1 Appendix

1.1 Compute

All model training and subsequent analyses were conducted on on-premises GPU computing servers. GPUs were orchestrated through the Run:AI software, which enabled fractional GPU allocations for containerized compute environments. For this work, we used at the most 0.5 of an A40 GPU, the equivalent of 24 GB of GPU RAM. CPU usage was minimal and all analyses downstream of model training were reproduced on Google Colab (CPU-only) and could run on a standard laptop. Our computing environment can be exactly reproduced by using our general-purpose Docker image available at: `https://hub.docker.com/r/talmo/tf-extras`

1.2 Supplementary Figures

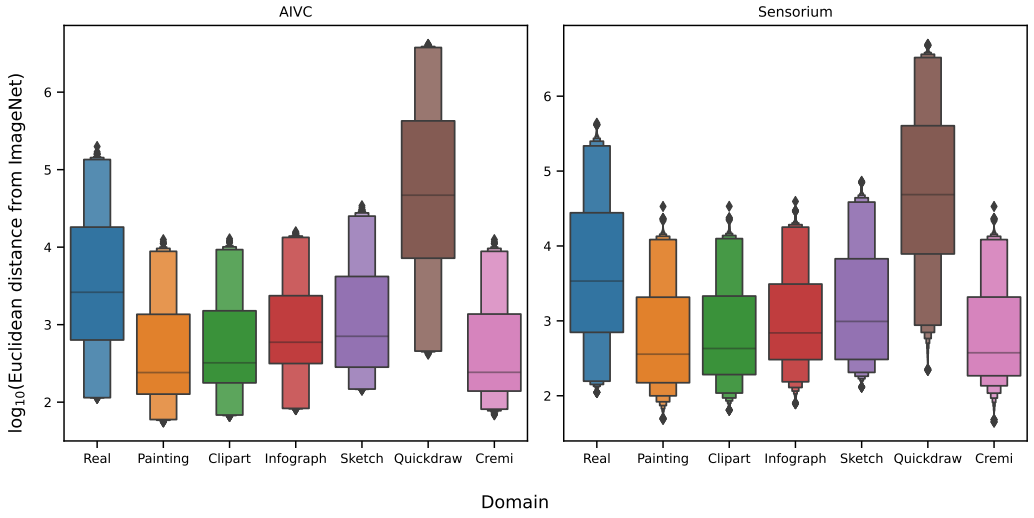

Figure 1: Euclidean distances of activations from each domain from ImageNet

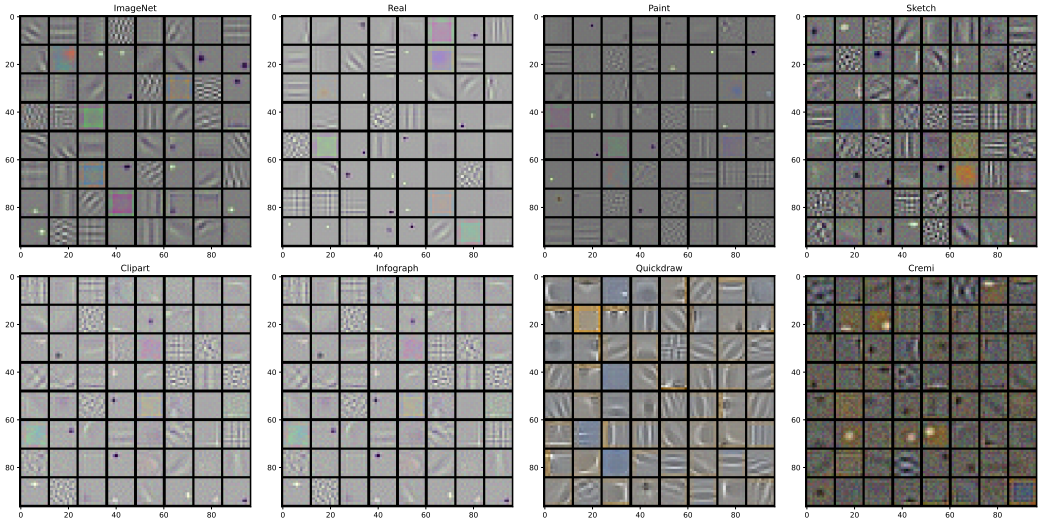

Figure 2: Filters learned by the first convolutional layer of each network

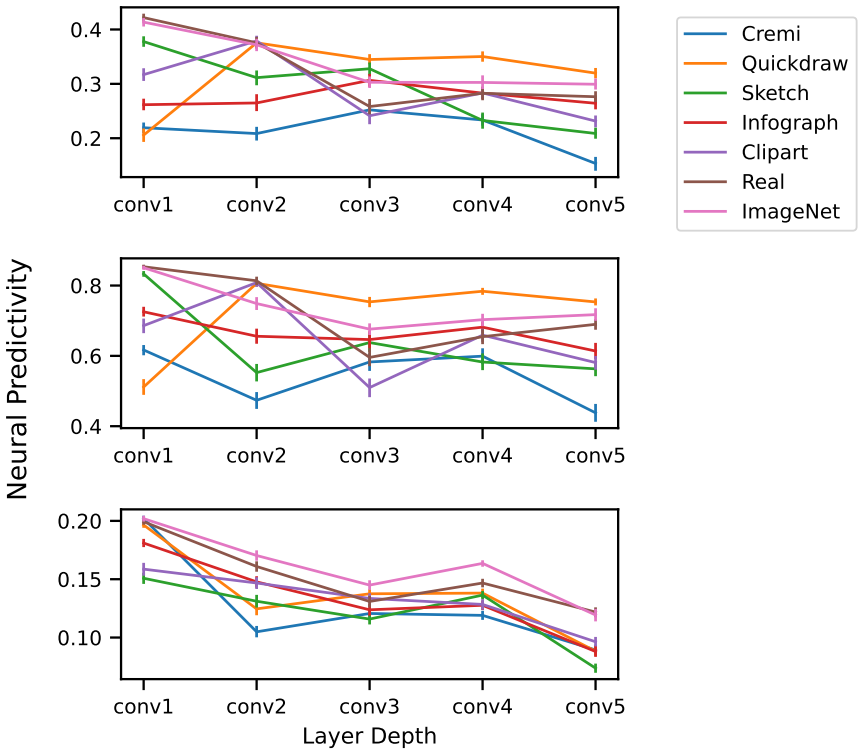

Figure 3: Neural Predictivity as a function of depth. We see a general trend that Neural predictivity decreases in lower layers but we also see that within the first convolutional layer, natural image domains have the higher neural predictivity, while it is less conclusive in later layers. The order is AIVC-calcium, AIVC-neuropixels, and sensorium from top to bottom.

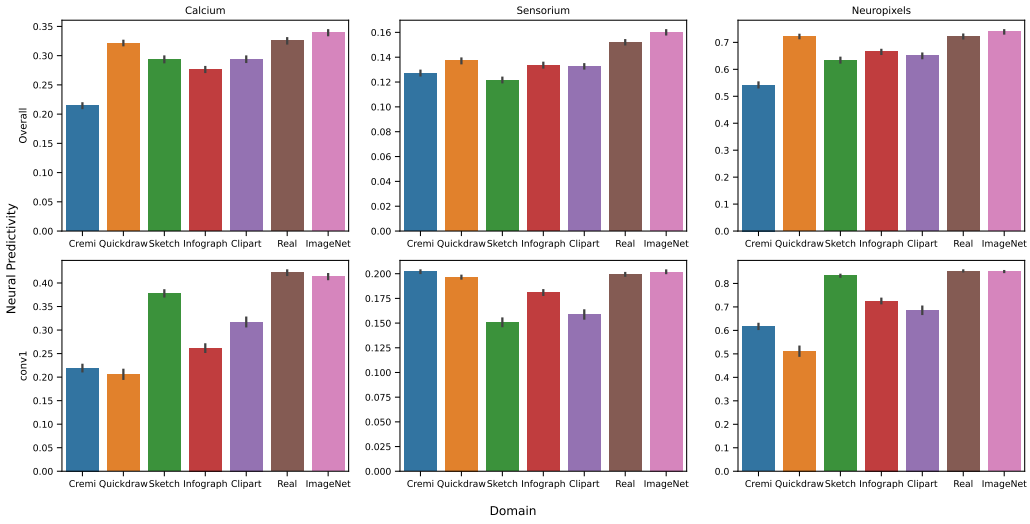

Figure 4: Comparison between the average predictivity over all convolutional layers(top row) and the neural predictivity of conv1 by itself(bottom row) and . We see that across the board conv1 has higher predictivity. We also see that the more natural domains have a higher level of predictivity