

Finding a Latent Space for the Virgin Mary

Nico Zevallos

Introduction

In the Catholic faith some images of the Virgin Mary are allowed to be venerated. These are images or sculptures to which miraculous power is ascribed. These images are copied and reproduced and in turn are ascribed their own miracles.

As a whole, they represent a deep archetype that crosses cultural boundaries and reaches beyond its scope in Catholicism.

My goal was to visualize a latent space for a complex archetype. What underlying structure exists for such a ubiquitous image? Is there enough variety for the computer to learn to produce new images? Would these images even be new?

Methods

All images downloaded using the google image search API

Images were queried using face++, an online face detection API and pruned to eliminate Johns and Cherubim

Images were cropped and normalized using facenet, as described in "FaceNet: A Unified Embedding for Face Recognition and Clustering"

Facenet Variational Autoencoder trained to generate latent space.

t-Distributed Stochastic Neighbor Embedding was used to further reduce the latent space to two dimensions.

Finally, the embedded space was relaxed into a square shape using Bilinear Coons Patch Image Warping

Results

A 100 dimensional latent space was created from the relatively small (1300 image) dataset. Outputs of decoded latent vectors throughout the space remain convincing at the small output resolution of 64 by 64 pixels.

The latent space seems to encode some perceptual variance such as head rotation quite well.

An interactive was created to allow users to move through the latent space and visualize its outputs.

Future work could include increasing the resolution of the output and finding new sources of data such as museum collections.

Conclusion


Overall this project was a success. A latent space was generated which was robust enough to create new images and encode structure in the images. Interpolating through the space is expressive and varied.

Stopping anywhere in the latent space results in a generic mary-like figure. Overwhelmingly she is white, often she has blue eyes and blonde hair. Clearly, the conglomerate mary isn't a visualization of a historical person. However, while watching the faces morph and warp through each other one is reminded of how an archetype is forged from others, such as the Aztec mother goddess Tonantzin who became Guadalupe, Mutemwiya in Egypt, or even Plato's mother Perictione in Grece.


Sample of dataset



Training VAE



Two dimensional latent space



Interpolated outputs

