 2009.
- [9] K Anders Ericsson. Deliberate practice and acquisition of expert performance: a general overview. *Academic emergency medicine*, 15(11):988–994, 2008.
- [10] David J Reinkensmeyer, Etienne Burdet, Maura Casadio, John W Krakauer, Gert Kwakkel, Catherine E Lang, Stephan P Swinnen, Nick S Ward, and Nicolas Schweighofer. Computational neurorehabilitation: modeling plasticity and learning to predict recovery. *Journal of neuroengineering and rehabilitation*, 13(1):42, 2016.
- [11] Olga C Santos. Artificial intelligence in psychomotor learning: Modeling human motion from inertial sensor data. *International Journal on Artificial Intelligence Tools*, 28(04):1940006, 2019.
- [12] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrilă, and Kai O Arras. Human motion trajectory prediction: A survey. *arXiv preprint arXiv:1905.06113*, 2019.
- [13] Jules Françoise, Olivier Chapuis, Sylvain Hanneton, and Frédéric Bevilacqua. Soundguides: adapting continuous auditory feedback to users. In *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pages 2829–2836. ACM, 2016.
- [14] Alvaro Sarasua, Baptiste Caramiaux, and Ataru Tanaka. Machine learning of personal gesture variation in music conducting. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 3428–3432. ACM, 2016.
- [15] S. Calinon, F. Guenter, and A. Billard. On learning, representing, and generalizing a task in a humanoid robot. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(2):286–298, April 2007.
- [16] Sylvain Calinon. A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service Robotics*, 9(1):1–29, 2016.
- [17] Jules Françoise and Frederic Bevilacqua. Motion-sound mapping through interaction: An approach to user-centered design of auditory feedback using machine learning. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 8(2):1–30, 2018.
- [18] Dana Kulić, Wataru Takano, and Yoshihiko Nakamura. Incremental learning, clustering and hierarchy formation of whole body motion patterns using adaptive hidden markov chains. *The International Journal of Robotics Research*, 27(7):761–784, 2008.
- [19] Dana Kulić, Christian Ott, Dongheui Lee, Junichi Ishikawa, and Yoshihiko Nakamura. Incremental learning of full body motion primitives and their sequencing through human motion observation. *The International Journal of Robotics Research*, 31(3):330–345, 2012.
- [20] Joëlle Tilmanne, Alexis Moinet, and Thierry Dutoit. Stylistic gait synthesis based on hidden markov models. *EURASIP Journal on advances in signal processing*, 2012(1):72, 2012.
- [21] Baptiste Caramiaux, Nicola Montecchio, Ataru Tanaka, and Frédéric Bevilacqua. Adaptive gesture recognition with variation estimation for interactive systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 4(4):18, 2015.
- [22] NM Rad and C Furlanello. Applying deep learning to stereotypical motor movement detection in autism spectrum disorders. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW)*, 2016.

- [23] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM Trans. Graph.*, 35(4), jul 2016.
- [24] Dan Zecha, Christian Eggert, Moritz Einfalt, Stephan Brehm, and Rainer Lienhart. 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*, pages 11–19. ACM, 2018.
- [25] Kosuke Kikui, Yuta Itoh, Makoto Yamada, Yuta Sugiura, and Maki Sugimoto. Intra-/inter-user adaptation framework for wearable gesture sensing device. In *Proceedings of the 2018 ACM International Symposium on Wearable Computers*, ISWC '18, pages 21–24, New York, NY, USA, 2018. ACM.
- [26] Ulysse Côté-Allard, Cheikh Latyr Fall, Alexandre Drouin, Alexandre Campeau-Lecours, Clément Gosselin, Kyrre Glette, François Laviolette, and Benoit Gosselin. Deep learning for electromyographic hand gesture signal classification using transfer learning. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 27(4):760–771, 2019.
- [27] Yujiao Cheng, Weiye Zhao, Changliu Liu, and Masayoshi Tomizuka. Human motion prediction using semi-adaptable neural networks. In *2019 American Control Conference (ACC)*, pages 4884–4890. IEEE, 2019.
- [28] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017.
- [29] Hongsong Wang and Jiashi Feng. Vred: A position-velocity recurrent encoder-decoder for human motion prediction. *arXiv preprint arXiv:1906.06514*, 2019.
- [30] Yan Duan, Marcin Andrychowicz, Bradly Stadie, OpenAI Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, and Wojciech Zaremba. One-shot imitation learning. In *Advances in neural information processing systems*, pages 1087–1098, 2017.
- [31] Chelsea Finn, Tianhe Yu, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot visual imitation learning via meta-learning. *arXiv preprint arXiv:1709.04905*, 2017.
- [32] Tianhe Yu, Chelsea Finn, Annie Xie, Sudeep Dasari, Tianhao Zhang, Pieter Abbeel, and Sergey Levine. One-shot imitation from observing humans via domain-adaptive meta-learning. *arXiv preprint arXiv:1802.01557*, 2018.
- [33] Liang-Yan Gui, Yu-Xiong Wang, Deva Ramanan, and José MF Moura. Few-shot human motion prediction via meta-learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 432–450, 2018.
- [34] J Zico Kolter, Pieter Abbeel, and Andrew Y Ng. Hierarchical apprenticeship learning with application to quadruped locomotion. In *Advances in Neural Information Processing Systems*, pages 769–776, 2008.
- [35] Chelsea Finn, Sergey Levine, and Pieter Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *International conference on machine learning*, pages 49–58, 2016.
- [36] Jonathan Ho and Stefano Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
- [37] Yijie Guo, Junhyuk Oh, Satinder Singh, and Honglak Lee. Generative adversarial self-imitation learning. *arXiv preprint arXiv:1812.00950*, 2018.
- [38] Karol Hausman, Yevgen Chebotar, Stefan Schaal, Gaurav Sukhatme, and Joseph J Lim. Multi-modal imitation learning from unstructured demonstrations using generative adversarial nets. In *Advances in Neural Information Processing Systems*, pages 1235–1245, 2017.
- [39] Yuke Zhu, Ziyu Wang, Josh Merel, Andrei Rusu, Tom Erez, Serkan Cabi, Saran Tunyasuvunakool, János Kramár, Raia Hadsell, Nando de Freitas, et al. Reinforcement and imitation learning for diverse visuomotor skills. *arXiv preprint arXiv:1802.09564*, 2018.
- [40] Ziyu Wang, Josh S Merel, Scott E Reed, Nando de Freitas, Gregory Wayne, and Nicolas Heess. Robust imitation of diverse behaviors. In *Advances in Neural Information Processing Systems*, pages 5320–5329, 2017.
- [41] D. Kulic, W. Takano, and Y. Nakamura. Online segmentation and clustering from continuous observation of whole body motions. *IEEE Transactions on Robotics*, 25(5):1158–1166, 2009.
- [42] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4346–4354, 2015.
- [43] César Lincoln C Mattos, Zhenwen Dai, Andreas Damianou, Jeremy Forth, Guilherme A Barreto, and Neil D Lawrence. Recurrent gaussian processes. *arXiv preprint arXiv:1511.06644*, 2015.

- [44] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016.
- [45] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*, pages 458–466. IEEE, 2017.
- [46] Philipp Kratzer, Marc Toussaint, and Jim Mainprice. Motion prediction with recurrent neural network dynamical models and trajectory optimization. *arXiv preprint arXiv:1906.12279*, 2019.
- [47] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018.
- [48] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1243–1252. JMLR. org, 2017.
- [49] Yanran Li, Zhao Wang, Xiaosong Yang, Meili Wang, Sebastian Iulian Poiana, Ehtaz Chaudhry, and Jianjun Zhang. Efficient convolutional hierarchical autoencoder for human motion prediction. *The Visual Computer*, 35(6-8):1143–1156, 2019.
- [50] Tyler R. Scott, Karl Ridgeway, and Michael C. Mozer. Adapted deep embeddings: A synthesis of methods for k-shot inductive transfer learning. In *Advances in Neural Information Processing Systems*, volume 2018-Decem, pages 76–85, 2018.
- [51] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [52] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Advances in Neural Information Processing Systems*, pages 7029–7039, 2017.
- [53] Yu-Xiong Wang and Martial Hebert. Learning to learn: Model regression networks for easy small sample learning. In *European Conference on Computer Vision*, pages 616–634. Springer, 2016.
- [54] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- [55] Aude G Billard, Sylvain Calinon, and Rüdiger Dillmann. Learning from humans. In *Springer handbook of robotics*, pages 1995–2014. Springer, 2016.
- [56] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012.
- [57] Reza Shadmehr and Ferdinando A Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of neuroscience*, 14(5):3208–3224, 1994.
- [58] Nicholas Stergiou and Leslie M Decker. Human movement variability, nonlinear dynamics, and pathology: is there a connection? *Human movement science*, 30(5):869–888, 2011.
- [59] Daniel A Braun, Ad Aertsen, Daniel M Wolpert, and Carsten Mehring. Motor task variation induces structural learning. *Current Biology*, 19(4):352–357, 2009.
- [60] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- [61] David J Herzfeld and Reza Shadmehr. Motor variability is not noise, but grist for the learning mill. *Nature neuroscience*, 17(2):149–150, 2014.
- [62] Matthew A Carland, David Thura, and Paul Cisek. The urge to decide and act: Implications for brain function and dysfunction. *The Neuroscientist*, 25(5):491–511, 2019.
- [63] Taisei Sugiyama, Nicolas Schweighofer, and Jun Izawa. Reinforcement meta-learning optimizes visuomotor learning. *bioRxiv*, 2020.
- [64] Anusha Nagabandi, Chelsea Finn, and Sergey Levine. Deep online learning via meta-learning: Continual adaptation for model-based rl. *arXiv preprint arXiv:1812.07671*, 2018.
- [65] David J Herzfeld, Pavan A Vaswani, Mollie K Marko, and Reza Shadmehr. A memory of errors in sensorimotor learning. *Science*, 345(6202):1349–1353, 2014.