

Hybrid Learning of Time-Series Inverse Dynamics Models for Locally Isotropic Robot Motion

Tolga-Can Çallar¹ and Sven Böttger¹

Abstract—Applications of force control and motion planning often rely on an inverse dynamics model to represent the high-dimensional dynamic behavior of robots during motion. The widespread occurrence of low-velocity, small-scale, locally isotropic motion (LIMO) typically complicates the identification of appropriate models due to the exaggeration of dynamic effects and sensory perturbation caused by complex friction and phenomena of hysteresis, e.g., pertaining to joint elasticity. We propose a hybrid model learning base architecture combining a rigid body dynamics model identified by parametric regression and time-series neural network architectures based on multilayer-perceptron, LSTM, and Transformer topologies. Further, we introduce a novel joint-wise rotational history encoding, reinforcing temporal information to effectively model dynamic hysteresis. The models are evaluated on a KUKA iiwa 14 during algorithmically generated locally isotropic movements. Together with the rotational encoding, the proposed architectures outperform state-of-the-art baselines by a magnitude of 10^3 yielding an RMSE of 0.14 Nm. Leveraging the hybrid structure and time-series encoding capabilities, our approach allows for accurate torque estimation, indicating its applicability in critically force-sensitive applications during motion sequences exceeding the capacity of conventional inverse dynamics models while retaining trainability in face of scarce data and explainability due to the employed physics model prior.

Index Terms—Dynamics, Model Learning for Control, Deep Learning Methods, Force and Tactile Sensing, Force Control

I. INTRODUCTION

THE availability of an inverse dynamics model (IDM) that infers the mechanical forces within a robotic system from its motion constitutes the foundation for a multitude of applications within the area of force control and motion planning. With the advent of multi-articulated robots employed in real-world tasks under uncontrolled conditions, there is an increased necessity for modeling strategies that are able to accurately and robustly represent the high dimensional dynamics of mechanically complex robots executing a kinematically and dynamically wide range of motions.

In real-world manipulation and physical interaction applications, the executed robot motion is often comprised of segments of low velocity and local isotropy in the form of frequent reversals of the joint-wise directions of rotation within small angular intervals for which we introduce the umbrella term of **Locally Isotropic Motion (LIMO)**. LIMO

Manuscript received: July, 29, 2022; Revised October, 12, 2022; Accepted November, 4, 2022.

This paper was recommended for publication by Editor Clément Gosselin upon evaluation of the Associate Editor and Reviewers' comments.

¹Tolga-Can Çallar and Sven Böttger are with the Institute for Robotics and Cognitive Systems, Universität zu Lübeck, Germany callar@rob.uni-luebeck.de

Digital Object Identifier (DOI): 10.1109/LRA.2022.3222951

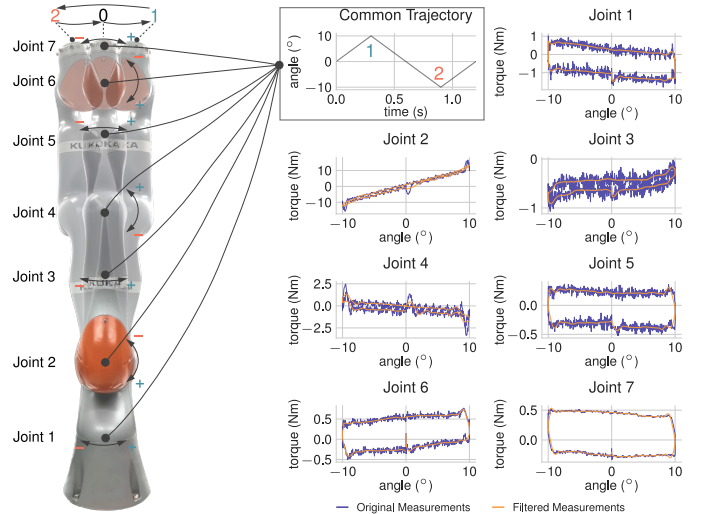


Fig. 1: Joint torque hysteresis on the KUKA LBR iiwa 14 during locally isotropic motion (LIMO) in the form of a joint-wise common cyclical linear trajectory. Also, note the presence of backlash torques at directional reversals.

can be encountered in collaborative, medical, social, or field robotics, where movements may predominantly manifest as irregular fine-scale motion sequences for precise compliant manipulation, e.g. in robotic sonography [1], and reactive contact control during continuous physical interactions with an uncontrolled environment and human agents [2]. This is distinct from locally anisotropic motion, e.g. in pick-and-place tasks [3] or large amplitude sinusoidal motion [4], where joint-wise rotation directions remain constant for larger angular intervals, i.e. saturating deformation states of flexible internal components and interfaces until, effectively, mechanical rigidity is approached up to a further directional reversal. During LIMO, an increased dynamic influence of hysteretic effects pertaining to low-velocity friction [5], joint flexibility, and backlash can be observed [2], [6]. As shown by the formation of characteristic joint torque hysteresis loops in Fig. 1, the sensed torques substantially depend on the rotation history, which is apparent from non-linearly increasing and saturating torques during rotation, and torque deflections at rotational reversals, which is consistent with mechanical hysteresis.

As such, LIMO quickly leads to the summation of multiple hysteretic torque deviations. In the context of IDMs, this introduces a non-negligible torque contribution that, albeit deterministic, results from the system's dynamic state history and is therefore difficult to model by conventional, steady-

state approaches [7]. Together with inherently complex non-linear dynamics of common robots, this aggravates the IDM problem for the deployment of robotics in critically force-sensitive applications where this motion type occurs.

IDM techniques are classically composed as parametric models [4] derived from the seminal work presented in [8], which formulates the Newton-Euler equations of motion linearly with respect to the inertial link parameters to identify a model by linear regression from dynamic measurements.

This so-called rigid body dynamics (**RBD**) model is often extended to also include basic frictional parameters [9].

For collaborative applications, lumped dynamic models in the form of momentum observers are proposed, disregarding dynamics related to acceleration trading faster model predictions and independence from noisy acceleration measurements for lower prediction accuracy [10]. Although providing good model explainability and robustness, parametric model formulations are inherently incomplete in face of a myriad of dynamic effects exceeding closed-form model definitions, and thus often unable to be encode temporal information and yield accurate estimates during challenging motion such as LIMO. On the other end of the spectrum, black-box techniques are investigated, ranging from non-parametric, Gaussian processes-based methods to artificial neural networks [11]–[14] employed to learn models directly from data, which in theory allows for the approximation of any non-linear dynamic system, although conversely introducing the risk of model variance and unexplainability. Thus, grey-box model learning approaches have been proposed, e.g. in form of neural network and Gaussian Process regression incorporating formulations of Newtonian, Lagrangian and Hamiltonian mechanics, where physical prior information is either directly embedded in a data-driven regression method, i.e. as a differentiable computational graph [11], [15]–[17] or used separately in an error-learning scheme [3], [11], [18]. However, time-dependency effects, innate to robot dynamics, have often been neglected. Advancements in this respect were made by casting the dynamic modeling problem as a time-series regression task and using recurrent neural networks for IDMs [7].

The commonality between these previous approaches is nonetheless that the investigated motion types are mainly intended for and validated on locally anisotropic motion, as defined above. To explore methods appropriate for the modeling of the dynamics of LIMO, we design dedicated techniques for LIMO modeling and present the following contributions:

- A hybrid learning base architecture combining a rigid body dynamics model prior learned by parametric identification with data-driven model learning using neural network topologies based on multilayer perceptrons and time-series networks in form of LSTM and Transformer topologies. (Sec. III)
- A novel joint rotation history encoding as a model input feature to reflect torque hysteresis states. (Sec. III-B1)
- A motion generation algorithm to obtain dynamically versatile training data for LIMO. (Sec. III-C)
- A comparative evaluation of various state-of-the-art IDM techniques on long-term proprioceptive time-series mea-

surements acquired on the KUKA LBR iiwa 14 during LIMO and conventional motion. (Sec. IV)

- Our model implementations and LIMO data sets are shared with the research community in [19].

The rationale behind our proposed methodology is based on two assumptions: Firstly, as evident from the mentioned related work on grey-box modeling the hybridization of a semantically pre-defined parametric model with a universal approximator should create synergetic effects regarding explainability, prediction robustness, and representation capabilities to eventually improve the modeling of complex dynamics occurring during challenging motion types as LIMO. Secondly, leveraging sequence history information on multiple levels, e.g. using different time-series neural network architectures and sequence encodings should benefit model accuracy, especially when dynamic states similar to LIMO are targeted.

Honing in on the outlined dynamic modeling problem, principal questions addressed by our work revolve around the infamous bias-variance tradeoff to design the optimal model architecture: **1)** Can rigid body dynamics models sufficiently represent the dynamics of LIMO? **2)** Can neural networks provide higher accuracy while generalizing to out-of-distribution dynamic states, given resource-intensive data acquisition in robotics? **3)** What is the benefit provided by a hybrid approach? **4)** How important is temporal information for model accuracy?

After delineating the fundamentals of the IDM problem, we present the details of our model architectures, feature design, data generation, and evaluation process.

II. ROBOT DYNAMICS

Considering a serial kinematic structure composed of $n = 7$ joints, in the case of our reference platform the KUKA iiwa 14, we define the dynamics by a composite model

$$\tau = \underbrace{H(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + g(\mathbf{q})}_{\tau_{RBD}} + \underbrace{\varepsilon(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \star)}_{\tau_{ADD}}, \quad (1)$$

expressing the joint torques $\tau \in \mathbb{R}^{n=7}$ as a function of the kinematic state variables of joint positions and their temporal derivatives, $\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}} \in \mathbb{R}^n$, along with an undefined set of additional variables subsumed under \star . With $H(\mathbf{q}) \in \mathbb{R}^{n \times n}$ as the symmetric and positive-definite joint-space inertia matrix, $C(\dot{\mathbf{q}}, \ddot{\mathbf{q}}) \in \mathbb{R}^{n \times n}$ as the combination of the Coriolis and centripetal forces and $g(\mathbf{q})$ as the gravity vector, the first terms grouped under τ_{RBD} represent the torques solely governed by the framework of rigid body dynamics [20]. Due to the presence of a plethora of effects that exceed those simplifying assumptions, e.g. manifestations of complex friction dynamics, e.g. *Stribeck* and *LuGre* friction [5], and joint flexibility, the model also includes a collective compensatory term $\varepsilon(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \star)$ to accommodate for joint torques τ_{ADD} resulting from the entirety of additional phenomena affecting the true robot dynamics. This also includes perturbations in the technical perception of the model variables caused by arbitrary changes of the actual or sensed dynamic state, e.g., deviations in the measurement of joint positions and computationally

derived velocities and accelerations, or the sensed joint torques themselves, i.e. mainly due to joint flexibility, friction or vibration [2].

We define the IDM problem as the identification of a dynamic model function $F(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \star)$ that approximates the observable true dynamics D such that the model estimates τ_{est} fulfill

$$\| \underbrace{D(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \star)}_{\tau} - \underbrace{F(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \star)}_{\tau_{\text{est}}} \| \rightarrow 0 \quad (2)$$

over a maximum range of the dynamic state space, where, in practical terms, \mathbf{q} and τ are observed via position sensors and strain-gauge sensors located at the output side of each joint of the KUKA LBR iiwa 14.

III. HYBRID MODEL LEARNING BASE ARCHITECTURE

Exploiting the compositionality of the model formulation in (1), we propose a hybrid model base architecture F_{HYB} consisting of two sequential elements, beginning with a parametric rigid body dynamics model F_{RBD} and a downstream neural network-based model F_{NN} as per Fig. 2 and:

$$\underbrace{F_{\text{HYB}}}_{\tau_{\text{HYB}}} = \underbrace{F_{\text{RBD}}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{k})}_{\tau_{\text{RBD}}} + \underbrace{F_{\text{NN}}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{r}, \tau_{\text{RBD}})}_{\tau_{\text{NN}}} \quad (3)$$

$F_{\text{HYB}} = \tau_{\text{HYB}}$ is the sum of the predictions made by both model components $F_{\text{RBD}} = \tau_{\text{RBD}}$ and $F_{\text{NN}} = \tau_{\text{NN}}$, where F_{NN} itself also depends on the prediction $F_{\text{RBD}} = \tau_{\text{RBD}}$ as an auxiliary input. As such, F_{HYB} constitutes a hybrid between a physics-based model prior and a data-driven, learned neural network. The internal structure of the respective components along with the definition of the utilized input features is the subject of the following.

A. Upstream Parametric Model Identification

Initially, we apply a parametric identification procedure based on the Newtonian formulation of mechanics to obtain a model F_{RBD} by statistical regression of the inertial link parameters implicitly embedded in the first terms of (1) pertaining to the torques τ_{RBD} that result from rigid body dynamics (RBD). For that, we extend the rotational component of the Newton-Euler equations to include viscous and Coulomb friction between adjacent links as a simplified actuation friction model [9], and define the torque τ_{ii} acting upon the i -th link due to its own motion as

$$\tau_{ii} = (\mathbf{g} - \ddot{\mathbf{p}}_i) \times m_i \mathbf{c}_i + \mathbf{I}_i \dot{\boldsymbol{\omega}}_i + \boldsymbol{\omega}_i \times (\mathbf{I}_i \boldsymbol{\omega}_i) + F_i^v \dot{\boldsymbol{\omega}}_i + F_i^c \text{sign}(\dot{\boldsymbol{\omega}}_i), \quad (4)$$




Fig. 2: Directed Graph of the Hybrid Base Architecture within which the dashed connections are ablated and the neural network configurations F_{NN} are varied as part of the evaluation experiments (see Sec. IV).




Fig. 3: (Left) Cartesian path of the end-effector during the execution of the optimized excitation trajectory. (Right) Excitation trajectory in configuration space.

where \mathbf{g} is the vector of gravitational acceleration, $\ddot{\mathbf{p}}_i$ is the linear link acceleration and $\boldsymbol{\omega}_i$ and $\dot{\boldsymbol{\omega}}_i$ are the angular link velocity and acceleration vectors. The parameters m_i , \mathbf{c}_i , \mathbf{I}_i , F_i^v and F_i^c , constitute the link mass, center of mass, symmetric inertia tensor, viscous and Coulomb friction coefficients respectively. We reformulate the Newton-Euler equations to express the joint torques linearly in relation to the link parameters as presented in [8], i.e.

$$\tau_{ii} = \mathbf{0} \begin{bmatrix} [(\mathbf{g} - \ddot{\mathbf{p}}_i) \times] & [\dot{\boldsymbol{\omega}}_i] + [\boldsymbol{\omega}_i \times] [\dot{\boldsymbol{\omega}}_i] & \dot{\boldsymbol{\omega}}_i & \text{sign}(\dot{\boldsymbol{\omega}}_i) \end{bmatrix} \times \underbrace{\begin{bmatrix} m_i & m_i \mathbf{c}_i & I_{xx} & I_{xy} & I_{xz} & I_{yy} & I_{yz} & I_{zz} & F_i^v & F_i^c \end{bmatrix}^T}_{\boldsymbol{\Phi}_i}. \quad (5)$$

With the Recursive Newton-Euler algorithm [20], given the kinematic structure \mathbf{k} , the link torques are transmitted backward along the kinematic chain to obtain a linear equation system

$$\tau_{\text{RBD}} = \mathbf{K}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{k}) \boldsymbol{\Phi}, \quad (6)$$

yielding τ_{RBD} for an observation matrix \mathbf{K} constructed from a measured kinematic state via the link parameters $\boldsymbol{\Phi} \in \mathbb{R}^{12n \times 1}$ [8]. Based on this, we construct a stacked version of (6) from N observations pairs of the kinematic state $(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}})$ and the joint torque observations τ_{meas} to estimate $\boldsymbol{\Phi}$ by the ordinary least squares method from

$$\bar{\tau}_{\text{meas}} = \bar{\mathbf{K}}(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}}, \mathbf{k}) \bar{\boldsymbol{\Phi}}, \quad (7)$$

where $\bar{\tau}_{\text{meas}} \in \mathbb{R}^{N \cdot n \times 1}$, $\bar{\mathbf{K}}(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}}, \mathbf{k}) \in \mathbb{R}^{N \cdot n \times 12n}$ and $\bar{\boldsymbol{\Phi}} \in \mathbb{R}^{N \cdot 12n \times 1}$. To optimize identifiability of $\boldsymbol{\Phi}$, the kinematic observations are made while executing joint-wise disjunct sinusoidal dynamic excitation trajectories that minimize $\text{cond}(\bar{\mathbf{K}}(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}}, \mathbf{k}))$ conversely maximizing parametric identifiability [4], [21]. The obtained excitation trajectory (Fig. 3) yields $\text{cond}(\bar{\mathbf{K}}) = 68$.

Before constructing (7), the measurements are low-pass filtered to reduce a) noise amplification from numeric differentiation of joint position measurements to retrieve the velocities and accelerations and b) torque vibrations caused by actuation transmission [2]. By rank reduction of $\bar{\mathbf{K}}(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}}, \mathbf{k})$ to $\bar{\mathbf{K}}_b$ with $\text{rank}(\bar{\mathbf{K}}_b) = 57$ via QR-decomposition [9], the reliably identifiable base parameters $\boldsymbol{\Phi}_b^{\text{LS}}$ are determined and solved for by pseudo-inversion of the full-ranked version of the linear equation system (7). This yields the model

$$F_{\text{RBD}} = \mathbf{K}_b(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}}, \mathbf{k}) \boldsymbol{\Phi}_b^{\text{LS}} = \tau_{\text{RBD}}, \quad (8)$$

which achieved a joint-wise average MSE of 0.2543 Nm on the data recorded during the excitation experiment.

B. Learning Downstream Models with Neural Networks

First, we introduce a motion encoding feature definition to alleviate the encoding of time-dependency information in the neural network input space. The subsequent sections III-B2 to III-B5 examine architectural variations of the neural network component F_{NN} utilized as the second stage of our hybrid model (3), for which we investigated and adapted several neural network types and topologies.

1) **Rotational History Encoding:** Our work focuses on dynamic sequence-history dependencies, i.e., hysteretic behavior, that occur during LIMO. From a topological perspective, learning of such sequence models can be accomplished either through the exploitation of recurrence (Sec. Section III-B3) or via an attention mechanism (Sec. Section III-B4). Still, the sequence length over which time-dependencies can be effectively modeled is limited. Conversely, generic feed-forward networks (Sec. Section III-B2) have no inherent means for sequence modeling.

Assuming that dynamic time-dependency effects mostly pertain to mechanical hysteresis during LIMO, we conjecture a correlation between observable joint torques and the history of joint positions that is defined by an associated hysteresis loop (Fig. 1). Hence, sufficient continuous rotation in one direction transitions the hysteresis into linear saturation, and contrarily, frequent reversals of motion direction, i.e., during LIMO, confine the associated dynamic effects to manifest with pronounced non-linearity inside the hysteresis loop. To abstractly express the hysteretic state of the joint torques depending on the history of past rotations, we introduce a novel rotational history encoding for which we define a feature vector

$$\mathbf{r} = [r_1 \quad r_2 \quad \dots \quad r_n], \quad r_j \in [-10^\circ, 10^\circ] \quad (9)$$

containing the signed total rotation angle of each joint since the respective last reversal of rotation direction. Before constructing \mathbf{r} for a time-step, the joint angles buffered since the last reversal are squashed into the range $\|\mathbf{r}^\circ\| < 10^\circ$ to exclude values redundantly encoding information for states where effects mechanical hysteresis is assumed to be saturated (Fig. 1). We hypothesize this auxiliary feature to aid sequence modeling networks with the extraction of long-term temporal dependencies within the data, as well as augment the performance of non-recurrent architectures by explicitly providing temporal information. The evaluation of the influence of this additional feature in the context of the generated data set for LIMO is presented in III-C.

\mathbf{r} is concatenated with the remaining input features of the kinematic state $(\mathbf{q}_{\text{meas}}, \dot{\mathbf{q}}_{\text{meas}}, \ddot{\mathbf{q}}_{\text{meas}})$ and the RBD model predictions $\boldsymbol{\tau}_{\text{RBD}}$ (3), yielding the full input feature space

$$\mathbf{x} = \text{vec}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{r}, \boldsymbol{\tau}_{\text{RBD}}) \in \mathbb{R}^{35} \quad (10)$$

for F_{NN} in (3) for which the network architecture variants are discussed next.

2) **Multilayer Perceptron (MLP):** Related to the topologies proposed by [3], [13], we construct a multilayer perceptron MLP_j for each robot joint j consisting of

- an input layer with 35 neurons accepting $\mathbf{x} \in \mathbb{R}^{35}$
- one hidden layer with 100 neurons with ReLU activation,
- an output layer with 1 neuron corresponding to the torque prediction $f_{\text{MLP}_j} = \tau_{\text{NN}}^j$ for the j -th joint.

The MLPs are arranged in a shared network topology such that the single-joint predictions are concatenated to $f_{\text{MLP}} = \boldsymbol{\tau}_{\text{NN}} = [f_{\text{MLP}_1}(\mathbf{x}) \quad \dots \quad f_{\text{MLP}_{n=7}}(\mathbf{x})]$ and the respective MLPs are trained jointly based on this output. For our implementation for the KUKA LBR iiwa with 7 joints, we abbreviate this architecture as **MLP-7**.

3) **Long Short-Term Memory (LSTM):** Consider the presence of temporal dependency effects extending over comparatively long time horizons in the dynamic system, that cannot be sufficiently encoded by temporal derivation of positional measurements or even time encodings. Hence, the model is extended from a single time-step prediction to the mapping of a time-series of kinematic states to a single dynamic state prediction using long short-term memory networks inherently suited for sequential data [7], i.e.,

$$f_{\text{LSTM}}(\mathbf{x}(t_1), \mathbf{x}(t_2), \dots, \mathbf{x}(T)) = \boldsymbol{\tau}_{\text{NN}}(T). \quad (11)$$

Here, T is the length of a sequence of observations $\mathbf{x}(t) = \text{vec}(\boldsymbol{\tau}_{\text{RBD}}(t), \mathbf{q}(t), \dot{\mathbf{q}}(t), \ddot{\mathbf{q}}(t), \mathbf{r}(t))$ for every time-step t_1, \dots, T , and $\boldsymbol{\tau}_{\text{NN}}(T) \in \mathbb{R}^7$ is the joint torque prediction made at the last time-step $t = T$ based on the preceding kinematic observation sequence. We compare two sequence lengths $T = 100$ and $T = 500$, to elucidate whether longer time-frames improve prediction accuracy, or put differently, whether dynamic effects with longer temporal dependency significantly influenced the model. The network is implemented as two consecutive LSTM cells with forget gates where

- the first LSTM cell has 35 input neurons accepting $\mathbf{x}(t) = \text{vec}(\boldsymbol{\tau}_{\text{RBD}}(t), \mathbf{q}(t), \dot{\mathbf{q}}(t), \ddot{\mathbf{q}}(t), \mathbf{r}(t))$, and 35 hidden neurons,
- the second LSTM cell has 35 input neurons and 7 hidden neurons, corresponding to the individual joint torques.

This network is abbreviated as **LSTM-2**.

Based on this same topology, we implement the architectures **LSTM-FCL 100** and **LSTM-FCL 500**, where a fully connected layer with 35 input and 7 output neurons with linear activation is used instead of the second cell to improve the regression performance by compensating for the squashing sigmoid output function of LSTM cells; the sequence lengths used for training are either $T = 100$ or $T = 500$.

4) **Transformer:** Originally proposed in [22], the Transformer network is investigated as a non-recurrent encoder-decoder topology, having achieved favorable results in sequence-to-sequence modeling applications. The input is concatenated by a sinusoidal encoding and fed through a series of encoding layers, each containing a multi-head scaled dot-product attention, normalization, and fully connected sub-layer. The encoded features are then fed through decoder layers and are subjected to a joint attention layer accepting both the encoding and the previous predictions to generate the next

output. The *vanilla* implementation of the architecture as presented in the seminal work is adopted and configured for input sequence and output dimensions as defined for the LSTM, yielding the network function $f_{TF}(\mathbf{x}(t_1), \mathbf{x}(t_2), \dots, \mathbf{x}(T)) = \tau_{NN}$.

5) **Input/Output Normalization:** Considering the largely varying numerical scales of the input features, all architectures are equipped with integral input and output normalization layers. Using the training data (see III-C), the mean \bar{x}_i and standard deviations $\sigma(x_j)$ are computed over all observations of the input features $x_i \in \mathbf{x}$ in (10) respectively to normalize each input x_j according to

$$x_i^{\text{norm}} = \frac{x_i - \bar{x}_i}{\sigma(x_i)}, \quad (12)$$

before being propagated forward through the network. The inverse operation is performed at the output where the prediction torques are subject to $\tau_j^{\text{pred}} = \sigma(\tau_j) \tau_j^{\text{norm}} + \bar{\tau}_j$.

C. Algorithmic Training Motion Generation with LIMOPA

Dynamic modeling is often performed for motion tasks within a fairly constrained effective workspace large-scale movement patterns [12], [14], e.g., palletization or pick-and-place tasks, characterized by mainly joint-wise anisotropic motion. To our best knowledge, the dynamic modeling of collaborative robots over a wide kinematic during LIMO and fine-scale movements was not attempted before. To provide a data set for the development and evaluation of IDMs for this motion type, we propose a novel algorithmic method for the parametric generation of randomized LIMO-paths (see Algorithm 1). The aims are 1) to provoke hysteretic effects at low velocities and directional reversals within narrow positional windows, i.e., backlash and non-linear friction, and 2) to maximize the kinematic range over which these are observed. LIMOPA partitions the motion in configuration space into linear *reaching*-phases where target configurations are approached, and subsequent sinusoidal *exploring*-phases where the robot performs locally isotropic reconfigurations around the prior target configuration. This is realized by the pre-calculation of a geometric path of length P composed as a joint configuration sequence $\mathbf{C} = \{\mathbf{q}_1, \dots, \mathbf{q}_P\}$ with $\mathbf{q}_p \in \mathbb{R}^{n=7}$




Fig. 4: Path geometry generated by LIMOPA. (*Left*) After interpolation between two scaffold configurations (see red markers) via random, joint-wise linear motion profiles, (*Right*) low-amplitude randomized sinusoidal joint trajectories are appended around the current scaffold configuration.

as the p -th waypoint of \mathbf{C} . The implementation is based on a randomly generated list of kinematically feasible and collision-free *scaffold*-configurations $\mathbf{S} = \{\mathbf{q}_1^{\text{scaf}}, \dots, \mathbf{q}_L^{\text{scaf}}\}$ of length L , that are identified in simulation. The scaffold configuration $\mathbf{q}_l^{\text{scaf}}$ represents the goal configuration to be reached in the l -th reaching-phase via joint-wise linear motion. This is followed by the l -th exploration-phase in the form of a time-unfolded joint-wise random sinusoidal path ${}_j e_l \in \mathbb{R}^T$ of length T parameterized by a discrete Fourier series

$${}_j e_l(t) = \sum_{k=1}^3 A_k^{\text{rand}} \sin(2\pi f_k^{\text{rand}} t + \phi_k^{\text{rand}}), \quad (13)$$

$${}_j e_l = \{{}_j e_l(1), \dots, {}_j e_l(T^{\text{rand}})\}. \quad (14)$$

$A_k^{\text{rand}} \in [0.5, 3]$, $f_k^{\text{rand}} \in [-4, 4]$ and $\phi_k^{\text{rand}} \in [0, 2\pi]$ are drawn randomly from the respective intervals for each joint j and each k -th summand of the Fourier series. The path length of the l -th exploration phase $e_l = ({}_1 e_l, \dots, {}_n e_l) \in \mathbb{R}^{(n=7) \times T^{\text{rand}}}$ is randomly determined by $T^{\text{rand}} \in [2000, 2500]$. This results in a configuration path

$$\mathbf{C} = \left\{ \underbrace{\mathbf{q}_1^{\text{scaf}}}_{7 \times 1}, \underbrace{e_1}_{7 \times T^{\text{rand}}}, \underbrace{\mathbf{q}_2^{\text{scaf}}}_{7 \times 1}, \underbrace{e_2}_{7 \times T^{\text{rand}}}, \dots, \underbrace{\mathbf{q}_L^{\text{scaf}}}_{7 \times 1}, \underbrace{e_L}_{7 \times T^{\text{rand}}} \right\}. \quad (15)$$

Joint-wise independent, random sinusoidal motions are chosen to excite a wide range of dynamic state combinations, also provoking hysteretic behavior via sinusoidal oscillations; the parameter intervals are defined to preserve nonlinearity by preventing saturation of potential hysteresis loops. Although neural network-based approaches rely on the availability of a rich and representative training data set, the high dimensionality of robot dynamics, encompassing at least $\dim(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \boldsymbol{\tau}) = 4 \times n$ state space dimensions for n joints, makes an exhaustive dynamic state space sampling intractable. To balance sampling efficiency and representative quality, we define an effective workspace W_{eff} as the kinematically reachable quadrant of a sphere radius 0.8 m in front of the robot, bounded by the y -axis and the positive x - and z -axes of the base coordinate system to be sampled as a subspace of the reachable workspace $W_{\text{tot}} \supset W_{\text{eff}}$ within which local manipulator dexterity is preserved [1], [23].

Algorithm 1 Locally Isotropic Motion Path Algorithm (LIMOPA)

```

while numberOfScaffoldConfigs < goalNumberOfScaffoldConfigs do
    randomConfig = getRandomConfig(jointIntervals)
    endEffectorPosition = forwardKinematics(randomConfig)
    feasible = checkFeasibility(randomConfig)
    if endEffectorPosition ∈ cartesianSampleRange and feasible == True then
        scaffoldConfigs[end] = randomConfig
    end if
end while
sortedScaffoldConfigs[1] = scaffoldConfigs[1]
for i=1 to length(sortedScaffoldConfigs) do
    queriedConfig = sortedScaffoldConfigs[i]
    remove queriedConfig from scaffoldConfigs
    neighborConfig[i+1] = kdNearestNeighbor(queriedConfig, scaffoldConfigs)
    remove neighborConfig from scaffoldConfigs
    sortedScaffoldConfig[i+1] = neighborConfig
    configPathPoints[end+1] = neighborConfig
    sinusoidals = generateRandomSinusoidals()
    append(configPathPoints, sinusoidals)
end for
return configPathPoints
    
```

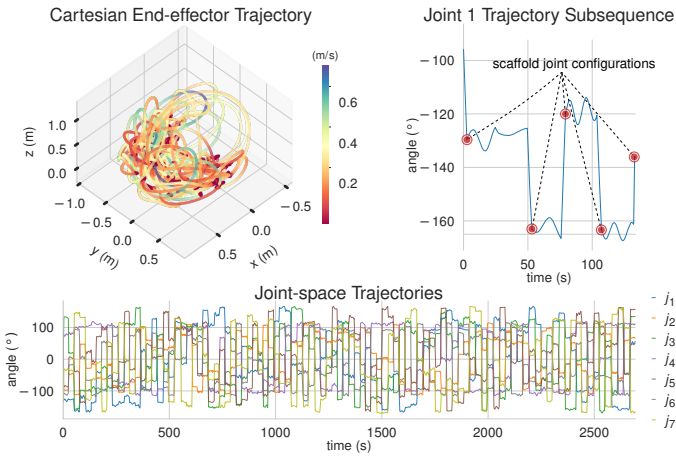


Fig. 5: (Top left) Trajectory traversed by the end-effector while executing the training motion path generated by LIMOPA (III-C). Note the red segments where joint-wise locally isotropic sinusoidal motions are executed. (Top right) A section of the sinusoidal trajectory executed on joint 1 shows locally isotropic motion in configuration space. (Bottom) The total training trajectory in configuration space.




Fig. 6: Visualization of the dynamic space of each joint spanning angle, speed, torque and rotation since direction reversal contained in the training data.

Given the robot joint limits, the joint spaces are uniformly sampled to obtain candidate configurations subject to a feasibility check for collisions and compliance with the sample space definition in simulation. This is repeated until, in our case, 150 so-called *scaffold* configurations are identified. To reduce the path execution time on the real robot, the scaffold configurations are sorted by distance in configuration space via a nearest-neighbor search based on a kd-tree representation of the configuration space. As described, the final geometric path is lastly constructed by inserting an exploration sequence after each scaffold configuration. With an average of 2250 waypoints for each exploration phase, this results in approximately, 337500 waypoints for the total path and an execution duration of roughly 1 hour. The waypoints on the geometric path are sequentially executed by the proprietary controller of the robot manufacturer such that the waypoint transitions are executed as linear configuration space interpolations with soft real-time constraints. At the beginning of each reaching phase, the joint-wise relative velocity and acceleration are randomly set between 60% and 90% with respect to the individual maximum joint velocities. During exploring phases, the relative velocities are randomly set between 10% and 30%.

Through this motion generation strategy, thorough coverage of the Cartesian and configuration workspace is achieved. Concurrently, the coverage of a wide range of dynamic state combinations in the acquired data is ensured by including multiple permutations of simultaneously occurring kinematic states, motion directions, and torques, yielding a comprehensive yet efficient sampling of the dynamic state space (Fig. 6 and 5). We provide our LIMOPA data set under [19].

D. Training Scheme

The neural networks are trained on the data set shown in Fig. 5 containing 804205 sequential joint positions and torques, where the velocities and accelerations are obtained by numeric differentiation.

We employ a random train-test-split of 80% to 20 % for the data set. With respect to the time-series networks, the data set is partitioned into sub-sequences, such that the shuffling affects only the order of the sub-sequences but not that of the samples within one such sub-sequence. The training of all networks is conducted for 10 runs with Xavier- and zero-initialization for 100 epochs using the AdamW optimizer [24] with an empirically determined initial learning rate $\epsilon = 0.001$ that is reduced when the training loss does not decrease for consecutive epochs. To prevent over-fitting, the batch size is set to 50. The mean squared error

$$L_{\text{MSE}}(\tau_{\text{HYB}}, \tilde{\tau}_{\text{meas}}) = \frac{1}{m} \sum_i^m \|(\tau_{\text{HYB}} - \tilde{\tau}_{\text{meas}})_i\|^2, \quad (16)$$

between τ_{HYB} and the low-pass filtered torque measurements $\tilde{\tau}_{\text{meas}}$, i.e., to reduce torque ripples due to vibrations induced by nonlinear actuator transmission [2], is used as the loss metric. The prediction modules are trained within the hybrid model structure (see (3)) either 1) with the full input feature space defined in (10), 2) as standalone networks without the input feature inclusion of τ_{RBD} and the output addition of the upstream RBD model estimates, 3) without the rotational history encoding input r , or in case of the time-series networks 4) with different input sequence lengths $T = \{100, 500\}$. As Transformers are intended for long sequence lengths exceeding the hypothetical capability of recurrent topologies, the Transformer is only trained once on sequence lengths of $T = 500$ and with the full feature space.

IV. EVALUATION OF THE MODEL ARCHITECTURES

To explore optimal IDM architecture and configuration for LIMO dynamics, we focus on the implications of various components of the base architecture presented in Section III.

The performance of the proposed model architectures and their variations are evaluated comparatively on the training and test data sets acquired during the execution of LIMO according to the trajectory generated by LIMOPA (Sec. III-C). The models are further validated on a separate data set acquired during the execution of a sequence of large-amplitude sinusoidal excitation trajectories (akin to the trajectory used for parametric regression) over the full workspace with decreasing base frequency to obtain dynamic and hysteretic states not contained in the original training data set, e.g. large accelerations, speeds, and rotations in one direction.




Fig. 7: Torque predictions of the evaluated model architectures during a LIMOPA-based trajectory for joints 2 and 7.

To this end, an inexhaustive grid search is performed by ablating parts of the input feature space and connections between the model components as per (10) and Fig. 2, and variation of the neural network base architecture. The tested modules are named after their base architecture (see Sec. III-B) and the specifiers "with r ", indicating usage of the rotational history encoding r , and "with τ_{RBD} ", denoting the model being configured as a hybrid architecture where τ_{RBD} is used both as a neural network input and for the final output addition (Fig. 3). Thus, the architectures are rendered equivalent to either the sole baseline RBD model, neural network component, or the full hybrid architecture.

On all data sets, models trained in hybrid configuration consistently achieved lower MSEs than their standalone counterparts, where hybridization of MLP-7 and LSTM-2 100 lead to MSE reductions by factors in the magnitude of 10^3 (see Table 1). This substantiates previous claims of a boosted absolute model performance through grey-box approaches, or at least better convergence on limited training data. In view of related hybrid modeling approaches carried out on dynamically similar robots as in [3], the problem of underfitting on highly dynamic robotic motion data seems common, necessitating the inclusion of model priors like an RBD model.

Interestingly, however, LSTM-FCL and Transformer archi-

tectures, converged to low MSEs on the LIMO data set even in non-hybrid configuration, whereby in hybrid configuration, LSTM-FCL achieved the lowest MSE. Specifically, the highest accuracy for the 500 variant, indicates that longer sequences may still contain additional exploitable information. The high accuracy of LSTM-FCL, may originate from its recurrent architecture and regression ability via the distal fully connected layer. The Transformer also evaluated on sequences of length $T = 500$ yielded no further improvement, which may indicate an advantage of recurrence over attention mechanisms.

A grid search for the optimization of topological hyperparameters for the best-performing model LSTM-FCL 500 with τ_{RBD}, r revealed both the addition of a second LSTM cell and the increase of the cell width to independently provide accuracy gains (Fig. 8). An exception was observed on the validation data where two layers with 15 neurons performed closely to two layers and 55 neurons and outperformed the other variants. As this is not seen in the LIMO data sets,

we suspect a lower dimensionality due to a width of 15 over two layers to cause a structural bias-variance balance allowing to rather generalize to non-LIMO dynamics contained in the validation data by restricting encoding of LIMO dynamics.

In terms of further utilization of sequence history information, each model trained with the rotational history encoding r outperformed their counterparts, demonstrating the benefit of sequence history information for non-recurrent and recurrent topologies alike. Respecting the performance on the non-LIMO validation data set, we report a reduced, albeit still significant, benefit of r . This can be regarded as evidence for our assumption, that sequence-history effects, i.e. hysteresis, play a more critical role during LIMO in contrast to conventional motion. This is also supported by a) the accuracy gains attributable to the use of time-series architectures on the LIMO data set, where these are able to leverage more temporal information, and, conversely, b) relative accuracy losses on the non-LIMO validation data set. Likewise, the RBD performed similarly on all data sets, but significantly worse than the remaining hybrid models during LIMO. This

TABLE I: MSE between model predictions and filtered measured torques. The results are sorted joint-wise by color.

Architecture	Joint 1	Joint 2	Joint 3	Joint 4	Joint 5	Joint 6	Joint 7	Joint Avg.
RBD	1.9499	25.2967	0.8434	2.7134	0.3373	0.0963	0.0442	4.4687
MLP-7	1.0649	202.5057	179.1369	77.0004	0.4084	0.5081	0.1949	63.8228
MLP-7 with τ_{RBD}	0.4154	0.9537	0.3475	0.3101	0.0468	0.1345	0.0701	0.3256
MLP-7 with τ_{RBD}, r	0.6039	82.8548	12.9735	12.1421	0.2604	0.4064	0.3088	15.6344
LSTM-2 100	0.0593	0.0952	0.0207	0.0208	0.0214	0.0252	0.0234	0.0380
LSTM-2 100 with τ_{RBD}	0.0518	0.0517	0.0163	0.0171	0.0129	0.0124	0.0111	0.0248
LSTM-2 100 with τ_{RBD}, r	0.1138	0.1530	0.0694	0.0662	0.0735	0.2152	0.1912	0.1260
LSTM-FCL 100	0.0770	0.2207	0.0586	0.0567	0.0270	0.0495	0.0413	0.0758
LSTM-FCL 100 with r	0.0385	0.0683	0.0188	0.0198	0.0161	0.0155	0.0115	0.0270
LSTM-FCL 100 with τ_{RBD}, r	0.0713	0.1468	0.0648	0.0698	0.0690	0.2011	0.1816	0.1150
LSTM-FCL 500	0.0317	0.1177	0.0237	0.0249	0.0284	0.0265	0.0300	0.0404
LSTM-FCL 500 with τ_{RBD}	0.0292	0.0507	0.0196	0.0199	0.0159	0.0133	0.0119	0.0230
LSTM-FCL 500 with τ_{RBD}, r	0.4293	0.1134	0.0459	0.0455	0.0190	0.0306	0.0216	0.0461
Transformer 500 with τ_{RBD}, r	2.0882	23.1839	0.0991	2.1749	0.4097	0.0985	0.0471	4.0145
RBD	1.9836	207.1595	178.9357	75.4605	0.4061	0.5070	0.196	66.2355
MLP-7	0.4081	0.9160	0.3758	0.3148	0.0463	0.1360	0.0710	0.3240
MLP-7 with τ_{RBD}	0.2747	0.8845	0.3709	0.2916	0.0252	0.0054	0.0385	0.2644
MLP-7 with τ_{RBD}, r	0.6095	80.8335	12.3105	11.5995	0.2737	0.4088	0.3088	15.7774
LSTM-2 100 with τ_{RBD}	0.0730	0.1297	0.0321	0.0338	0.0239	0.0291	0.0262	0.0497
LSTM-2 100 with τ_{RBD}, r	0.0663	0.0849	0.0274	0.0266	0.0151	0.0138	0.0121	0.0352
LSTM-FCL 100	0.1254	0.2114	0.0886	0.0922	0.0807	0.2282	0.1924	0.1455
LSTM-FCL 100 with r	0.1224	0.2955	0.1092	0.0777	0.0315	0.0595	0.0487	0.1064
LSTM-FCL 100 with τ_{RBD}, r	0.0526	0.0991	0.0258	0.0265	0.0171	0.0181	0.0124	0.0360
LSTM-FCL 500	0.0410	0.1380	0.0309	0.0319	0.0293	0.0273	0.0319	0.0472
LSTM-FCL 500 with τ_{RBD}	0.0400	0.0628	0.0253	0.0271	0.0168	0.0141	0.0124	0.0284
LSTM-FCL 500 with τ_{RBD}, r	0.1432	0.3179	0.1222	0.1284	0.0236	0.0390	0.0244	0.1142
RBD	3.487	31.0620	2.6918	4.7038	0.0639	0.0833	0.0177	5.5673
MLP-7	2.4092	234.5064	100.1022	120.5577	0.8988	0.6676	0.0627	66.6006
MLP-7 with τ_{RBD}	0.4618	10.3878	4.2824	4.0479	0.0644	0.1134	0.0387	2.7709
MLP-7 with τ_{RBD}, r	0.4289	7.4943	3.6548	3.8454	0.0386	0.1056	0.0207	2.2269
LSTM-2 100	0.5878	308.0968	370.908	88.5951	0.0740	0.2280	0.0325	0.1302
LSTM-2 100 with τ_{RBD}	0.5207	16.3673	4.0213	5.8706	0.0802	0.3689	0.0432	3.8960
LSTM-2 100 with τ_{RBD}, r	0.4342	15.1118	4.2252	5.2702	0.0743	0.2406	0.0627	3.6313
LSTM-FCL 100	1.9988	92.3633	43.3894	63.3062	0.8600	1.2037	0.1299	29.1216
LSTM-FCL 100 with r	2.4335	15.4618	4.6502	7.7476	0.2056	0.4692	0.2863	4.4649
LSTM-FCL 100 with τ_{RBD}, r	1.6668	12.8295	5.6193	7.7479	0.2303	0.1847	0.0616	4.0486
LSTM-FCL 500	2.1022	172.2180	52.5657	48.5196	0.9719	0.9582	0.1507	40.4978
LSTM-FCL 500 with τ_{RBD}	1.5170	18.0963	6.5467	5.7261	0.1450	0.2133	0.0810	4.6320
LSTM-FCL 500 with τ_{RBD}, r	1.7519	13.6668	6.2471	6.4873	0.1268	0.2433	0.0425	4.8088
Transformer 500 with τ_{RBD}, r	0.8202	20.3129	3.3142	3.6276	0.0561	0.1471	0.0448	4.0461




Fig. 8: MSEs of variations of LSTM-FCL-500 with τ_{RBD}, r investigated within ablation experiments.

again may be due to the proportionally higher influence of hysteretic effects being mainly provoked by LIMO dynamics outside the framework of RBD. Surprisingly, the hybrid MLP architectures performed best on the non-LIMO validation data in accordance with [3], [13] where this model class was successfully used for IDM of non-LIMO. We suspect predominant exposure to LIMO data during training leading the specialized time-series architectures to adapt to this motion type. Although re-training on data containing other dynamic ranges may be necessary to improve model generalization, the aforementioned topological optimization (Fig. 8) already yielded significant accuracy gains on the validation data set.

The temporal courses of the model predictions in Fig. 7 show the qualitative superiority obtained by model hybridism and time-series encoding capabilities by the congruence between predictions and measurements across all joints. In particular, this is apparent at inflection points in the torque profiles at directional reversals of joint 7, which is structurally isolated from gravitational torques, e.g. for the MLPs at $t \approx [2s, 12s]$: Without r the prediction curve is of rectangular shape with sudden jumps. This is mitigated by the rotation history encoding used by MLP-7 with r , which is able to approximate the slope of the torque curve. This qualitative advantage is present in the remaining time-series architectures as well, where the temporal information further reduces overshooting predictions at motion reversals and close tracking of the true torques, even for joint 2 with the highest dynamic range.

V. CONCLUSION

The inverse dynamics modeling problem is presented, emphasizing the complexity of modeling slow, locally isotropic, for which we postulated a relative increase in the contribution of a plethora of non-linear phenomena to the total dynamics originating from joint flexibility and low-velocity friction. Parametric model identification proved unsatisfactory in capturing the dynamics of motion sequences that provoke said effects. Likewise, conventional neural networks were shown to not yield satisfactory results in face of the data scarcity. Using time-series model learning based on LSTM and Transformer network topologies, we demonstrated that exploiting temporal information provides substantial model accuracy gains, that could be increased and transferred to non-recurrent architectures through the introduction of a rotational history encoding. We provided a hybrid base architecture that leverages the combination of a physics-based model prior with the universal approximation ability of neural networks, with three main advantages over conventional black or white box approaches for dynamic modeling: Firstly, the physical explainability of model predictions is retained. Secondly, the needed amount of training data could be reduced. Lastly, significant increases in joint torque estimation accuracy during complex and conventional motion were achieved. As a next step, we intend to investigate the integration of our model architecture within a real-time force control scheme.

REFERENCES

- [1] S. Böttger, T.-C. Callar, A. Schweikard, and E. Rückert, "Medical robotics simulation framework for application-specific optimal kinematics," *Current Directions in Biomedical Engineering*, vol. 5, no. 1, pp. 145–148, 2019.
- [2] V. Chawda and G. Niemeyer, "Toward torque control of a kuka lbr iiwa for physical human-robot interaction," in *IEEE International Conference on Intelligent Robots and Systems*, 2017, pp. 6387–6392.
- [3] K. Hitzler, F. Meier, S. Schaal, and T. Asfour, "Learning and adaptation of inverse dynamics models: A comparison," in *IEEE International Conference on Humanoid Robots (Humanoids)*, 2019, pp. 491–498.
- [4] Y. R. Stürz, L. M. Affolter, and R. S. Smith, "Parameter identification of the kuka lbr iiwa robot including constraints on physical feasibility," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6863–6868, 2017.
- [5] B. Bona and M. Indri, "Friction compensation in robotics: an overview," in *IEEE Conference on Decision and Control*. IEEE, 2005, pp. 4360–4367.
- [6] T. Allgeier and W. T. Evans, "Mechanical hysteresis in force transducers manufactured from precipitation-hardened stainless steel," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 209, no. 2, pp. 125–132, 1995.
- [7] E. Rueckert, M. Nakatenus, S. Tosatto, and J. Peters, "Learning inverse dynamics models in $O(n)$ time with lstm networks," in *IEEE International Conference on Humanoid Robotics (Humanoids)*, 2017, pp. 811–816.
- [8] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of inertial parameters of manipulator loads and links," *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 101–119, 1986.
- [9] W. Khalil and E. Dombre, "Dynamic modeling of serial robots," in *Modeling, Identification and Control of Robots*, W. Khalil and E. Dombre, Eds. Oxford: Butterworth-Heinemann, 2002, pp. 191–233.
- [10] S. Haddadin, A. De Luca, and A. Albu-Schäffer, "Robot collisions: A survey on detection, isolation, and identification," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.
- [11] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [12] S. Vijayakumar and S. Schaal, "Locally weighted projection regression: An $O(n)$ algorithm for incremental real time learning in high dimensional space," in *International conference on machine learning (ICML 2000)*, vol. 1, 2000, pp. 288–293.
- [13] N. Yilmaz, J. Y. Wu, P. Kazanzides, and U. Tumerdem, "Neural network based inverse dynamics identification and external force estimation on the da vinci research kit," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 1387–1393.
- [14] A. S. Polydoros and L. Nalpantidis, "A reservoir computing approach for learning forward dynamics of industrial manipulators," in *IEEE International Conference on Intelligent Robots and Systems*, 2016, pp. 612–618.
- [15] F. Díaz Ledezma and S. Haddadin, "First-order-principles-based constructive network topologies: An application to robot inverse dynamics," in *IEEE International Conference on Humanoid Robotics (Humanoids)*, 2017, pp. 438–445.
- [16] M. Lutter, C. Ritter, and J. Peters, "Deep lagrangian networks: Using physics as model prior for deep learning," *arXiv preprint arXiv:1907.04490*, 2019.
- [17] S. Greydanus, M. Dzamba, and J. Yosinski, "Hamiltonian neural networks," *Advances in neural information processing systems*, vol. 32, 2019.
- [18] R. Calandra, S. Ivaldi, M. Deisenroth, E. Rückert, and J. Peters, "Learning inverse dynamics models with contacts," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 3186–3191.
- [19] T.-C. Callar, "LIMODyn," 2022. [Online]. Available: <https://gitlab.roh.uni-luebeck.de/callar/limodyn>
- [20] R. Featherstone, "Inverse dynamics—the recursive newton-euler method," in *Robot Dynamics Algorithms*, 1987, pp. 65–77.
- [21] J. Swevers, C. Ganseman, D. Tukul, J. de Schutter, and H. Van Brussel, "Optimal robot excitation and identification," *IEEE Transactions on Robotics and Automation*, vol. 13, no. 5, pp. 730–740, 1997.
- [22] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017.
- [23] I. Kuhlmann, P. Jauer, F. Ernst, and A. Schweikard, "Robots with seven degrees of freedom: Is the additional dof worth it?" in *International Conference on Control, Automation and Robotics*, 2016, pp. 80–84.
- [24] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2019.