


Mapping the Multiverse of Latent Representations

What is a Multiverse?

Data driven analysis requires many different choices that may support different conclusions. Multiverse analysis embraces this variation, rather than keeping decisions hidden or implicit. We define a **latent-space multiverse** as a collection of latent-space models, parameterized by different algorithmic, implementation, and data choices.


Topological Multiverse Analysis

We introduce a topological distance between representations called **Presto**: Projected Embedding Similarity based on Topological Overlays. Capturing topological and geometric features of latent space projections, **Presto** provides a scalable tool for comparing pairs of representations. Because our method uses *persistence landscapes* we can even do statistical analyses on distributions of latent spaces that arise in a multiverse.




Latent-Space Variability

Presto can measure variation in latent structure across different Variational Autoencoders in a multiverse defined by common hyperparameter combinations. Despite similar generative performance, these models produce very different latent spaces.




Evaluating Parameter Sensitivity

We assess the latent-space variability introduced by specific hyperparameters. The latent space of InfoVAE is far less stable than that of its competitors when tuning key hyperparameters.




Assessing Parameter Compressibility

We measure the compressibility of hyperparameter grid searches using **Presto**, ensuring each model is sufficiently similar to a representative latent space in the compression.



Understanding Parameter Reusability

We compare latent-space relationships (*multiverse metric spaces*) that arise from identical grid searches on different image datasets. High correlations (*) indicate opportunities for hyperparameter transfer.



For any learning task, there are many possible models that arise from different **algorithmic, data, and implementation choices**.

Each model is a *universe* that we can learn from. Considering the ensemble gives rise to the **multiverse**.

We analyze the multiverse of latent representations for *generative models* and develop measures for (hyper-)parameter **sensitivity, compressibility, and reusability**.

