
Task-Linear Deep Representation of Physical Systems

Matthieu Blanke^{1,2} Marc Lelarge^{1,2}

Abstract

Machine learning methods can be a valuable aid in the scientific process, but they need to face challenging settings where data come from inhomogeneous experimental conditions. Recent meta-learning methods have made significant progress in multi-task learning, but they rely on black-box neural networks and suffer from a lack of interpretability. We introduce Task-Linear Deep Representation, or TLDR, a new meta-learning architecture capable of learning efficiently from multiple environments by incorporating the linear structure observed in many problems. Unlike other approaches, we prove that TLDR is able to learn the physical parameters of the system, hence enhancing interpretability. We show that our method performs competitively by comparing it to state-of-the-art algorithms on two systems derived from scientific modeling.

1. Introduction

The learning of physical systems is an essential application of artificial intelligence that can unlock significant technological and social progress. Physical systems are inherently complex, making them difficult to learn (Karniadakis et al., 2021). One particularly challenging and common scenario is multi-task learning, where observations of physical systems are collected under inhomogeneous experimental conditions (Caruana, 1997). In such cases, the scarcity of training data necessitates the development of robust learning algorithms that can efficiently handle environmental changes and make use of all available data.

Related work Multi-task learning has recently been addressed from a perspective of meta-learning (Hospedales et al., 2021). Building on the power of deep learning,

^{*}Equal contribution ¹INRIA ²DI ENS, PSL Research University. Correspondence to: Matthieu Blanke <matthieu.blanke@inria.fr>.

Accepted after peer-review at the 1st workshop on Synergy of Scientific and Machine Learning Modeling, SynS & ML ICML, Honolulu, Hawaii, USA. July, 2023. Copyright 2023 by the author(s).

gradient-based meta-learning methods such as the well-known MAML algorithm (Finn et al., 2017) and its ANIL variant (Raghu et al., 2020) achieve remarkable performances. Lately, there has been a growing interest in adapting these methods to architectures specifically designed for physical systems. Recent work has focused on the generalization and prediction of dynamical systems, such as DyAD (Wang et al., 2022), LEADS (Yin et al., 2021) and CoDA (Kirchmeyer et al., 2022).

These approaches model data with deep neural networks and hence benefit from high expressivity. However, deep neural networks also suffer from a lack of interpretability, which can be a problem when dealing with physical systems (Karniadakis et al., 2021). Typically, neural network parameters have no physical meaning, and the mathematical structure that many problems present is often not exploited. A common strong property that is common in physics is the linearity of equations with respect to certain parameters. Building on this observation, we introduce a learning architecture incorporating this structure, called Task-linear Deep Representation (TLDR).

Contributions We propose a formalization of gradient-based meta-learning algorithms and use it to compare the learning architectures of the main existing models. We introduce TLDR, a meta-learning architecture designed to leverage the linear structure of the system across learning environments. Our method learns and generalizes efficiently, and with low computational cost. We prove that it is able to identify the true system parameters with minimal supervision, unlike other approaches. We show experimentally that TLDR outperforms other state-of-the-art meta-learning algorithms on two physical systems. Furthermore, we show theoretically and experimentally that TLDR is able to recover the physical parameters of the system, bringing interpretability to the learning algorithm.

2. Gradient-based meta-learning

In this section, we describe the generic structure of gradient-based meta-learning algorithms. The learner is provided with a meta-dataset $D := \bigcup_{t=1}^T D_t$ consisting of T datasets from different tasks (or environments). We will assume for simplicity a classical supervised regression setting

Algorithm 1 Gradient-based meta-training

input datasets D_1, \dots, D_T , parametric-model $f(x; \theta, w)$,
 initial meta-parameters π , learning rate η , parametrizer ϕ
output learned meta-parameters $\bar{\pi}$
while not converged **do**
 for tasks $1 \leq t \leq T$ **do**
 instantiate parameters $(\theta_t, w_t) = \phi(\pi, D_t)$
 instantiate model $f_t(x; \pi) = f(x; \theta_t, w_t)$
 compute the task-specific loss $L_t[f_t(x; \pi)]$ (2.1)
 end for
 compute the meta-loss $L(\pi) = \frac{1}{T} \sum_{t=1}^T L_t[f_t(x; \pi)]$
 update $\pi \leftarrow \pi - \eta \nabla L(x; \pi)$
end while

where $D_t := \{x_{i,t}, y_{i,t}\}_{1 \leq i \leq N_t}$. The goal is to learn a predictor $x \mapsto y$ from D that is robust to task changes, in the sense that when presented with data from a new task, it can learn the underlying function from a few shots (Hospedales et al., 2021).

Architecture In gradient-based meta-learning a parametric model such as a neural network $f(x; \theta, w)$ is used for learning and adaptation to the environments. The parametrization consists in a task-agnostic component θ and a task-specific component w . In this framework, a meta-learning algorithm is characterized by a task parameter map $\phi : (\pi, D_t) \mapsto (\theta, w)$ specifying how the task-specific parameters of f are instantiated from the meta-parameters π and the dataset of task t . We give examples of this formalism for recent architectures in Table 1.

Training At training time, for each task t , the meta-learner instantiates a task-specific version of the model from the task dataset D_t , defining $f_t(x; \pi) = f(x; \phi(\pi, D_t))$. The error on dataset D_t is measured by the task-specific loss

$$L_t[f] := \sum_{i=1}^{N_t} \frac{1}{2} \|f(x_{i,t}) - y_{i,t}\|^2. \quad (2.1)$$

The parameters are trained to minimize the meta-loss defined as the aggregation of the L_t :

$$L(\pi) := \frac{1}{T} \sum_{t=1}^T L_t[f_t(x; \pi)]. \quad (2.2)$$

The training process is summarized in Algorithm 1.

Adaptation At test time, the learner is presented a new dataset D_{T+1} consisting of few samples (or shots). Using this adaptation data and the learned meta-parameters, the learner instantiates $f_{T+1}(x; \bar{\theta}, \bar{w})$ with $(\bar{\theta}, \bar{w}) = \phi(\bar{\pi}, D_{T+1})$. The task-agnostic component $\bar{\theta}$ is frozen, and

the task-specific component is tuned from the initial value \bar{w} by minimizing the prediction error on the adaptation dataset:

$$\min_w L_{T+1}(\bar{\theta}, w; D_{T+1}). \quad (2.3)$$

In all the above approaches, this minimization is performed by gradient descent. The meta-learning algorithm is successful if the adapted model has good prediction performance on new test samples of task $T + 1$.

3. Task-linear meta-learning

Many physical systems can be modeled mathematically as a sum of various functions representing different physical contributions, with scalar coefficients that may vary across experiments. We postulate that a task is represented by a vector of system parameter, or weight $w \in \mathbb{R}^r$, in which the target function is linear

$$f_\star(x; w) := w^\top v_\star(x), \quad (3.1)$$

with $v_\star(x) \in \mathbb{R}^r$. Hence the observations of the dataset have the form $y_{i,t} = f_\star(x_{i,t}; w_\star^*)$.

Although this structural assumption may not be valid for *all* physical systems, it does allow to model a broad class of problems as the following examples illustrate.

Example 1 (Electrostatic field). An electromagnetic field is typically the sum of several contributions (*e.g.* an ambient field and an object-specific field), which are neither known nor controlled by the experimenter and hence vary from an experimenter to another. However, their mathematical expressions across tasks are known to have a common structure, up to some linear coefficients. For instance, we consider the potential of an electrostatic dipole moment in the presence of an ambient field:

$$f_\star(x; w) = \frac{px_1}{2\pi\epsilon_0(x_1^2 + x_2^2)^{3/2}} - Ux_1, \quad (3.2)$$

This physical system has the form of (3.1) with the system parameters $w = (U, p)$ and features $v(x) = (x_1, x_1/(2\pi(x_1^2 + x_2^2)^{3/2}))$. The corresponding electric field is given by $E(x) = -\nabla f_\star(x) \in \mathbb{R}^2$ (Haus & Melcher, 1989).

Example 2 (Robot inverse dynamics). The Euler-Lagrange formulation for the dynamics of a robot has the form

$$M(q)\ddot{q} + C(q, \dot{q})\dot{q} + g(q) = Bu \in \mathbb{R}^d, \quad (3.3)$$

where q are the generalized coordinates, M is the mass matrix, C captures the Coriolis forces, $g(q)$ is the gravity vector and the matrix B maps the input u into generalized forces. It can be shown that equation (3.3) is linear with respect to the physical parameters (Nguyen-Tuong & Peters, 2010),

Table 1. Structure several meta-learning models. Here v and u denote parametric models, such as neural networks.

Meta-algorithm	MAML	LEADS	CoDA	TLDR
structure of the parametrization	additive in parameter space	additive in function space	additive in parameter space	multiplicative in function space
meta-parameters π	θ	$(\theta, w_1, \dots, w_T)$	$(\theta, \Theta, \xi_1, \dots, \xi_T)$	$(\theta, w_1, \dots, w_T)$
$\phi(\pi, D_t)$	$(\theta, -\eta \nabla L(\theta; D_t))$	(θ, w_t)	$(\theta, \Theta \xi_t)$	(θ, w_t)
$f(x; \theta, w)$	$v(x, \theta + w)$	$v(x, \theta) + u(x; w)$	$v(x; \theta + \Theta \xi)$	$w^\top v(x, \theta)$

and hence takes the form of (3.1). A simple, yet illustrative physical system of this form is the actuated pendulum:

$$m\ell^2 \ddot{q} + mg \sin q = u. \quad (3.4)$$

If several pendulums are observed in the training data, the system parameters $w := (m\ell^2, mg)$ of dimension $r = 2$ may vary across the datasets, but the observations will still be linear functions of the same feature map $v(x) = (\ddot{q}, \sin q)$. The goal is to learn both the task-agnostic and the task-specific components from observations of two different environments $y = w^\top v(x)$ and $y' = w'^\top v(x)$.

To learn a physical system of the form (3.1) across various environments, we propose the following Task-Linear Deep Representation architecture.

Definition 1 (TLDR architecture). The output is modeled as a linear function of task-specific weights $w_t \in \mathbb{R}^r$ with a task-agnostic feature map $v(x; \theta)$:

$$f(x; \theta, w) = w^\top v(x; \theta) \in \mathbb{R}. \quad (3.5)$$

The task-linear function (3.1) and the corresponding model of Definition 1 can be generalized to task-affine functions $f(x; w) = w^\top v(x) + c(x)$ by increasing the dimension and setting the last component of w to 1. It can also be generalized to multivariate observations $y \in \mathbb{R}^m$ with the parametrization $f(x; w) = V(x) \times w$, and $V(x) \in \mathbb{R}^{m \times r}$.

The function $v(x; \theta)$ is an arbitrary parametric model, which can be as complex as a deep neural network. It is meant to capture the task-agnostic features of the system. Our architecture makes a clear distinction between the task-agnostic representation v and the task-specific weights w . In this work, we assume that the dimension of the features $r = r_*$ is known in advance. In practice, this may not be true but r_* can be readily inferred during the training algorithm.

Meta-training The meta-parameters of our architecture are jointly trained by gradient descent as in Algorithm 1.

Adaptation The linear structure of our model allows us to solve (2.3) by ordinary least squares. Other meta-learning approaches, in contrast, require a number of gradient descent adaptation steps.

4. Parameter identification with TLDR

In this section, we study the interpretability of TLDR, *i.e.* the link between the trained parameters and the ground truth.

The first question is whether the trained \bar{v} and \bar{w}_t converge to the ground truth v_* and w_t^* . Unfortunately, this is not the case because the product $w^\top v$ and hence the loss function are invariant to matrix multiplication of the two factors v and w (Fu et al., 2018). For example, the feature-parameter pair $\{v, w\}$ and $v' = P^{-1}v$, $w' = P^\top w$ for some $P \in \text{GL}_r(\mathbb{R})$ produce the same output $y = w^\top v$. However, we show in the following lemma that if the dot product equality $w^\top v = w'^\top v'$ holds for a spanning set of \mathbb{R}^r , then there exists a linear transform that maps the two families of vectors.

Lemma 1. Assume that $T, N > r$. If $W, W' \in \mathbb{R}^{T \times r}$ and $V, V' \in \mathbb{R}^{T \times r}$ satisfy $WV^\top = W'V'^\top$ in $\mathbb{R}^{T \times N}$, and W and V are of full rank, then there exists $P \in \text{GL}_r(\mathbb{R})$ such that $V = V'P^\top$ and $W = W'P^{-1}$.

Assuming fixed training points $\{x_{i,t}\} = \{x_i\}$ across the task datasets, and applying Lemma 1 to the training features $V := (v(x_1; \theta), \dots, v(x_N; \theta))$ and parameters $W := (w_1, \dots, w_T)$, and V' and W' the corresponding ground truths, we obtain the following result.

Proposition 1. In the limit of vanishing training loss $L(\pi) = 0$, the meta-parameters recover the feature map and the parameters of the true system up to a linear transform.

It follows from Proposition 1 that we can identify both the task-agnostic and task-specific components of the system with little additional supervision information: the true solutions are known up to a matrix $P \in \text{GL}_r(\mathbb{R})$. For example, if a number of $s \geq r$ training parameters w_t^* are known, we can solve the ordinary least squares regression

$$\min_{P \in \mathbb{R}^{r \times r}} \frac{1}{2} \sum_{t=1}^s \|P^\top \bar{w}_t - w_t^*\|^2 \quad (4.1)$$

and the minimizer \hat{P} yields the estimates $\hat{w}_t := \hat{P}^\top \bar{w}_t$ for the training parameters. We can then estimate v_* according to Proposition 1 with $\hat{v}(x) = \hat{P}^{-1} \bar{v}(x)$. At test time,




Figure 1. Two-shot adaptation on two out-of-domain dipole environments. The adaptation points are represented by the \times symbols. The vector field is derived from the learned potential field.

a new system parameter w^* can be identified by minimizing $\|w^\top \hat{v}(x) - y\|^2$ in w over the adaptation set.

5. Experimentation on physical systems

We test the performance of TLDR on two physical systems: an electric dipole, as described in Example 1 and the inverse dynamics of an acrobot (or robot arm). The dipolar potential field is learned as a function of the position in a plane: $d = 2$. For the robot arm, the input space is of dimension $d = 8$, with x containing the sine and cosine of the two angles of the arm, and their first- and second-order time derivatives, and the output is the torque applied to the second joint. The equations of motion can be found in (Tedrake, 2022).

Baselines We compare TLDR with two state-of-the-art gradient-based meta-learning algorithms: MAML and CoDA (Kirchmeyer et al., 2022), whose architecture we have adapted to static regression. We implement MAML with the learn2learn Python package (Arnold et al., 2020) and CoDA using hypnettorch (von Oswald et al., 2020).

Evaluation The algorithms are evaluated with prediction error on several out-of-domain (*i.e.* new tasks) test datasets, from which a few labeled examples are provided for adaptation. For TLDR, we also define a feature identification error as $\|\hat{v}(\cdot; \theta) - v_*\|_{L^2}$ and a task parameter identification error as $\|w_t^* - \hat{w}_t\|_2$.

Experimental setup The systems are learned with a fully connected neural network f of 2 hidden layers of width 8 and tanh nonlinearity, and trained with the Adam optimizer (Kingma & Ba, 2015) with 5000 gradient steps. The data points $\{x_i\}$ are located a uniform grid of size 400. For training, $T = 9$ for the dipole and $T = 8$ for the robot arm. Our code is available at <https://anonymous.4open.science/r/meta-learning-1370/>.

Results The test error is plotted as a function of the number of shots in Figure 2. The task-linear meta-model is able




Figure 2. Minimal, median and maximal few-shot adaptation error across test environments as a function of the number of shots. For CoDA and MAML, adaptation is performed with 50 gradient steps.




Figure 3. System identification across the training steps.

to adapt well after only $r = 2$ or 3 shots, whereas the other approaches require more data. The electric fields obtained with MAML and TLDR after two-shot adaptation on two out-of-domain environments are shown on Figure 5. In this few-shot scenario, TLDR adapts accurately to the test environments while MAML has difficulty learning the new fields. Figure 3 shows the test error in the feature space and parameter space (averaged over the training tasks), obtained as in Section 4 with $s = r$. The curves show that TLDR effectively identifies the task-agnostic and task-specific components of the physical systems separately, with vanishing error.

6. Conclusion

We have proposed TLDR, a meta-learning architecture that leverages the task-linear structure of physical systems to learn and adapt efficiently from multiple environments. Unlike other black-box approaches, TLDR enables interpretability of the learned features and parameters.

Although we showed the interpretability of TLDR, the knowledge of multiple training parameters is a restrictive assumption and we would like to investigate a setting in which parameter recovery relies on more realistic prior knowledge. Interesting future research directions include learning from trajectories and forecasting dynamical systems, and validating the performance of TLDR for robotic systems on downstream control tasks using inverse dynamics control.

Broader impact

Although our method is a first step towards interpretable learning of physical systems, the theoretical result we provide is only valid in the vanishing training loss limit. Furthermore, the interpreted parameters can only be recovered after a supervised learning (*i.e.* regression) step, where the output is all the more accurate as supervision increases. Therefore, one should be careful when interpreting the output of TDLR as physical parameters and thorough checks should be performed before using them as inputs to some experiment downstream in the pipeline.

References

- Arnold, S. M. R., Mahajan, P., Datta, D., Bunner, I., and Zarkias, K. S. learn2learn: A library for Meta-Learning research. August 2020. URL <http://arxiv.org/abs/2008.12284>.
- Caruana, R. Multitask learning. *Machine learning*, 28: 41–75, 1997.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
- Fu, X., Huang, K., and Sidiropoulos, N. D. On identifiability of nonnegative matrix factorization. *IEEE Signal Processing Letters*, 25(3):328–332, 2018.
- Haus, H. A. and Melcher, J. R. *Electromagnetic fields and energy*, volume 107. Prentice Hall Englewood Cliffs, 1989.
- Hospedales, T., Antoniou, A., Micaelli, P., and Storkey, A. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021.
- Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., and Yang, L. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015. URL <http://dblp.uni-trier.de/db/conf/iclr/iclr2015.html#KingmaB14>.
- Kirchmeyer, M., Yin, Y., Dona, J., Baskiotis, N., Rakotomamonjy, A., and Gallinari, P. Generalizing to new physical systems via context-informed dynamics model. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 11283–11301. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/kirchmeyer22a.html>.
- Nguyen-Tuong, D. and Peters, J. Using model knowledge for learning inverse dynamics. In *2010 IEEE international conference on robotics and automation*, pp. 2677–2682. IEEE, 2010.
- Raghu, A., Raghu, M., Bengio, S., and Vinyals, O. Rapid Learning or Feature Reuse? Towards Understanding the Effectiveness of MAML. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rkgMkCEtPB>.
- Tedrake, R. *Underactuated Robotics*. 2022. URL <https://underactuated.csail.mit.edu>.
- von Oswald, J., Henning, C., Grewe, B. F., and Sacramento, J. Continual learning with hypernetworks. In *International Conference on Learning Representations*, 2020. URL <https://arxiv.org/abs/1906.00695>.
- Wang, R., Walters, R., and Yu, R. Meta-learning dynamics forecasting using task inference. *Advances in Neural Information Processing Systems*, 35:21640–21653, 2022.
- Yin, Y., Ayed, I., de Bézenac, E., Baskiotis, N., and Gallinari, P. Leads: Learning dynamical systems that generalize across environments. *Advances in Neural Information Processing Systems*, 34:7561–7573, 2021.