

Foundation Model for Real-Time Model Selection and Fitting

Fitting data to a variety of models is a fundamental challenge in the monitoring and control of dynamical systems across science and manufacturing domains. In this work, we present a compact foundation model designed for adaptive function selection and regression. The proposed architecture utilizes 1D convolutional neural networks (CNNs), augmented by physical constraints, to facilitate the selection and generation of functional forms that fit the underlying data. The model is capable of processing input data from spectroscopic measurements or time series and is grounded in a collection of 9 common mathematical functions derived from physical principles.

The data generation process for training involves randomly sampling parameters within predefined ranges for each function, evaluating these functions over a fixed interval of the independent variable (x). Models are trained on a large dataset with a sample size of 10,000 per function. The model is trained to accommodate up to five variable equations in the form $f(x)$, allowing it to handle diverse datasets where different subsets may be optimally described by distinct functional forms. The architecture is inherently extensible to accommodate additional functions as needed.

Once trained, the model can estimate the parameters for multiple functions in a single forward pass, generating predictions by evaluating each function with its corresponding parameter set. A built-in model selection mechanism ensures that the most appropriate function is chosen to represent the data. This approach offers an alternative to traditional methods like symbolic regression and sparse optimization of nonlinear dynamics, both of which are notoriously difficult to optimize due to high computational complexity and non-convexity of the solution space.

We further evaluate the performance of this model under varying noise levels, showcasing its robustness through stochastic averaging during batch processing. Despite its compact size, with only ~5,000 parameters, the model exhibits impressive scalability and performance improvements over time. Notably, the architecture is simple to convert and deploy on hardware platforms like FPGAs using the HLS4ML framework, making it highly suitable for real-time applications.

By leveraging this physics-constrained fitting framework, the model provides an efficient and lightweight solution for signal processing and real-time monitoring and control, paving the way for advanced, low-latency adaptive systems in a wide range of scientific and industrial applications.

Focus areas

Authors: TSAI, Cymberly; AGAR, Joshua

Presenter: TSAI, Cymberly