AI-assisted HPC simulations and application to astrophysics and cosmology

Tiziana Di Matteo, Rupert Croft and Yueying Ni (Carnegie Mellon), Simeon Bird (UC Riverside), Yin Li (Flatiron), Yu Feng (Google)

Dealing with finite resolution and resources. As with all HPC simulations of physical systems, cosmological simulations of galaxy formation [1, 8] are limited by insufficient computational resources. It is possible to directly follow the equations of gravity and gas dynamics either at high resolution in a small volume or at low resolution in a large volume. However, key processes of interest often result from the interplay of large scale environments and small scale physics. Drawing from the ongoing revolution in AI (specifically, Deep Learning, DL, e.g., [3]) it is becoming possible to bridge this gap from large to small scales. Neural Networks (NNs) have been developed that can learn from high resolution image data [2, 5], and then make accurate super resolution versions of different low resolution images. We are beginning to apply these techniques [6] to physical modeling in HPC simulations. The ideas are generally applicable, and our own use case is cosmological hydrodynamics, leading to super resolution modeling of galaxy and black hole formation.

Training Generative Networks. The heart of these AI-assisted simulations are *Generative Adversarial Networks (GANs)* [4, 7]. A GAN is a class of DL system in which two NNs contest against each other in a game: the generative network generates candidates while the discriminative network evaluates them. Given a training set, this technique learns to generate new data with the same characteristics as the training set. GANs can be used to generate entirely new data from initially random inputs, or in the case of super resolution simulations we use low resolution full hydrodynamics data as input to the generative network and produce data statistically consistent with high resolution simulations below the resolution scale. Training sets are generated by running simulations at different native resolutions, and results from different physics codes can be combined.

The path to hybrid AI/hydrodynamic modeling. Three different approaches are being tried, in increasing order of complexity (i) enhancement of simulated galaxy images; (ii) post processing augmentation of simulated three dimensional gas, dark matter and stellar density fields; and (iii) training and super resolution modeling inside running simulations ("on the fly") which allows for causal feedback between the generated and physically simulated scales. .

On-chip AI: The suites of super resolution models will make use of on-chip AI instructions becoming available in HPC, such as Intel's AVX-512 Vector Neural Network Instruction set available on the Cascade Lake processors used in the NSF's Frontera. Such technology holds great promise to study problems where large scales and high resolution are closely linked. In the case of astrophysics these super resolution simulations will bridge the gap in scales between the formation of stars, planets and black holes and the dark matter and dark energy that dominate the large-scale structure of the Universe. Our understanding of both ends of this range — from small to large — will increase as a result.

Impact on NSF Cyberinfrastructure We will be running AI algorithms inside advanced physics simulations, with training and use happening in real time, as the AI will affect the running simulation and learn from it at the same time. This level of coupling between AI and direct physical modeling is unprecedented and should lead to advances in the ways that AI can be used. Flexible tightly coupled systems of GPUs and CPUs with on-chip NN would be ideal resources for both NN training farms, and on the fly AI/fluid dynamics runs. If this work fulfills its promise it could revolutionize the use of AI-enabled chips in exascale high performance computing.

REFERENCES

[1] R. Croft, T. Di Matteo, N. Khandai, and Yu Feng. Petascale cosmology: Simulations of structure formation. *Computing in Science & Engineering*, 17(02):40–46, mar 2015.

[2] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *CoRR*, abs/1501.00092, 2015.

[3] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

[4] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.

[5] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. *CoRR*, abs/1609.04802, 2016.

[6] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.

[7] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.

[8] Mark Vogelsberger, Federico Marinacci, Paul Torrey, and Ewald Puchwein. Cosmological simulations of galaxy formation, 2019.