Meta-Learning for Avatar Kinematics Reconstruction in Virtual Reality Rehabilitation

Cristian Axenie
Audi Konfuzius-Institut Ingolstadt Lab
Ingolstadt, Germany
cristian.axenie@audi-konfuzius-institut-ingolstadt.de

Armin Becher
Technische Hochschule Ingolstadt
Ingolstadt, Germany
armin.becher@thi.de

Daria Kurz
Interdisciplinary Breast Center
Helios Clinic Munich West
Munich, Germany
daria.kurz@helios-gesundheit.de

Thomas Grauschopf
Technische Hochschule Ingolstadt
Ingolstadt, Germany
thomas.grauschopf@thi.de

***Abstract*—Virtual Reality (VR) sensorimotor rehabilitation is still in infancy but will soon require avatars, digital alter-egos of patients' physical selves. Such embodied interfaces could stimulate patients' perception in a rich and highly customized environment, where sensorimotor deficits, such as in Chemotherapy-Induced Peripheral Neuropathy, could be corrected. In such scenarios, motion prediction is a key ingredient for realistic immersion. Yet, such a task lives under hard processing latency constraints and the inherent variability of human motion. We propose a neural network meta-learning system exploiting the underlying correlations in body kinematics with potential to provide, within latency guarantees, personalized VR rehabilitation. The unsupervised meta-learner is able to extract underlying statistics of the motion data by exploiting data regularities in order to describe the underlying manifold, or structure, of motion under sensorimotor deficits. As avatars are patients' proxies in VR, and the direct extension of themselves into the virtual domain, their digital representations have to be naturally bound to their individual motion patterns and self-perception. Following this goal, we demonstrate, through preliminary experiments the potential of such a learning system for adaptive kinematics estimation in personalized rehabilitation VR avatars. *Index Terms*— Neural Networks, Virtual Reality, Inverse Kinematics, Meta-Learning, Rehabilitation**

I. Introduction

Chemotherapy-induced peripheral neuropathy (CIPN) is a common side effect of cancer treatment and has a large incidence (as high as 83% of breast cancer patients) [1]. Resulting sensory and motor dysfunctions often lead to functional impairments like gait or balance disorders. As an alternative to drug treatments, sensorimotor training has the potential to influence neuromuscular mechanisms for improved balance performance [2], [3]. A practical solution to screen sensorimotor degradation due to CIPN was recently proposed in [4]. Yet, using only physician guided verbal instructions and inertial sensory data, rehabilitation of balance disturbances was not capturing individual patient motion patterns. Focusing on a rehabilitation protocol exclusively based on visual computer-feedback balance training (VCFBT), [5] proved that drug based treatments are not effective for the treatment of CIPN symptoms such as weakness or loss of sensory modalities, including vestibular and proprioceptive. Visual-based treatment was more effective due to the visuo-proprioceptive training, capable of providing sensory support despite the vestibular deficit. Following the visual-based rehabilitation approaches, improvements in software, hardware and reduction in cost have made Virtual Reality (VR) a practical tool for immersive, three-dimensional (3D), multisensory experiences for chronic pain and CIPN patients [6]. Proven by many studies, VR-based training has the potential to influence neuromuscular mechanisms for improving balance and motion performance in patients with sensorimotor deficits [7]. In such applications there is a strong emphasis on motor behaviour by measuring motor output congruence between physical and virtual environments. It is assumed that the fidelity of virtual to physical world movements is vital for rehabilitation, in order to promote the recovery of movement quality, neuromuscular gait and balance properties. In this work we show that such an initial effort is possible through an individualized assessment of the deficit level and a corresponding correction signal for stimulation. We achieved that by employing modern machine learning algorithms capable of extracting underlying correlations in motion kinematics components. Converging to the underlying structure of the rehabilitation task, the meta-learning system extracts the structure and individual peculiarities of patient motion for a personalized, dedicated rehabilitation task, such as in CIPN. This is an initial effort to combine meta-learning with regression algorithms into a flexible platform targeting avatar kinematic reconstruction for effective CIPN VR-based rehabilitation.

II. Problem statement

Adaptive sensorimotor control for gait and balance, in rehabilitation, requires the integration of sensory information in order to assess the position of body in space and subsequently generate forces for controlling body position. This is a challenge in CIPN patients. A provoking issue in creating CIPN rehabilitation environments in VR is determining which

2471-7819/19/$31.00 ©2019 IEEE
DOI 10.1109/BIBE.2019.00117

aspects of the movements made in these environments need to be similar or better than those made in typical physical environments. In both worlds user-environment interaction provides a sense of timing and spatial location that leads to appropriate 3D movements. Such a spatio-temporal estimation task is a perfect candidate for meta-learning which is, basically, capable of learning the manifold on which possible body kinematics configurations lie. After training on multiple similar tasks, the meta-learner would adapt in few shots for a timely and robust position estimation. This is highly relevant, as, in VR avatar-based CIPN rehabilitation, rendering and processing latencies are crucial, impacting the perception of the patient and the results of the rehabilitation. Despite the proliferation of lightweight, high-resolution and high frame-rate VR systems, a remaining obstacle is how to get the user to feel truly immersed in the experience. Under harsh limits of motion-to-photon latency [8], such systems must solve high-resolution inverse kinematics chains of user motion. Yet, predictive markers for CIPN motor treatment strategies are not well-characterized [1], despite the potential VR systems have in providing a highly parametrized and flexible world designed by the physician. The objectives of our work are multiple:

- monitoring and personalized CIPN rehabilitation strategy as meta-learning extracts relations among the different sensory and motor quantities describing patient's motion from a limited number of samples;
- support personalized learning of motion patterns invariants for continuous neuromuscular diagnosis in CIPN;
- to use meta-learning to extract the structure in patient's motion characteristics in the controlled VR lab environment and then transfer to a personal, home-use VR system;
- support traditional neuropathy assessment algorithms to guide the neurotoxic chemotherapy treatments through a motor system assessment.

These objectives materialized in a series of (initial) contributions (i.e. introduced in the next sections):

- the development of a neural network based meta-learning system for sensorimotor correlations extraction;
- the integration of the neural learning system with the VR avatar generation system;
- the deployment of an end-to-end avatar reconstruction framework for precise full-body motion analysis and kinematics extraction;
- the initial tests of the framework in a simple test-bed similar to rehabilitation procedures.

III. Related work

A. Meta-learning approaches

Machine learning models require a large number of training examples, especially in complex tasks such as human kinematics learning. Humans, in contrast, learn new concepts and skills much faster and more efficiently [9]. Consider the problem of learning to control motion of a table tennis paddle. It would be useful if during practice, we learned a model that had a structure with hidden states that could also help us with learning control of a tennis racket. This should be possible because the two objects are rigid-body inertial systems with dynamics that are similar in structure. If we could somehow discover this structure while playing table tennis, it could vastly speed up our learning of tennis.

In this work, we consider the problem of structural learning or meta-learning for precise avatar kinematics reconstruction in CIPN rehabilitation VR. We expect that a good meta-learning model is capable of generalizing well to new tasks and new environments that have never been encountered during training time. This adaptation process is essentially a short learning session with a limited exposure to the new task configurations. Eventually, the adapted model can complete new tasks. In our scenario a good meta-learning model should be trained over a variety of motor tasks and optimized for the best performance on a distribution of tasks, including potentially unseen rehabilitation tasks or patients. Typically such tasks need to be described by both feature vectors and true labels, whereas the highly variable real-world scenarios, such as personalized motor CIPN rehabilitation, are governed by an unsupervised setting.

Despite its infancy, meta-learning comes in three flavours: metric-based, model-based, and optimization-based. The core idea in metric-based meta-learning is similar to nearest neighbours algorithms (i.e., k-NN classifier and k-means clustering) and kernel density estimation. The predicted probability over a set of known labels is a weighted sum of labels of support set samples. The weight is generated by a kernel function, measuring the similarity between two data samples. A successful implementation of the metric based meta-learning paradigm is the Siamese Neural Network (SNN). The architecture is composed of two twin networks and their outputs are jointly trained on top with a function to learn the relationship between pairs of input data samples. In a practical setting, [10] proposed a method to use the SNN for one-shot image classification. Another successful instantiation are the Matching Networks [11] capable to learn a classifier for any given (small) support set. Similar to other metric-based models, the classifier output is defined as a sum of labels of support samples weighted by attention kernel which should be proportional to a sample similarity. Moreover, similar to SNNs, in Relation Networks (RN) [12] the relationship is not captured by a simple L_1 distance in the feature space, but predicted by a CNN classifier and the objective function is MSE loss instead of cross-entropy. Finally, Prototypical Networks [13] use an embedding function to encode each input into a high-dimensional feature vector. A prototype feature vector is defined for every class as the mean vector of the embedded support data samples in this class.

From a different perspective, model-based meta-learning models make no assumption on the form of the probability distribution of the parameters given the input and output. Rather they depend on a model designed specifically for fast learning a model that updates its parameters rapidly with a few training steps. "Memory-Augmented Neural Network" (MANN) is a typical such model using an explicit storage buffer, making it is easier for the network to rapidly incorporate new information and avoid forgetting in the future. Training such a system [14] assumes that the memory is forced to hold information for longer until the appropriate labels are presented later. Another such model are the Meta Networks [15], or MetaNet, a meta-learning model with architecture and training process designed for rapid generalization across tasks. The rapid generalization of MetaNet relies on ”fast weights”. This faster way to learn utilizes one neural network to predict the parameters of another neural network and the generated weights are called ”fast weights”. MetaNet uses a parametric embedding function that encodes raw inputs into feature vectors similar to SNNs.

Looking at optimization-based approaches, deep learning models are representative candidates, learning through back-propagation of gradients. However, gradient-based optimization is neither designed to cope with a small number of training samples, nor to converge within a small number of optimization steps. A representative optimization based meta-learner is the LSTM Meta-Learner [16]. The meta-learner was modelled as a LSTM, because: a) there is a similarity between the gradient-based update in backpropagation and the cell-state update in LSTM; and b) knowing a history of gradients benefits the gradient update. One recent development, the MAML, short for Model-Agnostic Meta-Learning [17], is a fairly general optimization algorithm, compatible with any model that learns through gradient descent. The meta-optimization step relies on second derivatives. In response to recent work in meta-learning [17], [18] proposed training Reptile, a meta-learner on a distribution of similar tasks, in the hope of generalization to novel but related tasks by learning a high-level strategy that captures the essence of the problem it is asked to solve.

B. *VR kinematics tracking*

The kinematic validity of movements recorded with marker-less camera-based motion tracking systems is an area of growing interest because of the development of portable VR systems for upper and lower limb motor rehabilitation. Several of these systems use the Kinect (Microsoft Xbox) camera that reconstructs the user’s gestures using a simplified 15 joint skeleton in real time [19], [20], [21]. Compared to planar joint kinematics recorded with a 24 camera Vicon system, [20] reported that angle errors using Kinect camera tracking ranged from 7.1°to 13.2°for movements in three planes (sagittal, coronal, and transverse). Similarly, [22] determined the extent to which Kinect position tracking matched that recorded by an Optotrak motion tracking system. The mean error for tracking movement of the whole body was 3.9 cm. We complement this overview also in our analysis by comparative evaluation of other tracking systems.

IV. Materials and methods

In the following section we introduce the technical setup used in our preliminary experiments with the meta-learner for VR avatar reconstruction in rehabilitation tasks.

A. *Unsupervised Neural Meta-learning*

Deep neural networks excel in regimes with large amounts of data, but tend to struggle when data is scarce or when they need to adapt quickly to changes in the task. However, many recent meta-learning approaches [17], [18] are extensively hand-designed, either using architectures specialized to a particular application, or hard-coding algorithmic components that constrain how the meta-learner solves the task. We introduce a novel unsupervised meta-learner capable of extracting underlying, unknown, relations among the different sensory streams, by observing and characterizing an user’s avatar reconstruction in VR. The potential of such a system alleviates the need for large labelled datasets (e.g. unable to extract a relevant sample from clinical use) by using unsupervised learning. The core contribution of this work is proposing a novel meta-learning paradigm on top of [23]. The system extracts underlying correlations in incoming sensory streams describing the VR avatar kinematics for a certain motion task. The goal is to learn the structure of such a task in order to accelerate learning for tasks with similar structure. The meta-learning model is based on Self-Organizing Maps (SOM) [24] and Hebbian Learning (HL) as main ingredients for extracting underlying correlations in sensory data describing motion kinematics. In the following, we motivate our choices and design. We employed Self-Organizing Maps (SOMs) in

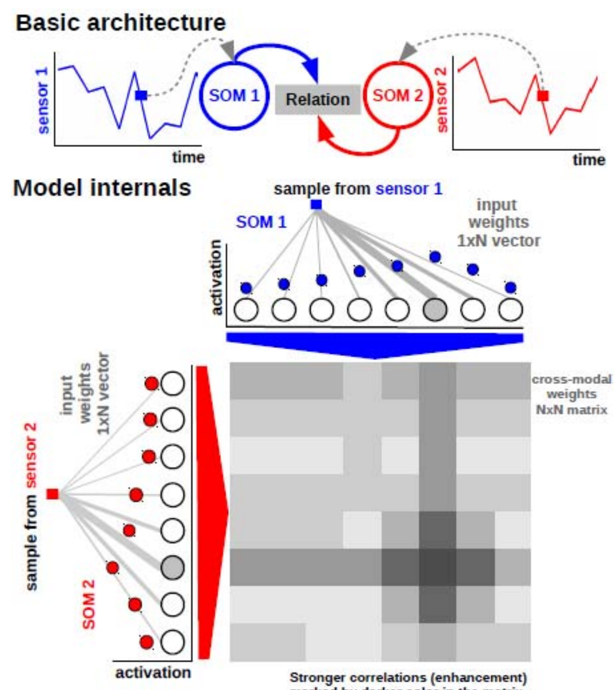

Fig. 1: Basic meta-learning model: SOM extract underlying statistics of the sensorimotor data describing patient’s motion, whereas the HL captures the temporal correlations of such sensorimotor data.

order to extract the statistics of the incoming data and encoding sensory samples in a distributed activity pattern, as shown in Figure 1. Using the SOM distributed representation, the model learns the boundaries of the input data, such that, after relaxation, the SOMs provide a topology preserving representation of the input space. In our approach each neuron not only specialises in representing a certain (preferred) value in the input space, but also learns its own sensitivity (i.e., tuning curve shape). Given an input sample, $s^p(k)$ at time step k, for each i-th neuron in the p-th input SOM, with the preferred value $w^p_{in,i}$ and $\xi^p_i(k)$ tuning curve width, the sensory elicited activation is given by

$$a^p_i(k) = \frac{1}{\sqrt{2\pi}\xi^p_i(k)} exp(\frac{-(s^p(k) - w^p_{in,i}(k))^2}{2\xi^p_i(k)^2}). \quad (1)$$

The winner neuron of the p-th population, $b^p(k)$, is the one which elicits the highest activation given the sensory input at time step k and together with neighbouring cells are updated under a decaying learning rate $\alpha(k)$ as,

$$\Delta w^p_{in,i}(k) = \alpha(k) h^p_{b,i}(k)(s^p(k) - w^p_{in,i}(k)). \quad (2)$$

The meta-learner can also learn the statistics of the input data and encode them in tuning curves. Each neuron's tuning curve is modulated by the spatial location of the neuron, the (Euclidian) distance to the input sample, the interaction kernel size, and the learning rate,

$$\Delta\xi^p_i(k) = \alpha(k) h^p_{b,i}(k)((s^p(k) - w^p_{in,i}(k))^2 - \xi^p_i(k)^2). \quad (3)$$

In a simple example, the meta-learner simultaneously extracts the hidden relation between two sensors and tuning curves (i.e. 3rd power law describing the motion of the filtered left hand VR controller), as shown in Figure 2. Figure 2b shows that higher input probability distributions, as shown in Figure 2a, are represented by a large number of sharp tuning curves, whereas lower or uniform probability distributions are represented by a small number of wide tuning curves. This allows our system to generalize and extract the structure of a task from the data distribution and its temporal dimension. The second component of the meta-learner is the Hebbian Learning network [25]: a fully connected matrix of synaptic connections between neurons in each input SOM, such that the projections propagate between the first SOM units and the second SOM units in the network. It is responsible for extracting the co-activation pattern between the input layers (i.e., SOMs), as shown in Figure 3, and for eventually encoding the learned relation between the sensors, as shown in Figure 2b. Functionally, connections between uncorrelated (or weakly correlated) neurons in each population are suppressed (i.e., darker color-lower value) while correlated neurons' connections are enhanced (i.e., brighter color-higher value). Formally, Hebbian connection weights, $w^p_{cross,i,j}$, between neurons i, j in each of the input SOM population are updated using

$$\Delta w^p_{cross,i,j}(k) = \eta(k)(a^p_i(k) - \overline{a}^p_i(k))(a^q_j(k) - \overline{a}^q_j(k)), \quad (4)$$

In our meta-learner, self-organisation and correlation learning processes evolve simultaneously, such that both representation

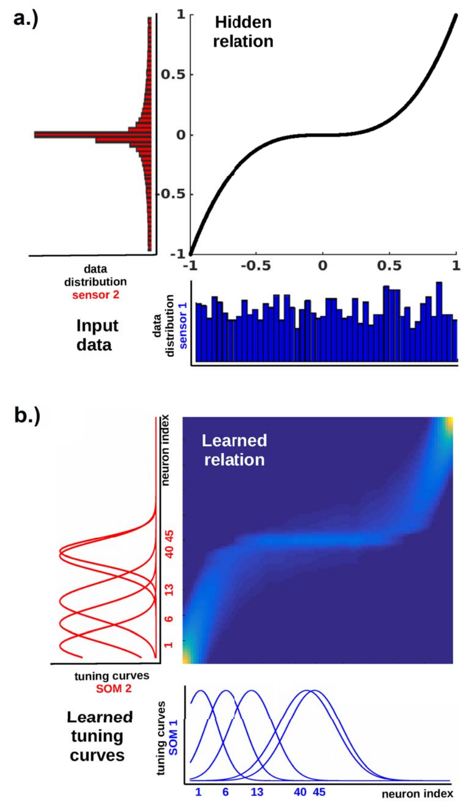

Fig. 2: Basic capabilities: a) learn underlying data distribution and encode it in tuning curves and learn the underlying functional correlation between sensory stream unsupervisedly; b) The structure is encoded in the learn pattern of neural activation.

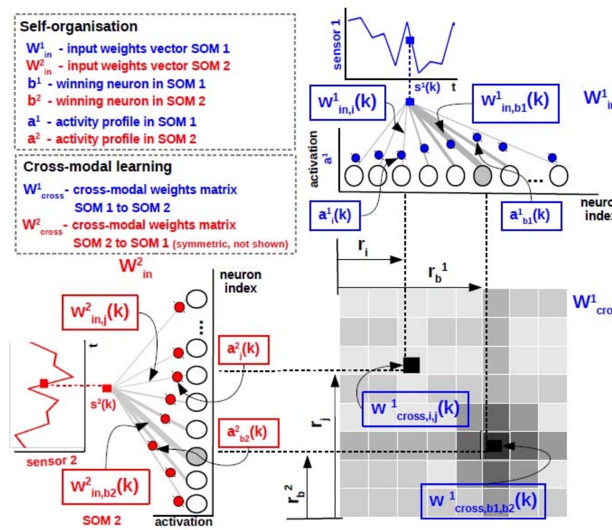

Fig. 3: Model internals.

and correlation pattern are continuously refined to maintain the learned structure. The core idea is to extract rules on how the user motion parameters covary as such variation may not span the entire possible space but lie on a low dimensional manifold - the structure.

B. Technical setup

The experimental setup is composed of several modules, as shown in Figure 4. We used a HTC Vive HMD which comes with two hand controllers. Orientation and position of these devices are tracked via an external infrared tracking setup. Altogether, 18 degrees of freedom (DOFs) are captured by HTCs external tracking system. The first and second modules use the data of the VR tracking system (i.e. global rotation of the headset and two vectors from the users head to the hand positions) to learn the inverse kinematics (IK) of the upper-body, global rotation and position of a user inside VR. The third and fourth modules focus on the RGB video data to reconstruct the lower body pose of the user's avatar. The last module is responsible to render the avatar inside the virtual world from the global trunk rotation, the local rotations for the neck, elbows, shoulders, and both collar joints of the IK regressor output. Quaternions were used to represent 3D-rotations as a point on a hypersphere in a R^4 space. Every rotation in 3D-space can be then expressed by one unambiguous point on the hypersphere. To meet this condition, an additional normalization layer is added right before the output layer of the IK regresor. The network is basically a ResNet-50 trained using the rotational difference between the predicted and the ground-truth rotations (i.e. roll, pitch, yaw) the CMU Database [26]. It contains over 2,500 motion capturing sequences. Within Unity 3D, all motions are mapped to the SMPL skeleton [27] with a fixed default skeleton size. The whole dataset is split in a test set (activities 1 to 30) and a training set (activities 31 to 144). To train the meta-learner, we used the motion data from the VR tracker system, namely VRT (HMD and two hand controllers IR tracker) and the VRI (HMD and two hand controllers inertial data), as shown in Figure 6. Note that the meta-learner is trained on live data from the VR tracking system (i.e. ample rotations, up to 50 degrees roll (0.9 rad) and 34 degrees pitch (0.6 rad)) whereas the IK regressor is trained on a large datasets. This is motivated by the fact that the IK regressor learns the kinematic constraints and underlying chain correlations whereas the meta-learner extract from few samples the structure of the task to make regression faster and comply with the motion-to-photon latencies the VR rehabilitation task requires. In order to evaluate the upper body kinematics our meta-learning adapted regressor is compared against three popular and widely used IK solvers. FABRIK is an IK solving algorithm which avoids the use of direct joint rotations [28]. It iteratively finds joint positions via the location of points on a line. The Cyclic Coordinate Descent CCD [29] algorithm is similar to FABRIK, but instead of finding a point on a line, every single joint in the IK chain gets bent towards the target. Like FABRIK, CCD is an iterative IK solver which terminates when the last joint in the IK chain aligns with the target position. The last IK algorithm we compared against is Limb (Final IK, Unity 3D plug-in). It employs a trigonometric IK solving which tries to heuristically keep the joint configuration in a natural and relaxed configuration. Our overall experimental system uses VIRTOOAIR: VIrtual Reality TOOlbox for Avatar Intelligent Reconstruction processing [23].

V. Experiments and evaluation

A. Single body-part experiment

For the experiments we use tracking data from the HTC Vive VR hand-controllers. In this initial experiment we look at the 10-shot sinewave regression learning task described in [17] where MAML is trained 500 separate models on 500 tasks. Each of their models was initialized randomly and trained on a large amount of data from its assigned task. Similarly, we feed our unsupervised meta-learner with a sinewave motion (the hidden relation between the two inputs in the dataset) of the HTC Vive controller. We remind the reader that we employ unsupervised learning and the meta-learner has no prior information about the underlying function in the data coming from the HTC Vive VR hand-controller, left hand. Our meta learner is able to extract the underlying relation from just few samples running 500 epochs (i.e. meta-iterations) training on 250 separate models (i.e. HTC Vive VR hand-controller generated sinewave parameters, amplitude and phase), as shown in Figure 5, lower panel. Our system is able to extract the data distribution simultaneously for each task, Figure 5 as shown in the tuning curves depicting neuron preferences for a certain task (e.g. 120) presented during training. The data was generated by HTC Vive VR hand-controller, left hand. In order to use the extracted underlying structure in adapting the IK regression, we decode the weight matrix encoding the relation (i.e. structure) using Brent's method [30].

B. Full-body experiment

For our real-world VR avatar reconstruction scenario, we examined how our meta-learner can extract avatar kinematics structure (i.e. rotation of the head angles: roll, pitch) given sensory data, Figure 6, tracking and inertial data from the HTC Vive system used in our setup. Now we extend the experiment by using all HTC Vive VR system data, hand controllers and HMD. The learned structures, for roll and pitch motions, resemble actually to the non-linear functions used in typical modelling approaches (i.e. roll is the arctangent transform applied to the ratio between inertial components on y and z axes). Yet they preserve the irregularities in the cross-sensory relations, as seen in Figure 7. We use the learnt structure from the meta-learner in order to "meta-train" on global rotations (i.e. roll, pitch) for new samples the ResNet-50 of our joint rotation IK regressor, and evaluated it in terms of Mean Per Joint Rotation Error (MPJRE), as shown in Table I. In our last experiment, we evaluate the performance of the meta-learner to improve our joint position IK regressor performance (Mean Per Joint Position Error (MPJPE)). Table II shows how our approach outperforms existing inverse kinematic solutions due

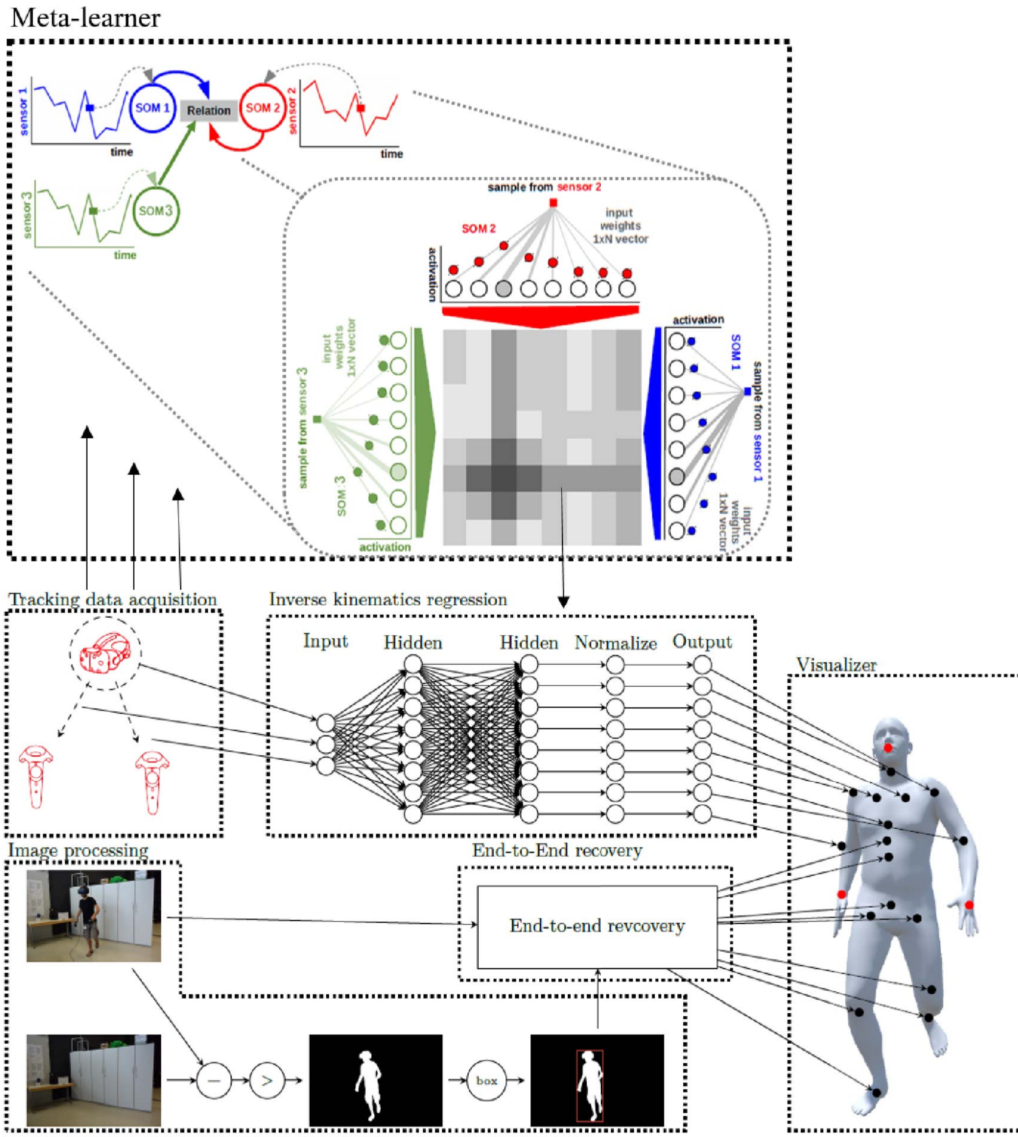

Fig. 4: Experimental system architecture. The meta-learner extracts the structure from few samples and drives the IK regressor to the underlying structure, speeding up adaptation.

TABLE I: Evaluating IK algorithms (joint rotation)

Method	MPJRE
Limb IK Solver	67.9°
CCD [29]	105.8°
FABRIK [28]	88.4°
Meta-Learner (ours)	**13.9°**

TABLE II: Evaluating IK algorithms (joint position)

Method	MPJPE
Limb IK Solver	29.5 mm
CCD [29]	54.7 mm
FABRIK [28]	43.7 mm
Meta-Learner (ours)	**25.8 mm**

to its learnt soft constrains of human motion (i.e. hidden in sensory correlations) learnt by the meta-learner from the manifold of real-world human poses. The broad range of experiments in this section were meant to demonstrate the capabilities of our unsupervised meta-learner to extract the underlying relations among the sensory data characterizing user's avatar motion kinematics in VR. The preliminary results show that such a meta-learning algorithm has potential to handle highly-nonlinear tasks supporting fast adaptation when facing with the inherent variability, such as in CIPN rehabilitation programs. [1].

[1]Source code available at: https://gitlab.com/akii-microlab/virtooair

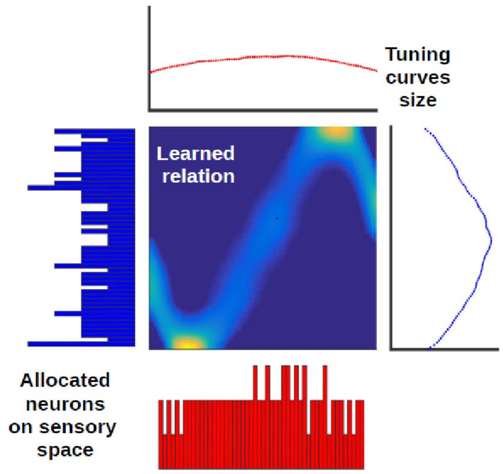

Fig. 5: Meta-learning regression task in 1D: Learning additional characteristics of the data, i.e. distribution encoded through tuning curves width.

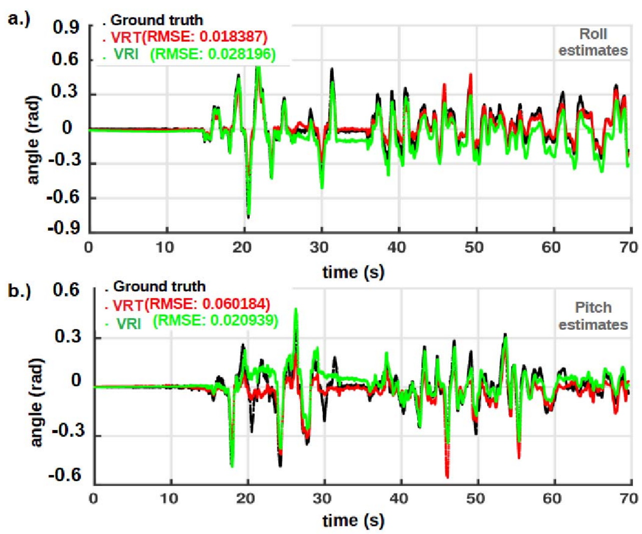

Fig. 6: Input data for learning rotation structure from VRT (VR controllers and HMD tracking) and VRI (VR controllers and HMD inertial): a) roll data from VRT and VRI vs Ground truth; b) pitch data from VRT and VRI vs Ground truth

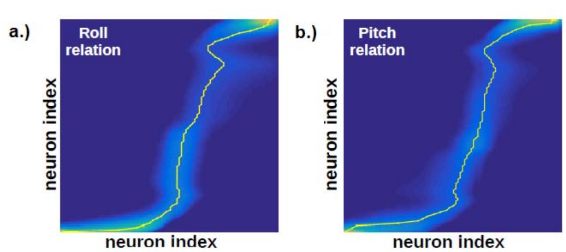

Fig. 7: Meta-learning task for rotation from VRT (VR controllers and HMD tracking) and VRI (VR controllers and HMD inertial): a) learned roll structure; b) learned pitch structure

VI. Discussion

In the proposed work we developed a neural network based meta-learning system for sensorimotor correlations extraction, as shown in Single body-part and Full-body experiments. The system is capable of attaining good performance for the initial tests of such a framework in a simple test-bed similar to rehabilitation procedures. This was achieved thorough the integration of the neural learning system, at the core of our system, with the end-to-end avatar reconstruction framework for precise full-body motion analysis and kinematics [23]. These contributions will guide the primary goal of motor CIPN rehabilitation, namely to help the individual return to functional performance of daily life activities through the recovery or compensation of lost motor skills [31] through personalized avatar-based rehabilitation. Our scientific rationale for using VR technology within CIPN rehabilitation is supported by the field of motor learning [32]. Therapeutic interventions using VR systems are attractive rehabilitation options because the motor learning variables underlying experience dependent neuroplasticity are inherent attributes of VR systems [33]. Moreover, our learning system capable of extracting the underlying sensory correlations in the patient's kinematics in VR can support personalized CIPN rehabilitation [34]. The proposed meta-learning system is capable of learning from few samples the underlying manifold on which the sensory data, describing patient motions, lies and use this to adapt to novel tasks. During learning our system doesn't only adjust the parameters of the default generative model but learns a new one, through a structure specific facilitation [9]. Our contribution is a novel unsupervised meta-learning system that improves end-to-end VR avatar kinematics tracking [23] in CIPN rehabilitation scenarios. Here, the latency constraints are high, the variability of the patients is high and the data is scarce.

VII. Conclusions

In Chemotherapy-Induced Peripheral Neuropathy(CIPN), the most common side effects of cancer treatments, the maintenance of an upright posture is highly dependent on feedback from the somatosensory and vision systems; therefore, sensory losses alone can greatly impair balance and gait. CIPN motor retraining involves the manipulation of variables which is where VR therapies prime. Since the introduction of VR interfaces into the rehabilitation toolkit, considerable work has been done to validate movements made in various environments. However, there is still much work to be done. Our system provides a platform for personalized motor rehabilitation (i.e. as in CIPN) including an objective measure of patient motion kinematics to support clinicians to better detect CIPN symptoms compared to relying solely on patient-reported measures. Functional measures, such as those described in our VR avatar experiments (i.e. for hand and full-body), may aid oncology rehabilitation clinics in reducing patients long-term motor deficits through learned sensorimotor correlations for individualized, personalized rehabilitation. The proposed meta-learning system speeds-up adaptation in the face of high patient motion variability and latency constraints towards a truly personalized experience. Our long-term goal is to train the meta-learning VR avatar system in the lab setup and deploy it for personal home-based rehabilitation. VR devices becoming commodity will allow for continuous monitoring and personalized rehabilitation. The clinical part of the team has developed an initial experiment design and validated several novel metrics to assess motor decline in breast cancer survivors CIPN. Employing the proposed meta-learner the study aims at determining and exploiting those rehabilitation relevant biomechanical variabilities for VR rehabilitation programs.

References

[1] S. M. Monfort, X. Pan, R. Patrick, B. Ramaswamy, R. Wesolowski, M. J. Naughton, C. L. Loprinzi, A. M. Chaudhari, and M. B. Lustberg, "Gait, balance, and patient-reported outcomes during taxane-based chemotherapy in early-stage breast cancer patients," *Breast cancer research and treatment*, vol. 164, no. 1, pp. 69–77, 2017.

[2] F. Duregon, B. Vendramin, V. Bullo, S. Gobbo, L. Cugusi, A. D. Blasio, D. Neunhaeuserer, M. Zaccaria, M. Bergamin, and A. Ermolao, "Effects of exercise on cancer patients suffering chemotherapy-induced peripheral neuropathy undergoing treatment: A systematic review," *Critical Reviews in Oncology/Hematology*, vol. 121, 2018.

[3] S. Kneis, A. Wehrle, K. Freyler, K. Lehmann, B. Rudolphi, B. Hildenbrand, H. H. Bartsch, H. Bertz, A. Gollhofer, and R. Ritzmann, "Balance impairments and neuromuscular changes in breast cancer patients with chemotherapy-induced peripheral neuropathy," *Clinical Neurophysiology*, vol. 127, no. 2, pp. 1481 – 1490, 2016.

[4] M. Zahiri, K. M. Chen, H. Zhou, H. Nguyen, B. T. Workeneh, S. V. Yellapragada, Y. H. Sada, M. Schwenk, and B. Najafi, "Using wearables to screen motor performance deterioration because of cancer and chemotherapy-induced peripheral neuropathy (cipn) in adults - toward an early diagnosis of cipn," *Journal of Geriatric Oncology*, 2019.

[5] S. Cammisuli, E. Cavazzi, E. Baldissarro, and M. Leandri, "Rehabilitation of balance disturbances due to chemotherapy-induced peripheral neuropathy: a pilot study." *European journal of physical and rehabilitation medicine*, vol. 52, no. 4, pp. 479–488, 2016.

[6] V. C. Tashjian, S. Mosadeghi, A. R. Howard, M. Lopez, T. Dupuy, M. Reid, B. Martinez, S. Ahmed, F. Dailey, K. Robbins, B. Rosen, G. Fuller, I. Danovitch, W. IsHak, and B. Spiegel, "Virtual reality for management of pain in hospitalized patients: Results of a controlled trial," *JMIR Ment Health*, vol. 4, no. 1, p. e9, Mar 2017.

[7] M. Schwenk, G. S. Grewal, B. Honarvar, S. Schwenk, J. Mohler, D. S. Khalsa, and B. Najafi, "Interactive balance training integrating sensor-based visual feedback of movement performance: a pilot study in older adults," *Journal of neuroengineering and rehabilitation*, vol. 11, no. 1, p. 164, 2014.

[8] M. E. St. Pierre, S. Banerjee, A. W. Hoover, and E. R. Muth, "The effects of 0.2hz varying latency with 20–100ms varying amplitude on simulator sickness in a helmet mounted display," *Displays*, vol. 36, pp. 1–8, 2015.

[9] D. A. Braun, A. Aertsen, D. M. Wolpert, and C. Mehring, "Motor task variation induces structural learning," *Current Biology*, vol. 19, no. 4, pp. 352–357, 2009.

[10] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition," in *ICML Deep Learning Workshop*, vol. 2, 2015.

[11] O. Vinyals, C. Blundell, T. Lillicrap, D. Wierstra *et al.*, "Matching networks for one shot learning," in *Advances in neural information processing systems*, 2016, pp. 3630–3638.

[12] F. Sung, Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, "Learning to compare: Relation network for few-shot learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 1199–1208.

[13] J. Snell, K. Swersky, and R. Zemel, "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4077–4087.

[14] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap, "Meta-learning with memory-augmented neural networks," in *International conference on machine learning*, 2016, pp. 1842–1850.

[15] T. Munkhdalai and H. Yu, "Meta networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2554–2563.

[16] S. Ravi and H. Larochelle, "Optimization as a model for few-shot learning," 2016.

[17] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1126–1135.

[18] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," 2018.

[19] T. Dutta, "Evaluation of the kinect sensor for 3-d kinematic measurement in the workplace," *Applied ergonomics*, vol. 43, no. 4, pp. 645–649, 2012.

[20] A. Fern'ndez-Baena, A. Susín, and X. Lligadas, "Biomechanical validation of upper-body and lower-body joint movements of kinect motion capture data for rehabilitation treatments," in *2012 fourth international conference on intelligent networking and collaborative systems*. IEEE, 2012, pp. 656–661.

[21] P. Zhang, K. Siu, J. Zhang, C. K. Liu, and J. Chai, "Leveraging depth cameras and wearable pressure sensors for full-body kinematics and dynamics capture," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 6, p. 221, 2014.

[22] G. Tao, P. S. Archambault, and M. Levin, "Evaluation of kinect skeletal tracking in a virtual reality rehabilitation system for upper limb hemiparesis," in *2013 international conference on virtual rehabilitation (ICVR)*. IEEE, 2013, pp. 164–165.

[23] G. T. Becher A., Axenie C., "Virtooair: Virtual reality toolbox for avatar intelligent reconstruction," in *Multimodal Virtual and Augmented Reality Workshop (MVAR) at 2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 2018.

[24] T. Kohonen, "Self-organized formation of topologically correct feature maps," *Biological cybernetics*, vol. 43, no. 1, pp. 59–69, 1982.

[25] Z. Chen, S. Haykin, J. J. Eggermont, and S. Becker, *Correlative learning: a basis for brain and adaptive systems*. John Wiley & Sons, 2008, vol. 49.

[26] Carnegie Mellon University, "Cmu graphics lab motion capture database." [Online]. Available: http://mocap.cs.cmu.edu/

[27] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "Smpl: a skinned multi-person linear model," *ACM Transactions on Graphics*, vol. 34, no. 6, pp. 1–16, 2015.

[28] A. Aristidou and J. Lasenby, "Fabrik: A fast, iterative solver for the inverse kinematics problem," *Graphical Models*, vol. 73, no. 5, pp. 243–260, 2011.

[29] L.-C. Wang and C. C. Chen, "A combined optimization method for solving the inverse kinematics problems of mechanical manipulators," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 4, pp. 489–499, 1991.

[30] R. P. Brent, *Algorithms for minimization without derivatives*. Courier Corporation, 2013.

[31] R. A. Schmidt, T. D. Lee, C. Winstein, G. Wulf, and H. N. Zelaznik, *Motor control and learning: A behavioral emphasis*. Human kinetics, 2018.

[32] M. K. Holden, "Virtual environments for motor rehabilitation," *Cyberpsychology & behavior*, vol. 8, no. 3, pp. 187–211, 2005.

[33] M. F. Levin, "Can virtual reality offer enriched environments for rehabilitation?" *Expert review of neurotherapeutics*, vol. 11, no. 2, pp. 153–155, 2011.

[34] C. Winstein, R. Lewthwaite, S. R. Blanton, L. B. Wolf, and L. Wishart, "Infusing motor learning research into neurorehabilitation practice: a historical perspective with case exemplar from the accelerated skill acquisition program," *Journal of neurologic physical therapy: JNPT*, vol. 38, no. 3, p. 190, 2014.