

# Weakly-supervised Learning of Human Dynamics

Petrissa Zell<sup>1</sup>, Bodo Rosenhahn<sup>1</sup>, and Bastian Wandt<sup>1</sup>

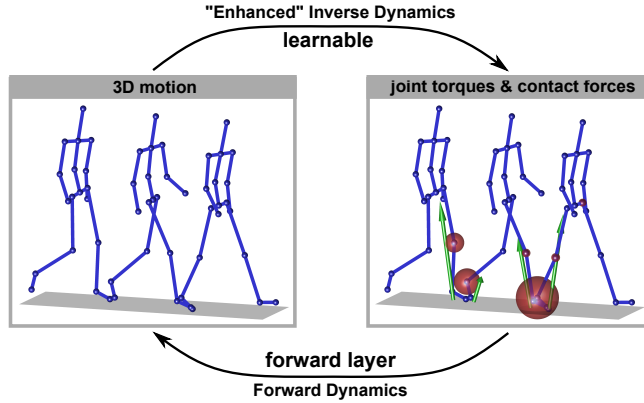
Leibniz University Hannover, 30167 Hannover, Germany  
`zell@tnt.uni-hannover.de`

**Abstract.** This paper proposes a weakly-supervised learning framework for dynamics estimation from human motion. Although there are many solutions to capture pure human motion readily available, their data is not sufficient to analyze quality and efficiency of movements. Instead, the forces and moments driving human motion (*the dynamics*) need to be considered. Since recording dynamics is a laborious task that requires expensive sensors and complex, time-consuming optimization, dynamics data sets are small compared to human motion data sets and are rarely made public. The proposed approach takes advantage of easily obtainable motion data which enables weakly-supervised learning on small dynamics sets and weakly-supervised domain transfer. Our method includes novel neural network (NN) layers for forward and inverse dynamics during end-to-end training. On this basis, a cyclic loss between pure motion data can be minimized, i. e. no ground truth forces and moments are required during training. The proposed method achieves state-of-the-art results in terms of ground reaction force, ground reaction moment and joint torque regression and is able to maintain good performance on substantially reduced sets.

**Keywords:** artificial neural networks, human motion, forward dynamics, inverse dynamics, weakly-supervised learning, domain transfer

## 1 Introduction

Inverse dynamics describes the process of estimating net moments of force acting across skeletal joints from the three dimensional motion of the human skeleton and a set of exterior contact forces and moments. The obtained net moments are generally called joint torques (JT) and are of central interest in important fields, such as diagnostics of locomotor disorders, rehabilitation and prostheses design [10], [12], [23]. The JT cannot be measured non-intrusively and therefore need to be derived using computationally expensive optimization techniques. The common approaches are inverse and forward dynamics optimization. In inverse dynamics optimization, the kinematics of a model are optimized to match the target motion, while JT are inversely calculated based on the interaction with the ground and the model kinematics. For this step, 2nd order time derivatives are necessary, making the approach sensitive to noisy motion capture data.



**Fig. 1.** An *Inverse Dynamics* method calculates joint torques based on an observed motion and contact forces. The enhanced task, that is implemented as a NN in this work, infers both, the joint torques and the exterior forces. The *Forward Dynamics* step yields a simulated motion based on the acting forces and moments. With our forward layer this step can now be integrated into NN training.

Forward dynamics optimization considers the reverse problem: The applied forces drive the motion of the human body. The equations of motion and the current motion state define an initial value problem, that can be solved by numerical integration, yielding a simulated motion. Included into an optimization framework, the forward dynamics step is utilized to find the optimal JT that generate a simulated motion with minimal distance to the captured motion. The necessary integration during each optimization step causes high computational complexity. Furthermore, both approaches, inverse and forward dynamics, require the measurement of exterior forces. In the case of locomotion, these exterior influences are the ground reaction forces and moments (GRF/M) acting at both feet. Hence, the analysis of human dynamics using conventional optimization techniques is restricted to a controlled laboratory setting with force plates to measure the GRF/M. To avoid the described problems and to achieve motion analysis in the wild, the exterior forces and the interior JT can be estimated directly from observed motions using machine learning techniques, e.g. neural networks (NN).

Deep learning of human dynamics requires large data sets of human kinematics (3D motion of a kinematic model), exterior GRF/M and the driving JT. Unfortunately, corresponding data sets are few and often very restricted in terms of size and included motion types. To solve this problem, we propose a weakly-supervised deep learning method, that realizes NN training on small dynamics sets and domain transfer to new motion types, even without any ground truth of the acting forces. More precisely, our goal is to design a training process that is less depending on a large amount of dynamics data than a supervised baseline NN that learns to regress forces and moments from motion input. The proposed approach is called **Dynamics Network** and includes two novel dy-

namics layers: The forward (dynamics) layer and the inverse (dynamics) layer. The network replaces the traditional inverse dynamics task and enhances it by not only inferring the JT but also the GRF/M, as illustrated in Figure 1. Subsequently, the forward layer solves the initial value problem given by the initial motion state and the network output. This results in a simulated motion that can be compared to the input motion, which gives rise to a loss purely defined on motion data. A full cycle, as illustrated in Figure 1 is executed. For further control over the learned forces and moments we include the inverse layer, that penalizes GRF/M which do not match the observed accelerations. In contrast to the forward layer, it considers the ground reaction independently from JT, which allows for a decoupled control of both variables during training.

We demonstrate the benefits of our Dynamics Network by reducing the part of the used training sets, that include dynamical information, relying on the networks’ own dynamics layers. Furthermore, we use our method to realize weakly- and unsupervised transfer learning between the related motion types, walking and running. We focus on locomotion since it represents the most important form of movement.

In summary, the contributions of this work are:

1. A data set of 3D human kinematics with force plate measurements including various motion types.
2. A novel forward dynamics layer that integrates the equations of motion and enables the minimization of a pure motion loss.
3. A novel inverse dynamics layer that propagates forces and moments along the kinematic chain to measure the match between GRF/M and segment accelerations.
4. Since the dynamics layers allow for training on pure motion samples without force information, we use this capacity in weakly-supervised learning and domain transfer between locomotion types.

## 2 Related work

A laboratory setting with embedded force plates enforces strong restrictions on captured motions. Therefore, researchers increasingly exploit artificial NN to estimate ground reactions. The exterior forces are predicted based on 3D motions [3], [5], [15], [22], accelerometer data [16, 17] or pressure insoles [25], [28]. The seminal work by Oh et al. [22] includes a fully connected feed forward NN to solve the GRF ambiguity during double support. Based on the complete ground reaction information, the JT are inferred using a standard inverse dynamics method. An interesting approach proposed by Johnson et al. [15] encodes marker trajectories as RGB-images to make use of pre-trained CNNs for image classification.

Compared to GRF/M regression, relatively few works dealt with machine learning for the estimation of JT in human motion [14], [33]. A method by Lv et al. [18] utilizes Gaussian mixture models to learn contact and torque priors that are included in a maximum a-posteriori framework. The priors help to lead the inverse dynamics optimization to realistic force and torque profiles. In a

previous work [34], different machine learning algorithms are compared for the GRF/M and JT regression. The focus lies on contrasting end-to-end models to a hierarchical approach that subdivides the task into gait phase classification and regression.

It is worth mentioning, that machine learning for the inverse dynamics problem is also part of robotics research. Here, the common goal is to learn robotic control, such as JT, for trajectory planning. The latest works predominantly rely on recurrent NN and reinforcement learning [8], [11], [21], [29].

In 2018, a differentiable physics engine for the inclusion into deep learning has been introduced and tested on simple simulated systems [2]. A related work proposes a comparable physics engine and includes it in an RNN for learning of robotics control [7]. These models, however, are not sufficient to represent the human kinematic chain. While robotics research focuses on the efficient execution of movement, the interaction with the environment and real-time application, the inverse dynamics problem for human motion is characterized by the complexity of the human locomotor system, which results in high data variability. Therefore, the focus primarily lies on the accurate estimation of forces and moments, in spite of the complex nature of the system, which makes regressors for human motion analysis highly dependent on large data sets.

The usage of cycle-consistency has already benefited visual correspondence tasks like temporal video alignment and cross-domain mapping of images, etc. [9], [13], [26], [31], [36]. The success of these methods has motivated us to apply cycle-consistency to human dynamics learning.

To the best of our knowledge, this work is the first to present differentiable NN layers for forward, as well as inverse dynamics of human motion. In contrast to all prior works, the proposed method can learn human dynamics in a weakly-supervised setting, without depending on complete dynamics ground truth. Our cycle consistent approach allows for the formulation of a pure motion loss, and thus drastically enlarges the usable data pool, even allowing for domain transfer between motion types.

### 3 Human motion data sets

Deep learning of human dynamics demands data of the observed 3D motion, the acting GRF/M and the driving JT. We recorded 195 walking and 75 running sequences performed by 22 subjects of different gender and body proportions (demographic information can be found in the supplementary material). The data was recorded using a Vicon motion capture system with synchronized AMTI force plates, embedded in the ground. An inverse kinematics algorithm was conducted to fit the skeleton of our physical model to the captured marker trajectories. In order to simplify the calculations of the forward layer, we apply a leg model with one torso segment to approximate the motion and the inertial properties of the whole upper body. This kind of model is typically used in locomotion analysis [32]. It's kinematics are completely represented by the

generalized model coordinates  $\mathbf{q}$  that consist of 6 global coordinates and 18 joint angles.

The GRF  $\mathbf{f}_c$  and GRM  $\mathbf{m}_c$  are calculated from the measured force plate data. Together they build the GRF/M vector

$$\begin{aligned}\mathbf{F}_c &= [\mathbf{f}_{c_l}, \mathbf{f}_{c_r}, \mathbf{m}_{c_l}, \mathbf{m}_{c_r}]^T, \\ \mathbf{m}_{c_i} &= \mathbf{d}_{cop_i} \times \mathbf{f}_{c_i} + \mathbf{t}_{z_i}, \quad i = l, r\end{aligned}\tag{1}$$

with the vector  $\mathbf{d}_{cop}$  pointing from the foot center of mass to the center of pressure of the applied reaction force and the torsional torque  $\mathbf{t}_z$ . The index  $i = l, r$  indicates the foot segment, where the respective part of  $\mathbf{F}_c$  is applied.

Based on the kinematics and GRF/M we execute a forward dynamics optimization to receive the non-measurable JT  $\boldsymbol{\tau}$ . The applied forward dynamics step is equivalent to our forward layer, so that the resulting JT conform with the layer dynamics during training. We employ a sliding window approach for this pre-processing step, as well as for the network training and application. This way, the parameter space is decreased and the convergence is accelerated.

To further reduce the number of parameters, we apply polynomial fits to all relevant data types. This approximation also facilitates the learning of inverse dynamics, since the networks are not required to explicitly model temporal context. Without the polynomials, additional smoothness losses would be necessary to achieve continuous forces and moments during forward dynamics optimization, as well as NN training. The kinematics  $\mathbf{q}$  and the GRF/M  $\mathbf{F}_c$  are approximated by 3rd order polynomials. For the JT  $\boldsymbol{\tau}$  we find, that linear approximations yield the best optimization results in a forward dynamics setting. The resulting polynomial coefficients are denoted by  $\gamma_q$ ,  $\gamma_f$  and  $\gamma_\tau$ , respectively. Apart from these coefficients, each forward dynamics simulation is depending on the subject specific segment dimensions  $\mathbf{l}_{sub}$ . Together these parameters build a sample in our human motion data set:

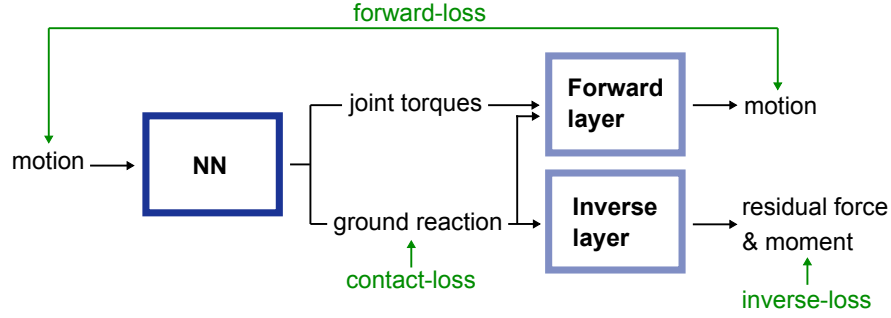
$$[\gamma_q, \mathbf{l}_{sub}, \gamma_f, \gamma_\tau].\tag{2}$$

Using these approximations, we still achieve a representation of non-continuous contact by means of an overlapping window approach, that allows a discretization below window length.

Due to the restrictions introduced by the localized force plate measurements, approximately a third of the data set contains GRF/M. If the forward dynamics optimization converges to an insufficient minimum the corresponding torques cannot be included, so that even a smaller part of the data includes JT. We divide our data set accordingly into a pure motion subset, a motion and GRF/M subset and a subset with the complete data. These subsets are referred to as *motion-set*, *ground-reaction-set* and *torque-set*, respectively. To investigate a further scenario, that requires less supervision and measurements, we define a *contact-set*, which contains motion states and binary information about the ground contact, i.e. which foot is in contact with the ground. The presented data sets are used in different training modes that represent various levels of supervision.

## 4 Dynamics Network

The structure of our Dynamics Network is presented in Figure 2. A fully connected NN executes the extended inverse dynamics task from motion to JT and GRF/M. The network has a bottle-neck structure with 5 fully connected layers, containing about 5800 parameters in total and Leaky-ReLu activations [19]. The forward dynamics step is implemented as a NN layer (forward layer) and yields simulated motion states based on the network output. A detailed description is given in Section 4.1. Combined, the NN and the forward layer build a cycle that enables the minimization of a loss between motion states. The additional inverse layer measures the consistency of the input motion with the predicted GRF/M by propagating forces and moments along the kinematic chain and calculating the residuals at the last segment. A complete description follows in Section 4.2. Furthermore, the network can be trained in a supervised manner, using a mean squared error (MSE) on the predicted GRF/M and JT. This approach will be used as a baseline. Since the corresponding ground truth data is not always available, our dynamics layers can be used to simulate and control it in a weakly-supervised setting. To gradually reduce the level of supervision, we implement a contact-loss, that penalizes forces during time frames with no ground contact. It only requires binary information.



**Fig. 2.** Schematic structure of the Dynamics Network. The output of the NN is processed using forward and inverse layer to accomplish training without GRF/M and JT ground truth.

Our methods operate with four different loss functions, that can be activated separately or in combination. On this basis, we define and compare different training modes that realize various levels of supervision and operate on different subsets of the data: Starting with full supervision, the baseline method uses the MSE of the predicted parameters  $\gamma_f^{\text{pred}}$  and  $\gamma_\tau^{\text{pred}}$  for GRF/M and JT, respectively:

$$L_{\text{mse}} = \|\gamma_f^{\text{pred}} - \gamma_f^{\text{true}}\|_2^2 + \|\gamma_\tau^{\text{pred}} - \gamma_\tau^{\text{true}}\|_2^2. \quad (3)$$

The baseline network is exclusively trained on the torque-set and the ground-reaction-set.

In the next training mode, the network is additionally trained using the contact-set with a reduced level of supervision. We assume, that only binary information about the contact state is available in this subset. The corresponding training mode is referred to as *contact-forward-inverse-training* or in short *cFI-training*. During this mode, the corresponding network (*cFI-net*) is trained in an alternating procedure: If data from the contact-set is chosen as input, the procedure switches between the contact-loss, the inverse-loss and the forward-loss. Based on Eq. (2), we define the contact-loss as

$$L_{\text{contact}} = \|\mathbf{F}_{c_i}(c_i = 0)\|_2^2, \quad (4)$$

with the contact state  $c_i = 1$  if foot  $i$  is in contact with the ground and  $c_i = 0$ , otherwise. If data from the torque-set or the ground-reaction-set is chosen, the MSE of Eq. (3) is minimized, similarly to the baseline method.

The third training mode is termed *forward-training* or *F-training*. In addition to supervised training on the torque-set and the ground-reaction-set, the network, *F-net*, is trained in an unsupervised setting, using the motion-set and minimizing the forward-loss.

In all training modes the included losses are minimized alternately. To balance the influence of different loss functions, we use adaptive weighting, that is updated after each epoch according to the ratio of the observed average values during the last epoch.

In Section 5 we compare the proposed training modes regarding their performance on small data sets and their capability to perform domain transfer from walking to running motions.

#### 4.1 Forward dynamics layer

In this section, we describe our forward dynamics simulation and the implementation as a NN layer. We choose a simple model and a basic numerical integration technique in order to maintain relatively low computational complexity. This is necessary to facilitate the integration in NN training.

Our human body model is a leg model with one additional segment that presents the mean upper body kinematics. As mentioned before, the kinematic state of the model is fully represented by the generalized coordinates  $\mathbf{q}$  and their derivatives  $\dot{\mathbf{q}}$ . The generalized coordinates include 18 joint degrees of freedom that are effected by the JT  $\boldsymbol{\tau}$ . The GRF/M  $\mathbf{F}_c$  introduced in Eq. (2) is applied at the center of mass of the respective foot segment. Each segment is associated with a mass and a tensor of inertia that are approximated from the segment shape and literature values of the population [32]. For this step, we model the segments as simple geometric shapes with constant density. The segment volume is scaled according to the length based on average scaling factors of the data set. With this assumption, the segment lengths  $\mathbf{l}_{sub}$  completely describe the body model.

We formulate the equation of motion using the TMT-method [24] similar to [4], [35]. The resulting 2nd order differential equation has the following form:

$$\mathcal{M}(\mathbf{q}_t, \mathbf{l}_{sub})\ddot{\mathbf{q}}_t = \mathcal{F}(\mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{F}_{c_t}, \boldsymbol{\tau}_t, \mathbf{l}_{sub}) \quad (5)$$

Here,  $\mathcal{M}$  is the generalized inertia matrix and the right hand side is the sum of all acting forces. For better readability we drop the time frame index  $t$  in the remaining part of this work. The problem is reformulated as a 1st order differential equation by introducing the state vector  $\mathbf{x} = [\mathbf{q}^T, \dot{\mathbf{q}}^T]^T$ :

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{q}} \\ \mathcal{M}(\mathbf{q}, \mathbf{l}_{sub})^{-1} \mathcal{F}(\mathbf{x}, \mathbf{F}_c, \boldsymbol{\tau}, \mathbf{l}_{sub}) \end{bmatrix} \quad (6)$$

Together with an initial state vector  $\mathbf{x}_0$  this equation gives rise to an initial value problem, that can be solved by numerical integration.

During integration an acceleration error propagates with the squared integration time. The same applies to an error in JT and GRF/M. In order to reduce this high sensitivity for NN training, we propose to apply a damping factor to  $\dot{\mathbf{x}}$ . The damping is based on the standard deviation  $\sigma_{\dot{\mathbf{x}}}$  of the absolute velocities and accelerations found in the training set and on the maximum absolute values  $\mathbf{m}_{\dot{\mathbf{x}}}$  that occur in the current sample. For each component  $j$  the damping is set to

$$d_j = \exp \left\{ -\max \left( \frac{|\dot{x}_j| - m_{\dot{x},j} - k\sigma_{\dot{x},j}}{k\sigma_{\dot{x},j}}, 0 \right) \right\}, \quad (7)$$

where  $\dot{x}_j$  is the undamped component of the current integration step. The resulting vector  $\mathbf{d}$  is included by building the Hadamard product to result in the damped equation of motion:

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{\mathbf{q}} \\ \mathcal{M}^{-1} \mathcal{F} \end{bmatrix} \odot \mathbf{d}. \quad (8)$$

The parameter  $k$  in Eq. (7) determines the steepness of the damping curve and the starting point of the decrease from one. In simple terms, it broadens the region of acceptable accelerations to  $k$  times of the standard deviation. The value is heuristically set to  $k = 10$ . We found, that this setting results in stable simulations that can still be optimized during training.

For numerical integration we apply Euler's method with constant step size to keep the computation time as small as possible. Consequently, our forward layer can be seen as a function  $FD$ , that executes  $n = (\text{window size} - 1)$  Euler steps. It receives the input parameters

$$\mathbf{p} = (\mathbf{x}_0, \mathbf{l}_{sub}, \mathbf{F}_{c_{1...n}}, \boldsymbol{\tau}_{1...n}, \mathbf{m}_{\dot{\mathbf{x}}}) \quad (9)$$

and outputs the simulated motion states

$$\mathbf{x}_{1...n}^{sim} = FD(\mathbf{p}). \quad (10)$$

Based on this result we define the forward loss as

$$L_{\text{forward}} = \text{MSE}(\mathbf{x}_{1...n}^{sim}, \mathbf{x}_{1...n}^{true}) + \alpha \|\mathbf{d} - \mathbf{1}\|_1, \quad (11)$$

with weighting factor  $\alpha = 1$  and an  $L_1$ -loss to penalize damping factors smaller than one.

For the back propagation of  $L_{\text{forward}}$  during NN training, the gradients of the output states with respect to the input of the layer need to be known. This can either be achieved using automatic differentiation included in most deep learning frameworks or by explicit calculation using sensitivity analysis. The latter approach is described in the supplementary material.

## 4.2 Inverse dynamics layer

The inverse layer receives the ground truth motion and the predicted GRF/M as input and propagates forces and moments along the kinematic chain in a bottom-up procedure. The calculation starts at the centers of mass of the model’s feet, where GRF/M are applied. Each segment is considered in a free body diagram to deduce the forces and moments at the proximal joint based on the acting forces and moments at the distal joint and the linear and angular acceleration of the segment. For segment  $s$  the force  $\mathbf{F}_p$  effecting the proximal joint is given by

$$\mathbf{F}_p = m_s(\mathbf{a}_s - \mathbf{g}) - \mathbf{F}_d, \quad (12)$$

with the distal force  $\mathbf{F}_d$ , the segment mass  $m_s$ , the segment acceleration  $\mathbf{a}_s$  and the gravitational acceleration  $\mathbf{g}$ . The moment  $\mathbf{M}_p$  acting across the proximal joint can be calculated in a related manner:

$$\mathbf{M}_p = \mathbf{I}_s \boldsymbol{\alpha}_s - \mathbf{M}_d - \sum_{j=p,d} \mathbf{r}_j \times \mathbf{F}_j, \quad (13)$$

where  $\mathbf{M}_d$  denotes the distal moment,  $\mathbf{I}_s$  is the tensor of inertia for the considered segment and  $\boldsymbol{\alpha}_s$  is its angular acceleration. The cross products account for moments resulting from the linear forces applied at the joints with  $\mathbf{r}_j$  being the vector from segment center of mass to the joint coordinates.

Based on these equations, the forces and moments are propagated along the kinematic chain, resulting in a residual force  $\mathbf{F}_{\text{res}}$  and a residual moment  $\mathbf{M}_{\text{res}}$  at the end of the chain, in our case, the center of mass of the upper body. If the model accelerations perfectly match the GRF/M these residuals are equal to zero, so that the inverse loss is defined as

$$L_{\text{inverse}} = \|\mathbf{F}_{\text{res}}\|_2^2 + \|\mathbf{M}_{\text{res}}\|_2^2. \quad (14)$$

The gradient calculation for the back propagation through the inverse layer is described in the supplementary material.

## 5 Experiments

In this section, the proposed methods are evaluated regarding their capability to learn exterior GRF/M and interior JT from motion. In particular, weakly-supervised learning on small training sets and domain transfer are investigated.

In the following experiments the root mean squared error (RMSE) and the relative RMSE (rRMSE), denoted by  $\epsilon$  and  $\epsilon_r$ , respectively, are used to quantitatively evaluate regression results. The rRMSE is normalized to average value ranges of the training set:

$$\epsilon_r = \frac{\epsilon}{\frac{1}{N} \sum_{i \in \text{train set}} [\max(v_t) - \min(v_t)]}. \quad (15)$$

Here,  $u$  is the predicted variable and  $v$  is the target variable. The number of training samples is denoted by  $N$ . All experiments are executed with three random splits into training, validation and test set. The splits are done subject wise, meaning, that sequences of the same subject are exclusively included in one set.

In addition to the experiments, presented in the following sections, an exemplary application to CMU data [6] and a noise experiment are included in the supplementary material.

### 5.1 Comparison to the state of the art

In a first experiment, we compare the proposed Dynamics Network and our baseline network to state-of-the-art methods for the inference of GRF/M together with JT. For this purpose, we use the gait set including slow and fast walking. Table 1 lists the corresponding results.

**Table 1.** RMSE  $\epsilon$  and rRMSE  $\epsilon_r$  of GRF/M and JT regression results of the gait data set.

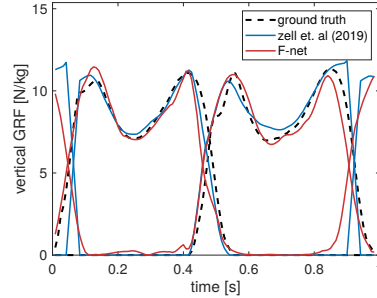
method	$\epsilon_f$ [N/kg]	$\epsilon_{r_f}$ [%]	$\epsilon_m$ [Nm/kg]	$\epsilon_{r_m}$ [%]	$\epsilon_\tau$ [Nm/kg]
Lv et al. [18]	0.700	20.3	0.077	28.1	-
Zell et al. [34]	<b>0.388</b>	<b>13.1</b>	<b>0.041</b>	<b>21.1</b>	0.055
Baseline net	0.591	14.4	0.056	21.2	0.055
F-net	0.626	14.9	0.059	22.1	<b>0.053</b>
cFI-net	0.733	14.7	0.064	22.5	0.054

While GRF/M ground truth can be calculated directly from force plate measurements and skeleton fits, the JT ground truth is a dynamics optimization result, associated with a higher uncertainty. Especially using forward dynamics, there are failure cases when the optimization does not converge to a satisfactory minimum. A comparison to these optimized sequences is not informative. Therefore, the JT evaluation is performed sample wise and only for the learning-based approaches to allow for a fair comparison. The GRF/M are evaluated as complete sequences.

The best performing method with regards to GRF/M regression [34], uses an SVM for gait phase classification and then regresses force and moment parameters on the resulting class subsets via Random Forests. In contrast to this approach the baseline and the Dynamics Network can be trained end-to-end

and yield competitive results. Regarding the JT results, the Dynamics Network slightly outperforms the other methods. Instead of minimizing a loss on the JT data directly, the forward loss regards the impact, the JT have on a simulated motion, which results in a stronger loss function. The method by Lv et al. [18] is a maximum a-posteriori inverse dynamics optimization, which incorporates a data-prior to guide the optimization towards realistic forces. The values were generated using our own implementation of the approach.

A qualitative comparison between F-net and [34], depicted in Figure 3, shows that although [34] achieves low error values, the primary gait phase classification may lead to rapidly changing forces at gait phase transitions. This physically implausible behaviour is inherent to the hierarchical structure of the method and does not occur using the proposed network.



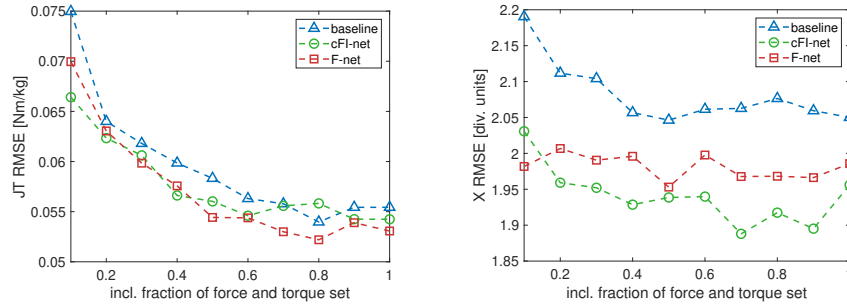
**Fig. 3.** Comparison of predicted vertical GRF by [34] and F-net

## 5.2 Learning dynamics on small data sets

The first goal of this work is to learn human locomotion dynamics on small data sets without overfitting to the training samples. In Section 4, we introduced two training modes with different degree of supervision to achieve this goal. These training modes are now compared to each other regarding their capability to operate on a training set with decreasing numbers of ground-reaction-set and torque-set samples. The motion-set, that only contains motion information, is always included to its full extend, so that the networks trained with forward and inverse layer, respectively, receive data of all subjects (contained in the training set). This way, we can show the effect of the higher data variability and the benefit of our dynamics layers. For cFI-training the contact-set is used, as well.

Figure 4 shows the RMSE values of predicted JT and their effect in terms of RMSE of simulated motion states. The error values are illustrated for different fractions of the used ground-reaction- and torque-set. A corresponding visualization of the mean regressed GRF/M and JT curves using only 10 % of the ground-reaction- and torque-sets can be seen in Figure 5. A percentage of 10 % corresponds to one subject included in the training data, whereas the test data contains motions of 5 subjects in each validation cycle.

It can be seen, that compared to the baseline, the Dynamics Network yields stable results with significantly decreased training sets, which indicates that the cyclic training is able to compensate the lack of ground truth data by learning from the larger motion-set. With increasing number of ground-truth GRF/M samples (cf. Figure 4 and Table 1), F-net performs slightly better than cFI-net. This is due to a gap between the model approximation and the measured



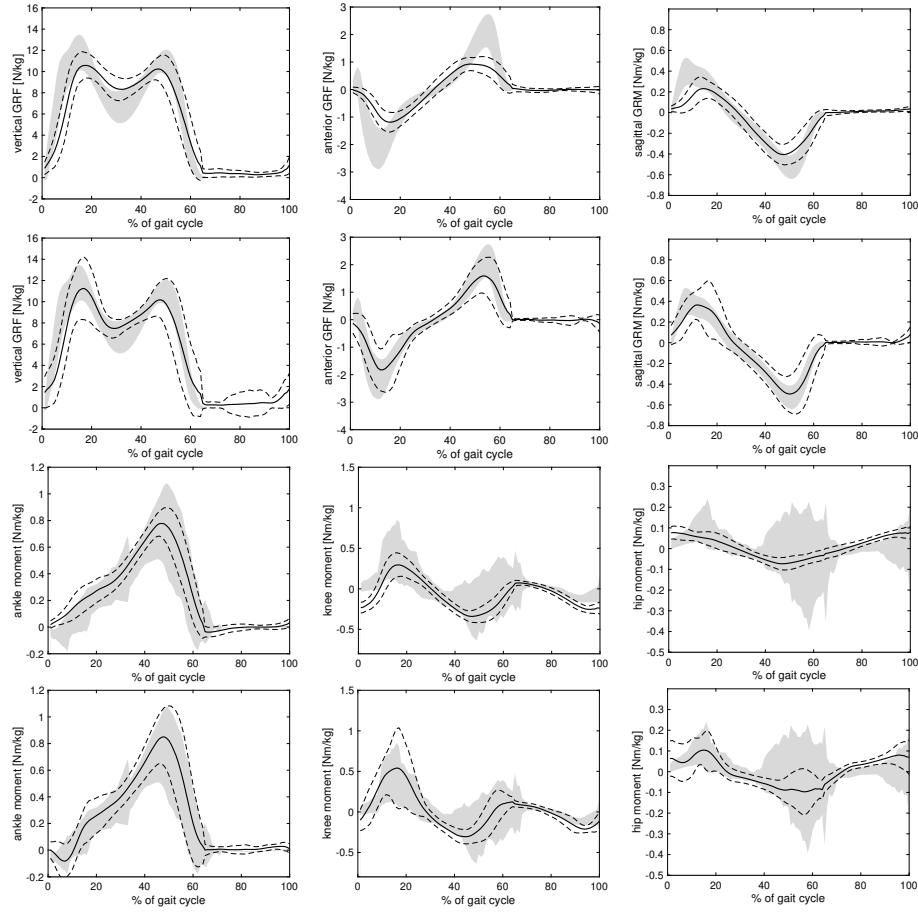
**Fig. 4.** Reduction of ground reaction and torque samples included in the training set. The left side shows RMSE of JT and the right side the associated RMSE of simulated motion states.

GRF/M. While F-net mainly learns ground truth GRF/M based on the included force-samples, cFI-net relies on the modelled forces to a similar extent, using the inverse-loss.

### 5.3 Domain transfer

The cyclic learning approach of the proposed method can be utilized for weakly- and unsupervised domain transfer. In the following experiment, we perform domain transfer between walking and running. For this purpose, we pre-train networks on the source motion type and apply the cyclic training modes introduced in Section 4 to achieve a transfer to the target motion type. During transfer learning no samples from the ground-reaction- and the torque-set are included, so that our methods operate completely independent from dynamics ground truth of the new motion type. This competence is only possible because of the novel dynamics layers. For cFI-training we use the motion-set in combination with the contact-set and minimize the corresponding loss functions alternately. During F-training, only the motion-set is used. For this reason, we refer to cFI-training as weakly-supervised and to F-training as unsupervised, although no ground truth data on GRF/M and JT are necessary in both cases. Table 2 lists the related results.

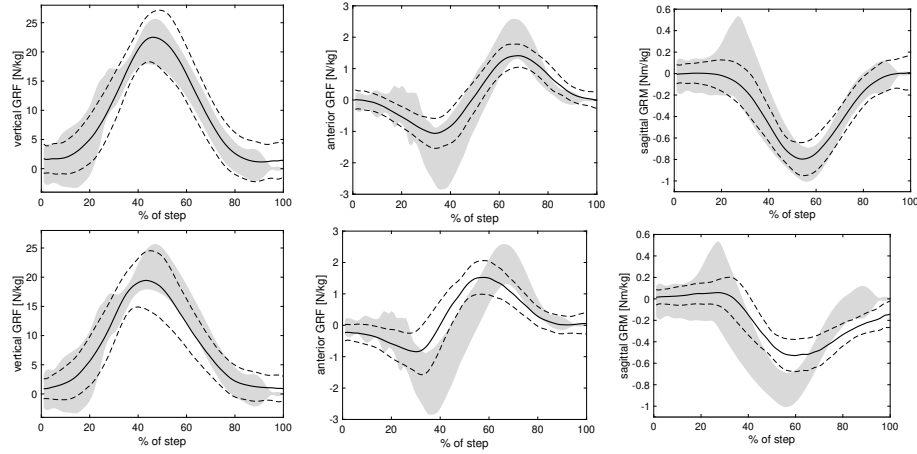
In this experiment, cFI-net clearly outperforms F-net, especially with respect to the GRF predictions. Most likely, the mutual dependency of GRF/M and JT is affecting the performance of F-net. The additional inverse layer, however, is able to solve this issue. In cFI-training the inverse layer is mainly responsible for the learning of realistic GRF/M, while the forward layer yields the matching JT. The contact-loss is necessary, since the minimization of the inverse-loss results in equally large gradients for the ground reaction of both feet. With the consideration of the contact state this overall equalisation is avoided. A visual comparison between the mean predicted GRF/M progressions for cFI-net and the supervised baseline network, is presented in Figure 6.



**Fig. 5.** Comparison of predicted mean curves using 10 % of the ground-reaction- and torque-sets. The first row shows baseline GRF/M results, the second row shows the corresponding predictions using F-net, the third row shows baseline JT results and the last row shows the corresponding results by F-net. The thick line is the mean value, the dashed lines represent standard deviations and the grey area displays the distribution of ground truth and optimized data, respectively.

**Table 2.** Domain transfer results for a transfer between walking and running in terms of RMSE  $\epsilon$  and rRMSE  $\epsilon_r$  for predicted GRF/M and JT.

transfer	method	$\epsilon_f$ [N/kg]	$\epsilon_{r_f}$ [%]	$\epsilon_m$ [Nm/kg]	$\epsilon_{r_m}$ [%]	$\epsilon_\tau$ [Nm/kg]
run	supervised	1.388	23.6	0.091	21.9	0.041
gait to run	F-net	3.942	35.6	0.178	26.1	0.062
gait to run	cFI-net	1.445	23.6	0.144	27.0	0.058
gait	supervised	0.591	14.4	0.056	21.2	0.055
run to gait	F-net	3.579	37.7	0.144	34.8	0.217
run to gait	cFI-net	0.685	13.5	0.076	33.2	0.153



**Fig. 6.** Domain transfer results for the target domain *running*. The top row shows baseline results, trained from scratch with supervision and the bottom row shows the results obtained using weakly-supervised transfer with cFI-net.

## 6 Conclusion

This paper proposes a weakly-supervised learning approach for the inference of human dynamics. Ground reaction forces, moments and joint torques are estimated from 3D motions by means of an artificial neural network (NN), that incorporates an inverse and a forward dynamics layer to allow the minimization of a pure motion loss. The method is designed for an optimal use of currently available data of human dynamics. While there exist large public motion capture data sets, the sets that include force plate measurements are often small in comparison. In our framework, the NN estimates joint torques, as well as exterior forces and moments from motion. Together with the forward dynamics layer, that executes an integration of the equations of motion, the network performs a full cycle from motion input to the underlying forces and back to a simulated motion. This way, a pure motion loss can be minimized, allowing for weakly-supervised learning and domain transfer. The experiments show, that the proposed method achieves state-of-the-art results and performs stable, even with few ground truth samples. The domain transfer between walking and running is realized without any ground truth on contact forces and moments. This experiment demonstrates the benefit of the inverse dynamics layer. Together with a loss on the binary contact state, it reliably constrains the exterior forces. The fact that the transfer is achieved with weak-supervision offers tremendous possibilities, especially in view of recent publications in pose reconstruction [1], [20], [27], [30] that make 3D human motion data readily available.

**Acknowledgement** Research supported by the European Research Council (ERC-2013-PoC). The authors would like to thank all subjects who participated in data acquisition.

## References

1. Arnab, A., Doersch, C., Zisserman, A.: Exploiting temporal context for 3d human pose estimation in the wild. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
2. de Avila Belbute-Peres, F., Smith, K., Allen, K., Tenenbaum, J., Kolter, J.Z.: End-to-end differentiable physics for learning and control. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 31, pp. 7178–7189. Curran Associates, Inc. (2018)
3. Bastien, G.J., Gosseye, T.P., Penta, M.: A robust machine learning enabled decomposition of shear ground reaction forces during the double contact phase of walking. *Gait Posture* **73**, 221 – 227 (2019)
4. Brubaker, M.A., Fleet, D.J.: The kneed walker for human pose tracking. In: *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*. pp. 1–8 (June 2008)
5. Choi, A., Lee, J.M., Mun, J.H.: Ground reaction forces predicted by using artificial neural network during asymmetric movements. *International Journal of Precision Engineering and Manufacturing* **14**(3), 475–483 (Mar 2013)
6. CMU: Human motion capture database (2014), <http://mocap.cs.cmu.edu/>
7. Degraeve, J., Hermans, M., Dambre, J., wyffels, F.: A differentiable physics engine for deep learning in robotics. *Frontiers in Neurorobotics* **13**, 6 (11 2019)
8. Devin, C., Gupta, A., Darrell, T., Abbeel, P., Levine, S.: Learning modular neural network policies for multi-task and multi-robot transfer. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 2169–2176 (May 2017)
9. Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A.: Temporal cycle-consistency learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
10. Federolf, P., Boyer, K., Andriacchi, T.: Application of principal component analysis in clinical gait research: Identification of systematic differences between healthy and medial knee-osteoarthritic gait. *Journal of Biomechanics* **46**(13), 2173 – 2178 (2013)
11. Finn, C., Xin Yu Tan, Yan Duan, Darrell, T., Levine, S., Abbeel, P.: Deep spatial autoencoders for visuomotor learning. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 512–519 (May 2016)
12. Fregly, B.J., Reinbolt, J.A., Rooney, K.L., Mitchell, K.H., Chmielewski, T.L.: Design of patient-specific gait modifications for knee osteoarthritis rehabilitation. *IEEE Transactions on Biomedical Engineering* **54**(9), 1687–1695 (Sept 2007)
13. Hoffman, J., Tzeng, E., Park, T., Zhu, J.Y., Isola, P., Saenko, K., Efros, A.A., Darrell, T.: Cycada: Cycle-consistent adversarial domain adaptation (2017)
14. Johnson, L., Ballard, D.H.: Efficient codes for inverse dynamics during walking. In: *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*. pp. 343–349. AAAI Press (2014)
15. Johnson, W., Alderson, J., Lloyd, D., Mian, A.: Predicting athlete ground reaction forces and moments from spatio-temporal driven cnn models. *IEEE Transactions on Biomedical Engineering* **PP**, 1–1 (07 2018)
16. Leporace, G., Batista, L.A., Metsavaht, L., Nadal, J.: Residual analysis of ground reaction forces simulation during gait using neural networks with different configurations. In: *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. pp. 2812–2815 (Aug 2015)

17. Leporace, G., Batista, L., Nadal, J.: Prediction of 3d ground reaction forces during gait based on accelerometer data. *Research on Biomedical Engineering* **34** (09 2018)
18. Lv, X., Chai, J., Xia, S.: Data-driven inverse dynamics for human motion. *ACM Trans. Graph.* **35**(6), 163:1–163:12 (nov 2016)
19. Maas, A.L.: Rectifier nonlinearities improve neural network acoustic models (2013)
20. von Marcard, T., Henschel, R., Black, M.J., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3d human pose in the wild using imus and a moving camera. In: *European Conference on Computer Vision* (Sep 2018)
21. Mukhopadhyay, R., Chaki, R., Sutradhar, A., Chattopadhyay, P.: Model learning for robotic manipulators using recurrent neural networks. In: *TENCON 2019 - 2019 IEEE Region 10 Conference (TENCON)*. pp. 2251–2256 (Oct 2019)
22. Oh, S.E., Choi, A., Mun, J.H.: Prediction of ground reaction forces during gait based on kinematics and a neural network model. *Journal of Biomechanics* **46**(14), 2372 – 2380 (2013)
23. Reinbolt, J.A., Fox, M.D., Schwartz, M.H., Delp, S.L.: Predicting outcomes of rectus femoris transfer surgery. *Gait & Posture* **30**(1), 100 – 105 (2009)
24. Schwab, A.L., Delhaes, G.M.J.: *Lecture Notes Multibody Dynamics B*, wb1413 (2009)
25. Seon Choi, H., Hee Lee, C., Shim, M., In Han, J., Su Baek, Y.: Design of an artificial neural network algorithm for a low-cost insole sensor to estimate the ground reaction force (grf) and calibrate the center of pressure (cop). *Sensors* **18**, 4349 (12 2018)
26. Shah, M., Chen, X., Rohrbach, M., Parikh, D.: Cycle-consistency for robust visual question answering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
27. Sharma, S., Varigonda, P.T., Bindal, P., Sharma, A., Jain, A.: Monocular 3d human pose estimation by generation and ordinal ranking. In: *The IEEE International Conference on Computer Vision (ICCV)* (October 2019)
28. Sim, T., Kwon, H., Eel Oh, S., Joo, S.B., Choi, A., Heo, H., Kim, K., Hwan Mun, J.: Predicting complete ground reaction forces and moments during gait with insole plantar pressure information using a wavelet neural network. *Journal of biomechanical engineering* **137** (06 2015)
29. Takahashi, K., Ogata, T., Nakanishi, J., Cheng, G., Sugano, S.: Dynamic motion learning for multi-dof flexible-joint robots using activepassive motor babbling through deep learning. *Advanced Robotics* **31**(18), 1002–1015 (2017)
30. Wandt, B., Rosenhahn, B.: Repnet: Weakly supervised training of an adversarial reprojection network for 3d human pose estimation. In: *Computer Vision and Pattern Recognition (CVPR)* (Jun 2019)
31. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
32. Winter, D.: *Biomechanics and Motor Control of Human Movement*. Wiley (2009)
33. Xiong, B., Zeng, N., Li, H., Yang, Y., Li, Y., Huang, M., Shi, W., Du, M., Zhang, Y.: Intelligent prediction of human lower extremity joint moment: An artificial neural network approach. *IEEE Access* **7**, 29973–29980 (2019)
34. Zell, P., Rosenhahn, B.: Learning inverse dynamics for human locomotion analysis. *Neural Computing and Applications* (Dec 2019)
35. Zell, P., Wandt, B., Rosenhahn, B.: *Physics-Based Models for Human Gait Analysis*, pp. 267–292. Springer International Publishing, Cham (2018)

36. Zhu, J., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 2242–2251 (2017)