Accelerating Bayesian inference in astrophysics through interpretable representations with normalizing flows

Dupourqué Simon

Statistical inference is at the heart of many publications in astrophysics. Today, the paradigm is shifting towards a generalisation of Bayesian inference, which is robust in providing parameter distributions given a model and a set of observations, and avoids local minimum traps. However, such approaches are slow and require a lot of computing power to be usable. Speeding up Bayesian inference is a key feature for the next decades of astrophysics and can be investigated using neural networks. In particular, simulation-based inference trains a normalising flow to learn a mapping between the parameter and observation space, and can be easily conditioned and inverted to provide the parameters for a given inference problem. Normalising flows are invertible transformation between probability densities. Their neural networks backend is mathematically motivated and leads to lightweight and fast-to-train architectures. Such approaches were originally designed to perform likelihood-free inference, but this aspect helps in designing interpretable inference pipeline, especially when dealing with high-dimensional observables. As they can learn arbitrary distributions, these networks take the best from a meaningful and compressed representation of the data, pushing us to carefully design features that are physically motivated for a given inference problem.