

L4DC 2025 notification for paper 23

1 üzenet

L4DC 2025 <l4dc2025@easychair.org>

2025. február 28. 16:44

Címzett: Bendegúz Máté Györök <gyorokbende@sztaki.hu>

Dear Bendegúz Máté Györök,

Your paper #23: Orthogonal projection-based regularization for efficient model augmentation, submitted for possible presentation at the 2025 Learning for Dynamics and Control Conference (L4DC) has now been reviewed. Copies of the reviews are enclosed.

I am pleased to inform you that your paper has been accepted for the 2025 L4DC. If the paper has been selected for oral presentation, you will be notified of this in a separate email by March 21.

We will send instructions for uploading your final paper in a separate email. Camera-ready papers are due on April 25.

The 2025 L4DC will be held as an in-person only conference at the University of Michigan, on 4-6 June, 2025. For the accepted manuscript to appear in the conference proceedings, at least one of the authors must register for the conference before the final version of the paper is uploaded. The paper must be presented in person at the conference venue. A full registration allows the registrant to upload up to two final papers.

The advance program, together with information about the conference format and related instructions, will be available on the conference website at

<https://sites.google.com/umich.edu/l4dc2025/>

and will be updated regularly as details are finalized. The final program for the conference will be available in early May.

Congratulations on the acceptance of your manuscript and thank you for helping us to achieve a high-quality technical program. We look forward to your participation in the conference in June.

Best regards,
2025 L4DC Organizing Committee

SUBMISSION: 23

TITLE: Orthogonal projection-based regularization for efficient model augmentation

----- REVIEW 1 -----

SUBMISSION: 23

TITLE: Orthogonal projection-based regularization for efficient model augmentation

AUTHORS: Bendegúz Máté Györök, Jan H. Hoekstra, Johan Kon, Tamás Péni, Maarten Schoukens and Roland Tóth

----- Overall evaluation -----

In the context of additive model augmentation structures for system identification, this paper introduces a novel method to enhance the learning pipeline by enforcing orthogonal projections. This method is more efficient, requires less parameters and works better in practice.

Comments:

C1 The introduction is written in such a way that it avoids delving into important details about the difference between f_{θ} and f^{ANN}_{η} . I believe it is important to clarify the differences between these two functions/architectures. That is to say, explicitly say what are the common choices and why f^{ANN}_{η} is needed. This is core in the formulation, and I believe the explanation is too short, which requires the reader to resort to other publications. As an example, if f_{θ} allows us to perfectly learn f , then f^{ANN}_{η} is not needed.

C2 Related to C1, the authors say that the goal is to find " η_* , θ_* parameter values that accurately describe the system." These might not be unique, or they may lie on a hyperplane, can the authors clarify the best choices for them?

C3 The paper reads "the goal of the optimization is to minimize the following objective function" related to equations (4). However, this is not an objective function, but rather a set of equalities. What is the objective function to be minimized? I would assume (4a), can the authors clarify.

C4 At the bottom of page 7, the authors say "this approximation is only valid if θ_0 is close to the optimal parameter values of θ ". This does not hold in practice, as parameters move away from the initial seed. What are the implications of this? How expensive is it to recalculate the projections?

C5 Table 1 only shows two values of SNR plus no noise. The choice seems arbitrary, why were only these two values chosen? Can the authors run more experiments showcasing the validity of the method?

C6 The authors claim in the abstract that "To overcome this challenge, this paper proposes an orthogonal projection-based regularization technique to enhance parameter learning, convergence, and even model accuracy in learning-based augmentation of nonlinear baseline models." However, there is no theoretical proof that this method improves convergence. Can the authors soften their claims or provide evidence of convergence?

----- REVIEW 2 -----

SUBMISSION: 23

TITLE: Orthogonal projection-based regularization for efficient model augmentation

AUTHORS: Bendegúz Máté Györök, Jan H. Hoekstra, Johan Kon, Tamás Péni, Maarten Schoukens and Roland Tóth

----- Overall evaluation -----

Summary

This paper discusses an augmented model approach to dynamic physical systems, consisting of a first principle (FP) model in parallel with a neural network (NN). This approach aims to capture physically interpretable information with the FP model and unknown nonlinear dynamics with the NN. To promote physical interpretability and to avoid competing models, it is necessary to ensure the NN captures only the dynamics not explained by the FP model. The authors address this with a regularization scheme that penalizes the loss function for NN outputs that make predictions in the same space as the FP model. For a FP model that is linear (or linearized) in the parameters, a basis for the model is computed with SVD. The output of the NN is then projected onto this space to compute the regularization term added to the model loss function. By promoting orthogonality between the two models, the output is expected to maintain physical interpretability while also capturing other unknown dynamics. Experiments of a digital twin of an electrical vehicle show that the FP + NN parallel model with regularization has increased model accuracy compared to the FP model alone or the FP + NN model without orthogonal regularization but decreased accuracy compared to a state of the art black box model.

Technical Soundness: 3

The use of a projection-based regularizer to promote orthogonality is an elegant way to promote two models in parallel to capture different dynamics. The math involved in this projection is sound.

As noted by the authors, SVD is required at each iteration for FP models where state measurements are not available and for nonlinear FP models where the linearization point is updated at each iteration step (necessary when nominal parameter values are unknown or not close enough to the optimal values), which is computationally demanding. The experiment has access to full state measurements from the data and uses nominal parameter values for a linearization point, so the SVD only needs to be performed once. The impact of performing SVD at each iteration is therefore not addressed experimentally.

The authors note that use of nominal values in the Taylor expansion (rather than updating the linearization point at each iteration, which would require SVD at each iteration) is valid if the optimum parameters are close to their nominal values. Optimal and nominal values are not reported for the experiments, so we cannot evaluate the validity of using these parameters as linearization points.

A physically meaningful FP model should have parameter values that are close to the true parameter values. For example, if the optimal mass (m) determined during training differs significantly from the true mass of the vehicle, the FP model may not be accurately capturing physical laws. A comparison between the model parameters and their true values (for the FP parameters whose true values are known) would help validate the claim of physical interpretability.

Clarity of presentation: 5

The structure of the paper presents the motivation, theory, and experiments in a rational, easy-to-follow manner. Relevant equations are presented and explained well.

Relevance: 4

The paper addresses a challenge that comes with modeling many physical systems: physical laws/first-principle models capture some but not all of the dynamics. These modeling problems would benefit from the ability to capture the dynamics not explained by first-principles while still retaining meaning of the first-principle model parameters.

Significance: 2

The projection-based regularizer discussed in both this paper and [1] is a novel approach to capturing known physical dynamics and unknown nonlinear dynamics. The contribution of this specific paper, the first order Taylor expansion of

a nonlinear FP model and an initialization scheme for the NN, is incremental with respect to [1].

The authors mention three significant advantages of this method: faster convergence, better accuracy, and physical interpretability. Convergence time is not addressed, and would be particularly important for cases where SVD is calculated at each iteration. Physical interpretability could be addressed more thoroughly (as discussed above). Without the validation of physically meaningful parameters and improved convergence time, it is difficult to validate whether this formulation is truly an improvement over the state of the art black box model, which achieves better accuracy in the experiments.

Other comments

Page 6, paragraph below Eq 14: I believe " $N \times N$ " should be " $N_{nx} \times N_{nx}$ "

Page 4, final paragraph states: "fully random initialization can decrease convergence time and even cause stability losses." I believe "decrease" should be replaced with "increase"

[1] Johan Kon, Dennis Bruijnen, Jeroen Van De Wijdeven, Marcel Heertjes, and Tom Oomen. Physics- Guided Neural Networks for Feedforward Control: An Orthogonal Projection-Based Approach. In Proc. of the American Control Conference, pages 4377–4382, 2022.

----- REVIEW 3 -----

SUBMISSION: 23

TITLE: Orthogonal projection-based regularization for efficient model augmentation

AUTHORS: Bendegúz Máté Györök, Jan H. Hoekstra, Johan Kon, Tamás Péni, Maarten Schoukens and Roland Tóth

----- Overall evaluation -----

The paper investigates an interesting approach to combining first-principle methods with artificial neural networks.

Some suggestions:

- Including more details on how the proposed method enhances interpretability would be valuable.
- The method involves a significant preprocessing step; it would be helpful to include a numerical study demonstrating how this preprocessing affects convergence, particularly in terms of wall clock performance.
- Investigating the potential of incorporating a linear regression preprocessing step could also be an interesting direction for further development.