A Physically informed Surrogate Approach to Causal System Modeling

Michael. J. Barrow*, Min S. Cho**,
Paul E. Grabowski, James A. Gaffney, Rushil Anirudh,
Jayaraman Thiagarajan, Joshua B. Kallman, Hai P. Le,
Howard Scott, Peer-Timo Bremer, and Jayram Thathachar

barrow9@llnl.gov*, cho28@llnl.gov**

In basic scientific research, complex systems are typically modeled using multiscale physical systems, which although accurate are barely computationally tractable. Surrogate models promise great speedups over traditional methods, with deep learning models being particularly interesting since they can model "arbitrary" inputs and outputs. However, traditional surrogates are hampered during training by non-physical error formulations which can cause spurious simulation crashes.

In so-called Inertial Confinement Fusion (ICF) problems, the most computationally intensive step of multiscale simulation is often the non-local thermodynamic equilibrium (NLTE) calculation (up to 90% of wall clock time). NLTE input includes scalars such as the temperature and density of plasmas in space and time. And we typically seek to model non-scalar output quantities, i.e., opacity and emissivity histograms of plasma resulting from a laser drive to the system. Unfortunately, building a surrogate using the L2 type losses, typical in deep network training, can introduce unstructured errors in which outputs are self-inconsistent, or even unphysical (do not preserve energy).

Here we develop a new Physics Informed Loss function (PIL) that targets inconsistency issues by reformulating a surrogate's output representation. Our PIL reformulation introduces two physical quantities transformed by opacity and emissivity, which enforce physical constraints to the surrogate model in terms of energy and radiation transport:

Firstly we enforce an energy conservation constraint by requiring outputs be consistent with one another in a physical way. In particular, we define system energy histogram bins, and within each bin, we tie the emissivity and opacity outputs together in such a way that the surrogate targets energy conservation during training, rather than an arbitrary correlation between emissivity and opacity. These quantities are appropriately scaled so that no one bin is over-represented during training.

Secondly we regularize the training loss with a physically informed filtering that concentrates on spatiotemporal regions where the NLTE code is a key part of the modeling stack within the multiscale simulation system. The filter has two effects on the 'mean free path' of an energy bin, where the 'mean free path' energy (MFP) is indicative of the energy state of plasma gas to be modeled. The first effect is to de-emphasize the large dynamic range that the MFP can assume throughout system evolution, such that the network is better able to represent this physical quantity. The second is to remove regions where the MFP is too low to be of interest (indicates that the region contains material that is not plasma), and to remove regions where MFP is too high to be of interest (indicates that the region is in vacuum). It is important to remove such regions during training since representing them is complex, consuming surrogate modeling resources, yet the modeling power is wasted since NLTE is not useful to the multiscale simulation as a whole in these regions.

With our techniques, we are able to model the plasmas dynamics of ICF in a physically plausible way due to the energy conservation constraint, and an informationally efficient way due to the physically informed filtering.

Manuscript prepared by LLNL under Contract DE-AC52-07NA27344

When training our networks we have sought to balance the learning objective of surrogate NLTE models between a physics informed PIL training loss, which emphasizes the physics isolated by our energy conservation and plasma state constraints, with a traditional learning objective that seeks only to accurately predict input and output correlations.

Our preliminary results focus on "λ_{PIL}". λ_{PIL} is a hyper parameter knob controlling the emphasis of PIL on Surrogate model learning and can be used in concert with other network specification parameters to tune overall model loss on a ICF NLTE dataset built for copper Hohlraum simulations. Our design space hyper parameter sweep of surrogate network architecture indicates a relationship between a large λ_{PIL} and low NLTE prediction error. In other words, even though the PIL does not specifically target learning accurate NLTE prediction, using it seems to guide our networks to learn a better correlation in the embedded space to reconstruct NLTE data than just trying to learn such a correlation through traditional optimization.

In ongoing work we are exploring integrating PIL surrogates into a complete multiscale ICF model to explore the tradeoffs of our surrogates in computational efficiency and stability.

In conclusion, PIL simulations provide more physically direct ways to model the underlying processes involved in plasma hydrodynamics than traditional methods which tackles issues in traditional surrogates that can lead to system crashes.