Few-shot Deep Generative Transfer

Prathosh A P and Chandra Sekhar Seelamantula

Deep generative models learn to sample data from an unknown distribution. Most of the modern generative models do so by learning a transformation from the samples of a latent or noise variable that has a known distribution to the samples from distribution of interest. The family of Auto-Encoder-based generative models use variational inference to maximize a bound on the data likelihood, the family of adversarial generators such as GANs learn to sample by solving a min-max optimization game and the normalizing flow-based methods utilize tractable transformations between the latent and data distributions.

Training of generative models requires a large amount of data points from the unknown distribution. This may not be feasible in certain applications (such as clinical image generation) which makes it hard to learn generative models. However, it is reasonable to assume that there exists data from a distribution (call it source) that is "close" to the distribution of interest (Target distribution). In this setting, we wish to address the following problem - ***How to learn a deep generative model using a few examples (10-1000) of a target distribution given there exists abundant data from the source distribution?*** We term this problem as few-shot generative transfer.

We assume that divergence metric between the source and target distributions are upper-bounded and there exists a good-sampler for the source data. With these assumptions, the objective adopt the source-generator to the target. Multiple approaches including latent space transformations, sparsity regularizers, restriction of optimization landscapes, gradient regularization etc. may be explored. Further, we aim to come up with theoretical bounds on quality of the generated target in terms of its distance from the source. We also wish to draw parallels from the studies in neuroscience on information transfer in the human brain to come up with techniques. The application of such models can span a across lot of domains especially in clinical image processing. For instance, one could aim to generate samples of liver CT images given a large amount of chest CT, generate samples of histopathalogical images of one biopsy tissue given data from other organ etc.