Project Title:
An adversarial approach to self-consistent astrophysical simulations.

Summary

I propose to hire a graduate student to work on deep generative adversarial architectures for constructing initial realizations and run-time validation of simulations. This approach is novel and may lead to entirely new leadership in computer-aided astrophysics.

GANULATIONS

Initial conditions for simulations are generally motivated by mathematical principles describing the system's macroscopic characteristics. The actual starting point is then realized by random sampling from initial distribution functions. After the simulations are finished, the result is compared to observations or analyzed in a more theoretical context. To further study the initial parameter space, small variations in the initial conditions are mapped in a systematic, human-guided fashion or via MCMC until the best match with the observations is achieved.

This entire process can be automated using machine learning through a generative adversarial neural network (GAN, discovered in 2014). Generative adversarial neural networks are deep architectures comprised of two networks in a generator-discriminator setup. For GANULATIONS, the discriminator space (the observations) is known, and the generator's latent space (the initial conditions) are constrained by the underlying physics. Both networks compete; one generates initial conditions (the generator), and the other tries to discover this impostor from observations (the discriminator) as depicted in Fig. 1. The GAN network iteratively takes physically motivated realizations and maps these into conditions until the discriminator cannot tell whether the data originated from the observations or the generator. A GAN is ideal for this semi-supervised approach, and training improves continuously while more simulation and observational data become available.

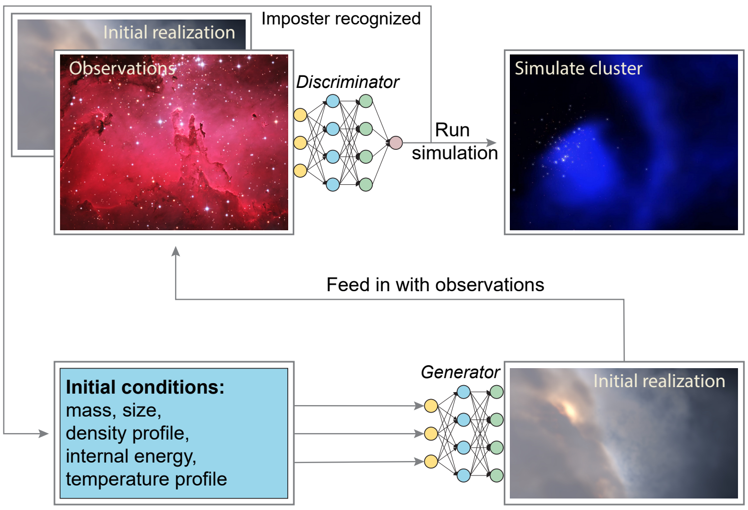

Figure 1: A GAN network setup generating realizations for simulations. The *discriminator* compares observations with physically motivated initial conditions (top left). If the impostor is recognized new initial conditions are selected (bottom left) to produce a new realization using the *generator* network (bottom right). Simulations start (top right) when the *discriminator* network is unable to recognize the initial conditions as the impostor.

The advantage of using a GAN is that initial conditions can be motivated directly by observations: Observed maps of column densities or ionization levels, stellar positions, velocities, etc. The GAN can subsequently, generate a self-consistent realization that is directly motivated by the observations and constrained by the theoretical framework (such as continuity, pressure gradients, kinematic structure, etc.). A GAN can deal with incomplete observational data, such as limited spectral coverage, discrete evolutionary stages, partial spatial coverage or the reduced dimensionality of the observations.

Figure 2: A photo of Ghent city (left) is used together with a painting of van Gogh (right) to construct a new middle image in which scenery (from Ghent) and style (from van Gogh) are combined [1].

In figure 2 we give an example of how this works. From an astrophysical perspective, one takes a numerically generated set of initial conditions (represented with Ghent's city to the left) and an observed temperature or density structure of a molecular cloud (represented by a painting of van Gogh to the right). The GAN will make a new distribution (middle image) consistent with the simulation results and the observations.

Eventually, the simulation results can also be compared with observations using a similar GAN setup. This approach can be extended to an entire chronology of machine-validated simulation-results. This could lead to initial conditions that are directly motivated by some observation, but in which also the temporally resolved simulation results are validated with observations. This cascade of GAN validation steps at run-time leads to a time-resolved consistency between simulations and observational data.

This is complicated, in particular, because the observed systems may intrinsically have different initial conditions, which should not necessarily (and probably not desirable) lead to a converged GAN. On the other hand, intrinsic variations in observational data, uncertainties in the initial conditions, freedom in the selected parameter space and systematic errors in the observations will provide inherent freedom allowing the GAN to find a Nash equilibrium.

The proposed setup of a single GAN has been demonstrated to arrive at a Nash equilibrium, but this does not guarantee that they also converge in this complicated task. However, if GANs fails to converge we will implement reinforcement-learning instead of unsupervised learning. This would be a revolutionary (high-risk) approach, and the outcome cannot be predicted without trying.

References

[1] Gatys, L. A., Ecker, A. S., & Bethge, M. 2016, in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2414–2423