The APEX project - Artificial Intelligence for Policy Excellence in the Climate Crisis: initial results and outlook

J.P.S. Böger and **O.B. Lassen** and **F. Madaleno** and
C.L. Azevedo and **G. Cantelmo** and
F. Rodrigues and **S.Agriesti** and **F.C. Pereira**
Department of Technology, Management and Economics, Technical University of Denmark samaa@dtu.dk

Abstract

Policy-making requires comparing multiple scenarios involving many variables and uncertainties, assessing trade-offs between economic, environmental and social factors. To make informed decisions, policymakers often rely on data from simulators, to frame the impacts of choices on a set of possible futures. These simulations, though, can be computationally intensive and run over the course of days, if not weeks. This problem can become an insurmountable barrier, for example, in tackling the climate crisis, which requires long-term planning and the comparison of different policy trajectories on a large scale. We argue that to solve this problem, the generalizing power of simulators has to be coupled with the computational efficiency of machine learning solutions. In this position paper we detail how simulators and non-parametric approaches can be enhanced through each other's strengths and how the results may increase the amount of scenarios that can inform policy-making, while reducing uncertainties.

Keywords: Surrogate Modeling, Generalizability, Robustness

1 Introduction

We present here the results from the first year of the APEX (Artificial Intelligence for Policy Excellence in the Climate Crisis) project and the overarching vision. APEX aims at providing policy-makers with more powerful tools to assess not only few scattered data-points in the future, but complete trajectories and a variety of scenarios not constrained by computational limits (for example, in the case of environmental or mobility applications). It does so by combining the domain knowledge and reliability of simulators with the speed of Machine Learning (ML) solutions. The intuition behind the project is that ML solutions are trustworthy if enough data is available to properly frame the dynamics of a phenomenon in distribution (ID). Still, this is rarely the case for long-term implementable policies, for which data cannot be obtained until the policy is implemented (and thus, the related investment has been carried out). Besides, for policies designed to tackle climate crisis scenarios (e.g. historical flooding in Copenhagen [Vandervoort et al., 2025]), data simply won't be available until it is too late to act and the dynamics of the problem won't be the same as the current ones (i.e. out of distribution - OOD[1]). This is where the generalizing power of simulators can be harnessed to improve ML solutions. We argue that ML solutions can be fully relied upon even when data is scarce, if the domain knowledge and algorithmic structure of simulators is embedded in the ML architecture [Xu et al., 2019], making them able to handle distribution shifts. To achieve these results, the APEX project is built over three research streams:

- Embedding in the surrogate the first-principle dynamics (e.g. physics laws) of the underlying phenomenon will improve its performance OOD
- Designing the architecture of the surrogate to mimic the algorithmic structure of a simulator will guarantee higher generalizing power
- By expanding the theory and applicability of causal abstraction we can minimize Interventional Consistency Loss between surrogate and simulator

Each stream focuses on generalization but can then be exploited to speed up experiments, by allowing to replace computationally expensive simulations with surrogate runs, trusting the latter to perform OOD. In the following, we detail the initial outcomes from each of the three research directions, with a specific focus on neural network (NN) architectures.

2 Embedding dynamics to complement data

Physical laws and problem specific dynamics are usually framed through ordinary (ODEs) or partial differential equations (PDEs). In hydrodynamics, for example, one of the foundational relations is described through the Navier-Stokes equations [Constantin and Foiaş, 1988], while diffusion phenomena can be capture by the Bateman–Burgers PDE [Bonkile et al., 2018]. Still, it is not intuitive how to embed these physical laws into surrogates in a way that does not degrade the performance ID. We plan to address this gap by combining Koopman operator theory [Budišić et al., 2012] with a neural architecture (MixFunn [Farias et al., 2025]) that is designed to integrate multiple parameterized nonlinear functions, augmented with quadratic neurons to capture input interactions. At the present stage, we applied Koopman theory to the SIR model [Hethcote, 2000] to predict the population susceptible to a specific illness and the share of infected and recovered. It does so by providing a linear representation of nonlinear dynamics in an infinite-dimensional space, which is approximated through a finite set of dictionaries of functions. While the results with state-of-the-art methods are promising (Fig. 1), the dictionary selection remains ad hoc and problem-specific. We plan to use MixFunn to build Koopman dictionaries from activation functions inspired by known solutions of ODEs and PDEs, increasing interpretability and simplifying the embedding of physical laws and dynamics.

[1] Depending on the field, this phenomenon could also be defined as out of domain [Wald et al., 2021]

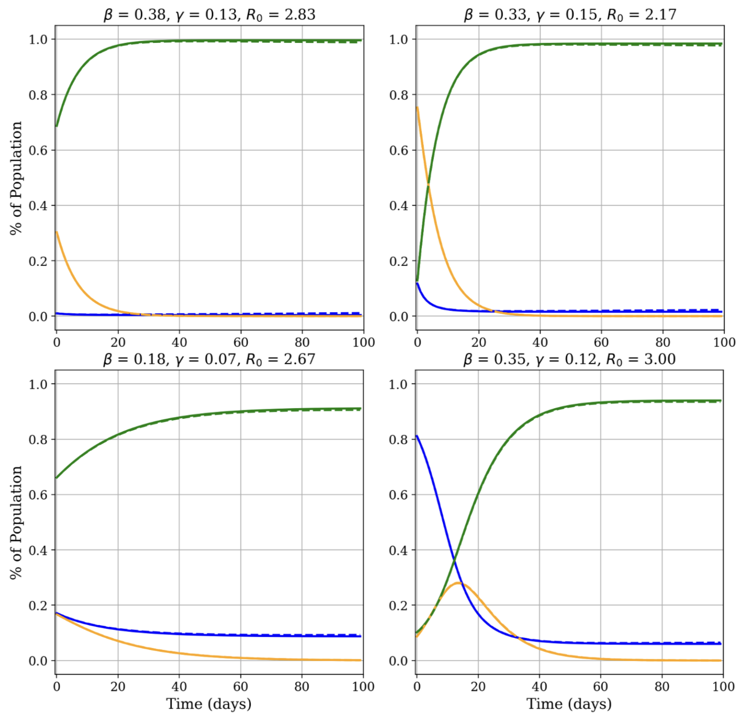

Figure 1: Application of the Koopman operator to predict the SIR outputs 3 time steps in advance

3 Designing a NN with a specific algorithmic structure

Another approach to make surrogates more robust, is to align their architecture with the algorithmic structure of the surrogate. The approach has already been tested [Xu et al., 2019], but it has yet to be applied to a complex simulator. We chose the Traffic Assignment Problem (TAP) as our first experimental use case for a simulator and identify Graph Neural Networks (GNNs) as NNs whose structure better aligns. Our first set of experiments indeed shows promising results [Lassen et al., 2025]. In these experiments, a Message-Passing Neural Network is tested OOD by running the surrogate in areas of the input space where the state variables of capacity, free flow speed and demand are outside the training intervals. The results reported in Fig. 2 show that the surrogate succeeds in predicting flows when capacity is OOD. Still, the chosen MPNN (with the GatedGCN as specific layer architecture [Bresson and Laurent, 2018]) struggles when the demand extends OOD. Additionally, the performance of GNNs degrades when applied to modified graphs because they fail to capture the true underlying dynamics. We plan to build an improved MPNN architecture that mimics the simulator's iterative structure such that the performance does not degrade OOD.

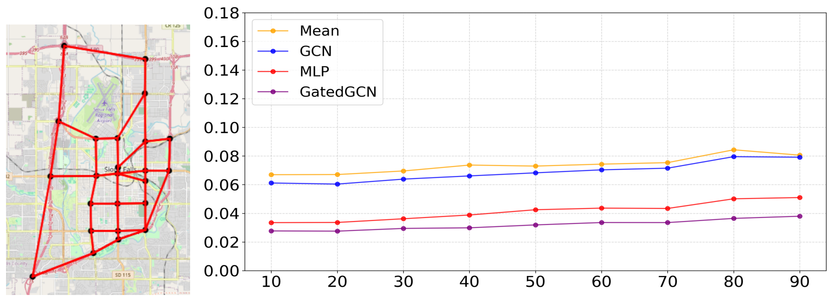

Figure 2: Left: The road network, Right: MAE (y-axis) as a function of the number of edges with capacity values OOD (x-axis)

4 Exploiting causal structures to ensure interventional invariance

One key application of simulators is to assess and compare different interventions [Christiansen et al., 2021], to identify the best policy. NNs cannot easily replicate this behavior, as they cannot guarantee a consistent behavior under interventions when different portions of the input space are modified. This makes their application for policy comparison challenging. We plan to address this limitation by expanding the theory and applicability of causal abstraction, an approach that aims at minimizing Interventional Consistency Loss [Dyer et al., 2025] between surrogate and simulator (Fig. 3).

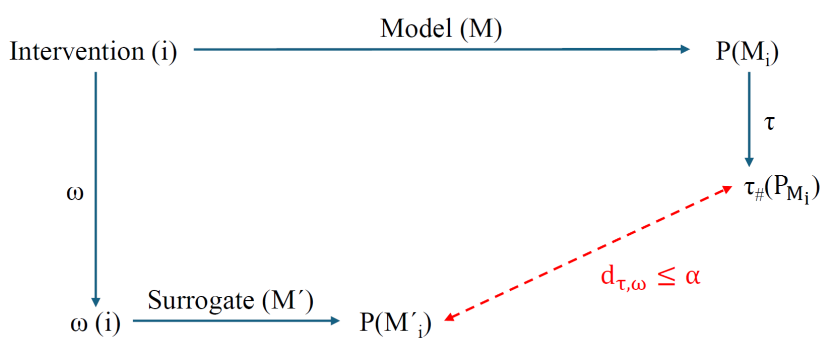

Figure 3: Formal framework for Causal Abstraction, adapted from [Dyer et al., 2025]. τ is a transformation mapping corresponding states of M and M' while ω is a transformation mapping corresponding interventions. d is the abstraction error.

The methodological improvement will focus on identifying the most promising interventions by: ensuring interventional consistency, extrapolating from partial sets of interventions and designing the optimal subset for maximum extrapolation. Once the surrogate behaves upon unseen interventions with the same dynamics captured ID, the surrogate shall be applied to test multiple policies without the need to verify at each run that the surrogate has not steered towards a portion of the OOD space where it generalizes badly.

5 Conclusions and outlook

The results from the first year of the APEX project are promising but the long term vision, to be developed in the next 4 years, is wider. The three building blocks will be developed to each tackle a limitation of surrogate modeling by embedding different aspects of domain knowledge from simulators. Once this is accomplished, the three streams of research shall converge to one and all the developed solutions shall result into a methodology for developing reliable surrogate models of a large-scale and computationally expensive simulator.

References

M. P. Bonkile, A. Awasthi, C. Lakshmi, V. Mukundan, and V. S. Aswin. A systematic literature review of burgers' equation with recent advances. *Pramana*, 90(6), 2018.

X. Bresson and T. Laurent. Residual gated graph convnets, 2018. arXiv preprint arXiv:1711.07553.

M. Budišić, R. Mohr, and I. Mezić. Applied koopmanism. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4), 2012.

R. Christiansen, N. Pfister, M. E. Jakobsen, N. Gnecco, and J. Peters. A causal framework for distribution generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10), 2021.

P. Constantin and C. Foiaş. *Navier-stokes equations*. University of Chicago press., 1988.

J. Dyer, N. Bishop, Y. Felekis, F. M. Zennaro, A. Calinescu, T. Damoulas, and M. Wooldridge. Interventionally consistent surrogates for complex simulation models. In *Advances in Neural Information Processing Systems*, volume 37, pages 21814–21841, 2025.

T. D. S. Farias, G. G. de Lima, J. Maziero, and C. J. Villas-Boas. Mixfunn: A neural network for differential equations with improved generalization and interpretability, 2025. arXiv preprint arXiv:2503.22528.

H. W. Hethcote. The mathematics of infectious diseases. *SIAM review*, 42(4): 599–653, 2000.

B. Lassen, S. Agriesti, M. Eldafrawi, D. Gammelli, G. Cantelmo, G. Gentile, and F. Pereira. Learning traffic flows: Graph neural networks for metamodelling traffic assignment. In *In 2025 9th International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS)*, 2025.

A. P. K. M. Vandervoort, M. N. da Costa, M. W. Petersen, M. Drews, S. Haustein, K. M. Morrissey, and F. C. Pereira. Using reinforcement learning to integrate subjective wellbeing into climate adaptation decision making. In *In ICLR 2025, Tackling Climate Change with Machine Learning workshop*, 2025.

Y. Wald, A. Feder, D. Greenfeld, and U. Shalit. On calibration and out-of-domain generalization. In *Advances in neural information processing systems*, volume 34, pages 2215–2227, 2021.

K. Xu, J. Li, M. Zhang, S. S. Du, K. I. Kawarabayashi, and S. Jegelka. What can neural networks reason about?, 2019. arXiv preprint arXiv:1905.13211.