

# RoboMorph: In-Context Meta-Learning for Robot Dynamics Modeling

Manuel Bianchi Bazzi<sup>1</sup>, Asad Ali Shahid<sup>2</sup>, Christopher Agia<sup>3</sup>, John Alora<sup>3</sup>, Marco Forgione<sup>2</sup>,  
Dario Piga<sup>2</sup>, Francesco Braghin<sup>1</sup>, Marco Pavone<sup>3</sup> and Loris Roveda<sup>2,3</sup>

<sup>1</sup>*Department of Mechanical Engineering, Politecnico di Milano, Italy*

<sup>2</sup>*Istituto Dalle Molle di Studi Sull'Intelligenza Artificiale (IDSIA USI-SUPSI),  
Scuola Universitaria Professionale della Svizzera Italiana, DTI, Italy*

<sup>3</sup>*Stanford University, U.S.A.*

**Keywords:** Transformers, In-Context Learning, Meta Learning, Transfer Learning, Deep Learning, Isaac Gym, Robot Dynamics Modeling.

**Abstract:** The landscape of Deep Learning has experienced a major shift with the pervasive adoption of Transformer-based architectures, particularly in Natural Language Processing (NLP). Novel avenues for physical applications, such as solving Partial Differential Equations and Image Vision, have been explored. However, in challenging domains like robotics, where high non-linearity poses significant challenges, Transformer-based applications are scarce. While Transformers have been used to provide robots with knowledge about high-level tasks, few efforts have been made to perform system identification. This paper proposes a novel methodology to learn a meta-dynamical model of a high-dimensional physical system, such as the Franka robotic arm, using a Transformer-based architecture without prior knowledge of the system's physical parameters. The objective is to predict quantities of interest (end-effector pose and joint positions) given the torque signals for each joint. This prediction can be useful as a component for Deep Model Predictive Control frameworks in robotics. The meta-model establishes the correlation between torques and positions and predicts the output for the complete trajectory. This work provides empirical evidence of the efficacy of the in-context learning paradigm, suggesting future improvements in learning the dynamics of robotic systems without explicit knowledge of physical parameters. Code, videos, and supplementary materials can be found at project website.

## 1 INTRODUCTION

The field of Deep Learning has undergone a significant transformation with the widespread adoption of Transformer-based architectures (Vaswani et al., 2017), particularly impacting Natural Language Processing (NLP) in generating text (Touvron et al., 2023). Recently, new applications have emerged to solve partial differential equations and image vision tasks. However, in complex areas such as robotics, (Shahid et al., 2022), where systems are highly non-linear, the implementation of Transformer-based solutions remains limited. Transformers have shown success in providing high-level task knowledge to robots, but there has been little progress in using them for learning system dynamics or performing system identification. This paper introduces a new approach to learning a meta-dynamical model for high-dimensional physical systems, such as the Franka robot arm, using a Transformer-based architecture

without prior knowledge of the system's physical parameters. The goal is to accurately predict key quantities (end-effector pose and joint positions) based on the torque signals for each joint. Such predictions are valuable for integration into Deep Model Predictive Control frameworks, which are increasingly utilized in robotics. The meta-model is given an initial context, that establishes the relation between torques and positions and predicts the output for the complete trajectory. Using massively parallel simulations, large datasets, representing different robot dynamics, are generated in a simulated physics environment (Isaac Gym) to train the meta-model. The effectiveness of this learned model is demonstrated across various types of control inputs. This work demonstrates the use of transformer-based models in learning robot dynamics without any explicit knowledge of physical parameters.

## 1.1 Related Work

In the domain of Partial Differential Equations (PDEs) resolution, Transformers are gaining momentum compared to Physics Informed Neural Networks (PINNs) and Recurrent structures such as LSTMs (Schmidhuber et al., 1997). Transformer-based architectures, coupled with Fourier-Neural-Operators (Li et al., 2020), demonstrate promising capabilities (Yang et al., 2023). Despite the limited existing literature on meta-learning in robotics, (Gupta et al., 2022) demonstrates the adaptability of Transformers in learning general controllers. Furthermore, Transformers and meta-learning have notably found applications in the domain of Fault Diagnosis (Chen et al., 2023). (Lin et al., 2022) discusses considerations regarding Transformers and highlights a key challenge: their inefficiency in processing long sequences, primarily due to the computational and memory complexities associated with the self-attention module.

## 1.2 Contribution

This paper proposes an approach that utilizes transformer-based architectures to learn a meta-dynamical model of a robot arm. Diverse training datasets are generated using massively parallel simulations in the Isaac Gym simulation framework (Makovychuk et al., 2021). To generate trajectories for robots in simulation, specific families of Operational Space Control (OSC) tasks were selected, with domain randomization applied to certain parameters. After learning the meta-model, this work investigates the following aspects:

1. **Hyper-Parameters and Context:** Evaluation of how hyper-parameters, such as the number of multi-attention heads, dimension of compressed information ( $d_{\text{model}}$ ), and context length influence prediction accuracy.
2. **In-Distribution and Out-of-Distribution Performance:** Examining the meta-model's ability to generalize in an out-of-distribution regime, and the required training dataset dimension to achieve acceptable results.
3. **Transfer Learning on Out-of-Distribution Scenarios:** Analysis of the impact of pre-training on the transfer of knowledge across different distributions of control actions, including zero-shot and few-shot learning scenarios.

## 2 METHODS

This paper aims to develop an approach that relies on a black-box model-free simulation of a system without any information regarding the robotic system.

### 2.1 Preliminaries

**Black Box Model:** Black box models describe the input-output behavior of a system without explicitly modeling its internal mechanisms. These models are estimated directly from experimental data.

**Simulation Model:** Simulation models predict system outputs solely from inputs, without relying on past measured data. They are valuable for designing controllers and emulating physical systems.

**Model-Free Simulation:** In this approach no specific model is needed for the system at hand; the meta-model  $\mathcal{M}_\phi$  receives input/output sequences  $(u_{1:m}^{(i)}, y_{1:m}^{(i)})$  up to time step  $m$  (*context*) and a test input sequence (*query*)  $u_{m+1:N}^{(i)}$  from time  $m+1$  to  $N$  to produce the corresponding output sequence  $\hat{y}_{m+1:N}^{(i)}$ :

$$\hat{y}_{m+1:N}^{(i)} = \mathcal{M}_\phi(u_{1:m-1}^{(i)}, y_{1:m-1}^{(i)}, u_{m+1:N}^{(i)}). \quad (1)$$

The meta-model indeed is trained by minimizing the mean squared error (MSE) between the simulated open-loop prediction and the ground truth.

### 2.2 Model Architecture

The standard Transformer architecture has been adapted to handle real-valued input/output sequences generated by dynamical systems, instead of the sequences of symbols (word tokens) typically used in natural language modeling. Specifically, compared to plain GPT-2 (Radford et al., 2019), the initial token embedding layer is replaced by a linear layer with  $n_u + n_y$  inputs and  $d_{\text{model}}$  outputs, while the final layer is replaced by a linear layer with  $d_{\text{model}}$  inputs and  $n_y$  outputs. This modification allows the Transformer to effectively process continuous-valued sequences, typical of dynamical systems, aligning the model architecture with the requirements of system identification tasks rather than natural language processing.

The architecture is visualized in Figure 1 and consists of (i) an encoder that processes  $u_{1:m}, y_{1:m}$  (without causality restriction) and generates an embedding sequence  $\zeta_{1:m}$ ; (ii) a decoder that processes  $\zeta_{1:m}$  and test input  $u_{m+1:N}$  (the latter with causality restriction) to produce the sequence of predictions  $\hat{y}_{m+1:N}$ .

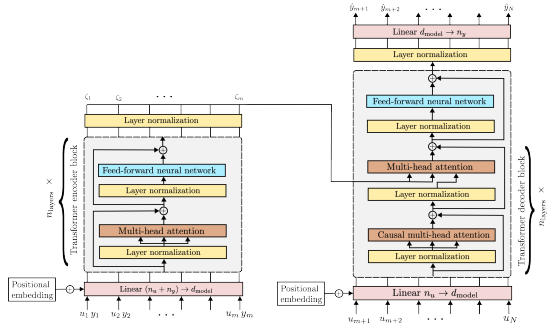


Figure 1: Encoder-decoder Architecture (Forgione et al., 2023).

## 3 EXPERIMENTAL FRAMEWORK

### 3.1 Overview

The strength of the proposed approach lies in its flexibility, allowing us to arbitrarily choose the dimensions of inputs and outputs. In the analyzed experimental framework, we learn a mapping from a 7-dimensional input to a 14-dimensional output:

- **Input:** joint torques applied to the 7 degrees of freedom (DoFs) Franka robot in  $Nm$ .
- **Output:**
  - end effector (EE) Cartesian coordinates  $(x, y, z)$  in [m];
  - EE quaternion  $(X, Y, Z, W)$ ;
  - joint positions  $(q_0 \dots q_6)$  in [rad].
- **Context:** 20% of the entire simulation.
- **Total Time Window:** 1000 steps  $\rightarrow$  16.7 s.

The core idea is to train a meta-model by varying the range of parameters influencing the dynamics of the Franka robot. Figure 2 shows the context and the input used to generate the predictions for complete trajectory.

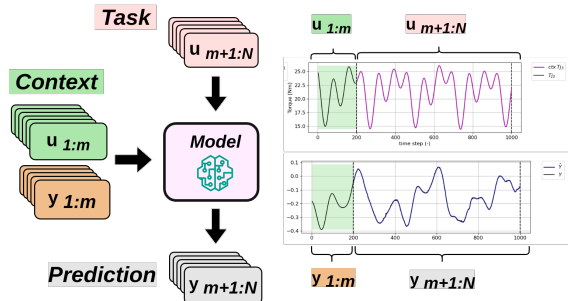


Figure 2: Proposed meta-model uses context and input to perform the prediction. Context length is highlighted in green color.

### 3.2 Domain Randomization

The classic domain randomization approach simulates a broad distribution across domains or Markov Decision Processes (MDPs), aiming to train a model robust enough to perform well in real-world conditions (Hospedales et al., 2020). Our proposed approach focuses on modeling a real Franka robot arm using this consideration: Handling real-world scenarios involves addressing deviations from nominal values, such as uncertainties in joint friction and damping characteristics. Below, we highlight some foundational aspects of the proposed domain randomization:

- Each link’s mass is randomized according to a uniform distribution. This strategy increases the complexity of the problem, prevents overfitting in the model, and enhances overall robustness.
- Initial joint positions are randomized to facilitate meta-learning of the system. This variation aims to enhance the explorability of the robot’s workspace, rather than strictly mapping specific control actions or trajectories of the robotic arm.
- Unlike typical robotic controllers in Franka Control Interface (FCI), which compensates for gravity terms and internal joint frictions by default, our control actions consider those components. We virtually handle the torque from the motor side, emphasizing that the model needs to gain a deeper understanding of the underlying physical system.

#### 3.2.1 Object of Randomization

**Initial Joint Positions:** Initial joint positions are randomized around the mid-positions.

**Mass of the Links:** Each link’s mass is uniformly distributed around its *nominal value*, with a symmetric variation of a certain percentage ( $\pm x\%$ ).

**Joint Stiffnesses:** Stiffness and damping of joints are randomized and constrained to fixed values.

**Center of Masses:** The center of mass is varied along three dimensions.

**Frequency (for Random Inputs):** This parameter serves as the main frequency component of the overall signal. Its value ranges from 0.1 and 0.25 Hz.

### 3.3 Task Definition

Tasks are categorized into two main families: synthetic inputs or trajectories derived from Operational Space Control (OSC), spanning from direct torque commands to desired Cartesian positions. Figure 3 shows the example cartesian trajectories generated with two types of control inputs.

### 3.3.1 Synthetic Random Input

In this case, the torque is generated directly by a specific mathematical function and applied in  $Nm$ . Two types of torque inputs, namely *multi-sinusoidal* and *chirp*, are considered to generate torque profiles for each joint:

- *Multi-sinusoidal*:

$$u_i = [A_0 \quad A_1 \quad A_2 \quad A_3]_i \begin{bmatrix} \cos(w_0 t) \\ \sin(w_1 t) \\ \cos(w_2 t) \\ \sin(w_3 t) \end{bmatrix}, \quad (2)$$

$$\begin{bmatrix} w_0 \\ w_1 \\ w_2 \\ w_3 \end{bmatrix} = \begin{bmatrix} w_0 \\ 1.5 \cdot w_0 \\ 2 \cdot w_0 \\ 3 \cdot w_0 \end{bmatrix}, \quad (3)$$

$$w_0 = 2\pi f_0, \quad f_0 \in \left[ \frac{f_m}{1.5}, 1.5 \cdot f_m \right], \quad (4)$$

$$A_0 \cdots A_3 \in [-f_m \cdot 15, f_m \cdot 15]. \quad (5)$$

- *Chirp*:

$$u_i = A_i \cos(w_1 \cdot (1 + \frac{1}{4} \cdot \cos(w_2 \cdot t)) \cdot t + \phi), \quad (6)$$

$$\phi \in [-\pi, \pi], \quad q_0 \in [-0.5, 0.5], \quad A_i \in [-4, 4], \quad (7)$$

$$w_1 = 2\pi f_1, \quad w_2 = 2\pi f_2, \quad (8)$$

$$f_1 \in [f_m, f_m \cdot 1.5], \quad f_2 \in \left[ \frac{f_m}{1.5}, f_m \cdot 2 \right]. \quad (9)$$

Figure 4 illustrates random control actions for two joints generated with multi sinusoidal, with each color representing a different robot. It is important to note that joint 1 experiences higher gravity compensation compared to joint 0. As a result, the randomness appears less pronounced but it also varies considerably across the workspace.

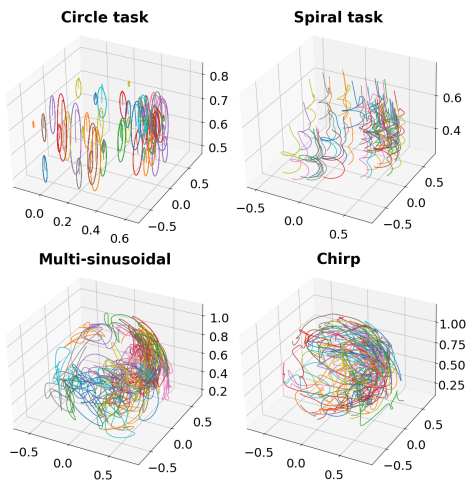


Figure 3: 3D trajectories in Cartesian coordinates [m].

### 3.3.2 Operational Space Control (OSC)

OSC tasks depend on two control gains,  $\mathbf{K}_p$  and  $\mathbf{K}_d$  which regulate the responsiveness of the control action in relation to the error with respect to the desired trajectory  $x^d$ . The control input  $u$  is defined as:

$$u = g(q) + J_{ee}^T(q) \left( M_{ee} \cdot \ddot{x}^d + C_{ee} \cdot \dot{x}^d + K_p \cdot (x^d - x) + K_d \cdot (\dot{x}^d - \dot{x}) \right).$$

Two operational space tasks are defined as illustrated in Figure 3:

- *Circle Task*: circle trajectories on YZ plane, with different radii and *same* frequencies.
- *Spiral Task*: spiral trajectories on the XY plane performed top-down or down-top, at different frequencies and with different radii.

### 3.4 Data Generation in Isaac Gym

Isaac Gym (Makoviychuk et al., 2021) is NVIDIA's prototype physics simulation environment for RL research. It allows developers to experiment with end-to-end GPU-accelerated RL for physical systems. To the best of the author's knowledge, there is no previous contribution mentioning Isaac Gym as a simulation environment outside of the RL domain, indeed the use of Isaac Gym in this work to generate large datasets for transformer-based model training is quite novel. Ensuring a trade-off between randomness and feasibility regarding workspace position and required torque is crucial to prevent issues during the training stage. For these reasons, three features/sub-functions have been incorporated to handle spurious robot instances:

- **Self-Collision and Floor-Collision Detection:** To ensure self-collision and floor-collision scenarios are managed effectively, a net contact force

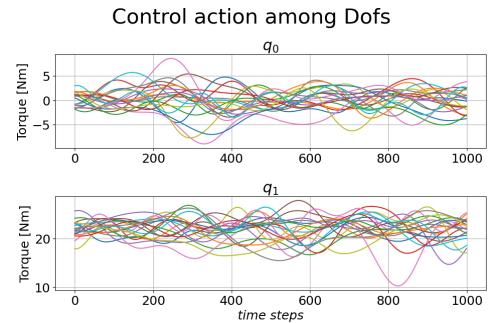


Figure 4: Multi-sinusoidal torque profiles for joint 0 and joint 1, respectively, and for 20 robots each.




Figure 5: Collision detection visualization.

tensor is employed to filter out colliding robots from the data buffer. Figure 5 shows the detection of such collisions.

- **Position and Torque Saturation Check:** In this scenario training data should ideally consist of numerical functions that are free from singularities or unrealistically high values. At each time step, simulations that result in joint positions reaching their limits or torques which saturate, are excluded from the training dataset.
- **Exclusion of Quaternion Error:** In certain simulations, unexpected changes in quaternion acquisition have been observed, possibly due to internal reference shifts within Isaac Gym or singularities. These sudden changes occurred infrequently and were straightforwardly excluded from the dataset.

The data generation algorithm is shown in Algorithm 1.

```

Data: num_robots, timesteps
Result: Generated data, black-listed robots
Setting up simulation parameters;
Generating random torques for all robots' DoFs;
for  $i \leftarrow 1$  to num_robots do
    | Randomization of the dynamical parameters;
end
while  $t \leq$  timesteps do
    | Apply torques and step simulation;
    if saturation or collision or self_collision
        then
            | add to the black-list of robot_idx;
        else
            | Store full pose and torques in buffer;
        end
    end
    | Remove black-listed robots;
    | Save tensors;
    
```

Algorithm 1: Data Generation - Higher Level.

### 3.5 Dataset Composition

Figure 6 illustrates the composition of datasets used for training and fine-tuning based on the type of control action and the number of simulations. Each

dataset has specific mass and joint randomization bounds.

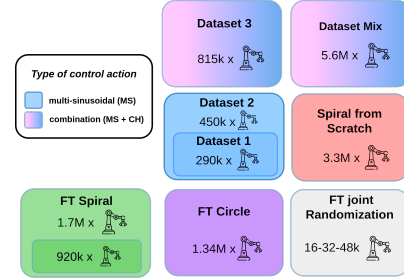


Figure 6: Dataset composition.

### 3.6 Training Loss

As highlighted by (Kirsch et al., 2024), Transformer training often exhibits a substantial plateau problem. One of the solutions proposed by the authors is to increase the batch size to reduce the plateau problem. In this work, each “simulation batch” is composed of over 3000 classes of Franka robots (each with different dynamics), while each training batch consisted of 16 robots. This approach aimed to promote learning-to-learn as opposed to mere memorization. The loss function for an example training is shown in Figure 7.

## 4 EVALUATION APPROACH

Each model trained on a specific dataset is tested on different simulations, which may be In or Out of Distribution with respect to the training dataset. Different metrics have been used to compare these prediction performances.

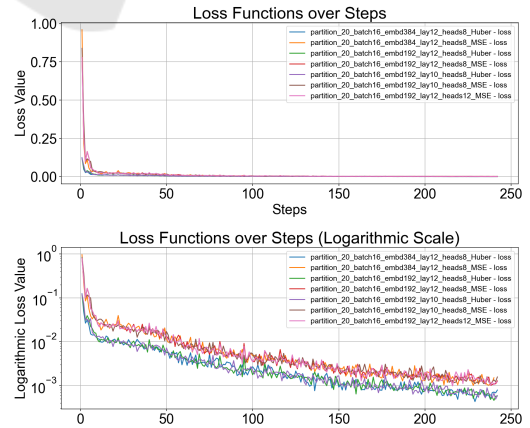


Figure 7: Loss function of an example dataset (for every 100 iterations).

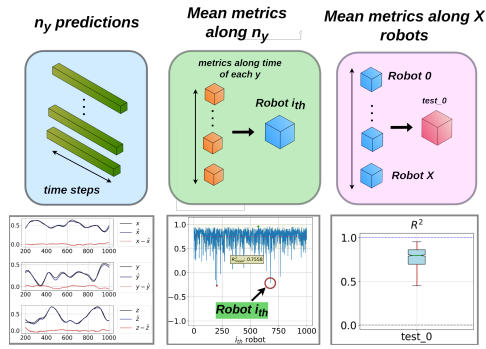


Figure 8: First evaluation approach.

## 4.1 Metrics

Indexes are calculated over time for each coordinate, then averaged across the output dimensions  $n_y$ , and finally averaged across a batch of  $X$  robots. This type of evaluation is illustrated in Figure 8. The following indexes are used:

- Coefficient of Determination:  $R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$ ;
- Root Mean Square Error:  $RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}$ ;
- Normalized Root Mean Square Error:  $NRMSE = \frac{RMSE}{\sigma_y}$ ;
- Fit Index:  $FI = 100 \cdot \left( 1 - \frac{\sqrt{\|\sum_{i=1}^n (y(t) - \hat{y}(t))\|^2}}{\sqrt{\sum_{i=1}^n (y(t) - \bar{y}(t))^2}} \right)$ .

The main drawback of this approach is that  $R^2$  for “flat” and relatively small values tends to penalize the single robot prediction metrics. In the OSC tasks, most of the joints experience small variations compared to multi-sinusoidal and chirp signals and for this reason, averaging over the output dimensions  $n_y$  and then over batches leads to really low  $R^2$  mean values. For this reason another *evaluation approach* was used to investigate the performance of OSC tasks. Instead of averaging over output dimensions  $n_y$  and then over different robots, several trajectories are merged sequentially coordinate by coordinate, and then indexes are computed on these merged trajectories as shown in Figure 9. Predictions are performed separately but they’re considered merged only for the calculation of metrics. The aim is to better demonstrate the overall capability in different scenarios rather than evaluating the local precision as in the first evaluation approach.

$R^2$  tends to assume values very close to 1.0 on the coordinates which show already good capabilities locally. To better evaluate the prediction performance




Figure 9: Second evaluation approach.

of the models, the fit index has been chosen as the primary metric, due to its better differentiation even when  $R^2$  values are particularly high and similar.

## 5 RESULTS

**Influence of the training context:** Tests were performed for contexts ranging from 5% to 50%, with 20% used as a reference for all tests. The quadratic dependence of computational and memory complexity on sequence length in self-attention is a limiting factor for Transformers. Generally, Transformers can accept context lengths different from those used during training, provided that the test context lengths are smaller than the training ones. Accordingly, it is also possible to use the same context length while varying the prediction horizon to a smaller number of steps, effectively changing the context length in percentage.

Varying the prediction horizon while keeping the context window the same as in training reveals an initial transition error. However, overall performance improves as the test horizon approaches the training horizon as shown in Figures 10 and 11, indicating that the model requires the knowledge of whole trajectory patterns to perform optimally.

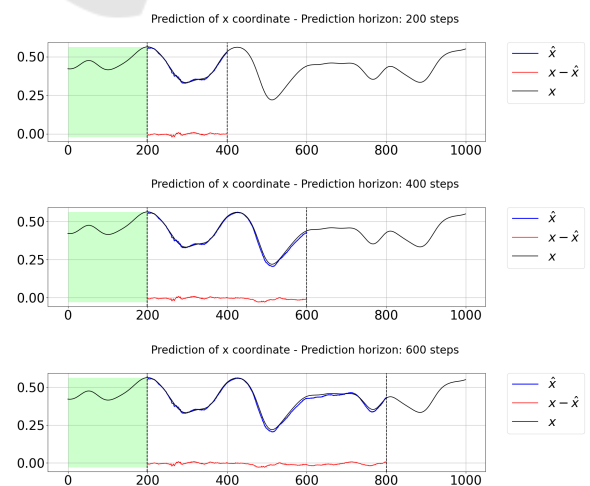



Figure 10: Different prediction horizons and with the same test context as in training.

Figure 11:  $R^2$  in a 200 steps context.

**Hyper-Parameters Tests:** The hyper-parameter values tuned through trial and error are given below. Most other parameters were fixed as per (Forgione et al., 2023). Batch size is not mentioned as an independent parameter, due to computational constraints.

- **Loss Function:** MSE / Huber;
- **Model Dimension** ( $d_{\text{model}}$ ): 192 / 384;
- **Number of Multi-Attention Heads:** 8 / 12;
- **Number Layers:** 10 / 12 / 16.

Given the training and test dataset from the same distribution, it is possible to investigate how these parameters affect the quality of predictions. Table 1 shows an example demonstrating the impact of the number of layers on performance, using the first evaluation approach. The results highlight the best performance achieved with the 12 layers model. Varying the model dimension and number of multi-attention heads showed comparable results without any significant performance differences. The results are available in the appendix on project website.

**Analysis of Model Performance:** The performance analysis of the optimal model is presented in Table 2 (with related color-convention in Table 3), with further details provided below:

- **In Distribution:** For bounded tasks, good results can be achieved on in-distribution tests within 2 hours. However, as the number of tasks and frequencies increases, the required training time and

Table 1: Impact of the number of layers on prediction accuracies.

10 Layers				
TestA	$R^2$	$\sigma_{R^2}$	RMSE	$\sigma_{RMSE}$
test0	0.682	0.343	0.0530	0.0236
test1	0.750	0.292	0.0465	0.0218
test2	0.755	0.293	0.0458	0.0212
12 Layers				
TestA	$R^2$	$\sigma_{R^2}$	RMSE	$\sigma_{RMSE}$
test0	0.713	0.287	0.0521	0.0242
test1	0.787	0.219	0.0454	0.0219
test2	0.793	0.212	0.0448	0.0213
16 Layers				
TestA	$R^2$	$\sigma_{R^2}$	RMSE	$\sigma_{RMSE}$
test0	0.700	0.303	0.0531	0.0258
test1	0.761	0.262	0.0468	0.0239
test2	0.767	0.255	0.0462	0.0234

Table 2: Comparison between zero-shot, fine-tuning on the spiral, and training from scratch.

	Zero-Shot		FT Spiral (2.5h)		Scratch (9h)	
	$R^2$	Fit Index	$R^2$	Fit Index	$R^2$	Fit Index
x	0.98	86	0.998	96.09	0.988	88.92
y	0.99	89.78	1	97.89	0.996	93.47
z	-3.014	-100.34	0.867	63.57	-1.254	-50.13
X	0.986	88.01	0.999	97.56	0.982	86.54
Y	0.989	89.51	1	98.04	0.998	95.72
Z	0.137	7.11	0.953	78.37	0.493	28.76
W	0.191	10.05	0.965	81.39	0.572	34.54
q0	0.999	96.77	1	98.71	0.998	95.15
q1	0.966	81.56	0.983	86.8	0.875	64.67
q2	0.847	60.94	0.997	94.55	0.991	90.37
q3	-0.669	-29.18	0.969	82.44	0.627	38.9
q4	0.988	88.96	1	97.87	0.996	93.44
q5	0.871	64.13	0.999	96.77	0.838	59.81
q6	0.981	86.11	1	98.05	0.997	94.28

Table 3: Fit index color convention.

Fit Index Range
$\geq 90$
80 - 89.99
60 - 79.99
30 - 59.99
< 30

model dimensions naturally increase.

- **Slightly out of Distribution:** The model showed good results, supporting the hypothesis of generalization rather than mere memorization of the training dataset. Specifically regarding mass randomization, its impact appears to be secondary to the frequency of the tested signal. Transformers demonstrate greater effectiveness when applied to higher frequencies compared to their training frequencies, as opposed to lower frequencies.
- **Out of Distribution:** The first column of Table 2 displays the performance of the model, which was trained on a mixed dataset containing both multi-sinusoidal and chirp signals, and tested on a spiral task. The performance varies significantly across different output dimensions, with some showing good results and others performing poorly.
- **Fine-Tuning on Variable Spiral:** Subsequent examinations involved a model trained extensively for 16 hours on a range of different frequencies, including both multi-sinusoidal and chirp signals. This model was then fine-tuned on a Spiral task for 2.5 hours (following the ideas in (Piga et al., 2024)). Results were compared to a model trained from scratch for 9 hours (Table 2, third column). The results show that fine-tuned model perform significantly better than the one trained from scratch.

**Additional Analysis:** Even after fine-tuning the Spiral task, challenges persist in the z-coordinate, likely due to the intrinsic complexity of the problem. Spirals, varying in both upward and downward directions, exhibit distinct characteristics in terms of fre-

quencies and elevation rates: combinations of larger radii and slow elevations differ significantly from tasks involving fast elevation changes and small radii.

Fine-tuning on the same task family (multi-sinusoidal) with controlled joint randomization led to a significant increase in overall prediction accuracy, depending on the number of examples provided.

Fine-tuning on Spiral tasks and zero-shot testing on the Circle task resulted in poor performance. This underscores the difficulty in predicting OSC tasks generally, particularly during initial transient phases. Direct fine-tuning on the Circle task yielded better results compared to fine-tuning on Spiral tasks and testing on Spiral. This observation suggests that the model's effectiveness is constrained to in-distribution (ID) and slightly out-of-distribution (OOD) tasks.

## 6 CONCLUSIONS

This work tackles the problem of learning a meta-model of robot dynamics using an encoder-decoder Transformer architecture. The challenge lies in the simulation domain, where the meta-model accurately predicts complex systems over long sequences based on a 20% context and 80% prediction of the overall trajectory. The results indicate that Transformer-based models can learn dynamics in a zero-shot or few-shot fashion within control action distributions, suggesting their potential for use in robotics. The results also highlight fine-tuning is advantageous in these scenarios and more practical compared to training such models from scratch.

Current limitations are mainly related to generalization. Following a black-box approach to generalize a single robotic arm independently of the type of control action appears structurally unfeasible. Results show distinctions between in-distribution (ID) and out-of-distribution (OOD) control actions, highlighting the critical role of the model's inputs.

For future work, the results presented pave the way for pre-compensating control actions in unknown systems, particularly where estimating parameters such as payload, joint stiffness, and damping is challenging. This approach can be extended beyond control distributions to encompass Transfer Learning across diverse robot morphologies.

## ACKNOWLEDGMENTS

This paper has received funding from the Hasler Foundation under the GENERAI (GENerative Robotics AI) Project.

## REFERENCES

- Chen, C., Wang, T., Liu, C., Liu, Y., and Cheng, L. (2023). Lightweight convolutional transformers enhanced meta-learning for compound fault diagnosis of industrial robot. *IEEE Transactions on Instrumentation and Measurement*, 72:1–12.
- Forgione, M., Pura, F., and Piga, D. (2023). From system models to class models: An in-context learning paradigm. *IEEE Control Systems Letters*, 7:3513–3518.
- Gupta, A., Fan, L., Ganguli, S., and Fei-Fei, L. (2022). Metamorph: Learning universal controllers with transformers. *arXiv preprint arXiv:2203.11931*.
- Hospedales, T., Antoniou, A., Micaelli, P., and Storkey, A. (2020). Meta-learning in neural networks: A survey.
- Kirsch, L., Harrison, J., Sohl-Dickstein, J., and Metz, L. (2024). General-purpose in-context learning by meta-learning transformers. *arXiv preprint arXiv:2212.04458*.
- Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A. (2020). Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.
- Lin, T., Wang, Y., Liu, X., and Qiu, X. (2022). A survey of transformers. *AI Open*.
- Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., and State, G. (2021). Isaac gym: High performance gpu-based physics simulation for robot learning. *CoRR*, abs/2108.10470.
- Piga, D., Pura, F., and Forgione, M. (2024). On the adaptation of in-context learners for system identification. In *Proceedings of the 20th IFAC Symposium on System Identification*, Boston, MA.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. (2019). Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Schmidhuber, J., Hochreiter, S., et al. (1997). Long short-term memory. *Neural Comput*, 9(8):1735–1780.
- Shahid, A. A., Piga, D., Braghin, F., and Roveda, L. (2022). Continuous control actions learning and adaptation for robotic manipulation through reinforcement learning. *Autonomous Robots*, 46(3):483–498.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. (2023). Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*, 30.
- Yang, L., Liu, S., Meng, T., and Osher, S. J. (2023). In-context operator learning with data prompts for differential equation problems. *Proceedings of the National Academy of Sciences*, 120(39):e2310142120.