Analyzing Coarse-Grained Model Assumptions Through Variational Autoencoders

Jacob Monroe [C, S] and Vincent Shen
Chemical Sciences Division, NIST, Gaithersburg, MD, U.S.A.
jacob.monroe@nist.gov

Bottom-up coarse graining of molecular simulations leads to faster sampling in a reduced configuration space, but innately requires a loss of information. Structural and thermodynamic consequences related to this lossy encoding are currently difficult to assess. Even when back-mapping to atomistic coordinates, there is no guarantee that the resulting structures are not biased by the lower-resolution model or that these structures satisfy the original all-atom statistical ensemble of interest. We explore this problem from an information theory perspective, employing variational autoencoders (VAEs) to learn coarse-grained models and back-mappings of a simple 2D lattice gas. Training a VAE represents a coarse-graining problem in which we seek the most useful coarse-grained representation for reconstructing an original input by passing through neural networks that perform encoding and decoding operations. Our work reveals that molecular coarse-graining based on relative entropy optimization is a special case of the encoding portion of a VAE, which we demonstrate in practice by training VAEs under different assumptions concerning the coarse-grained and full-space probability densities. As a direct consequence of modeling probability densities in a Bayesian framework, Monte-Carlo (MC) moves achieving large configurational changes in the fine-grained ensemble arise naturally. We show that these moves efficiently transition the simulation along the discovered coarse-grained space and will always be accepted in the limit of a perfectly trained VAE. For various VAE models, we examine the limitations of such MC moves and explore the thermodynamic and structural consequences of the assumptions entering into encoding and decoding probability distribution models.