Naturally Computed Scale Invariance in the Residual Stream of ResNet18

Supplementary Material

1. Appendix

1.1. Feature Visualization Regularization

We use an identical regularization method as described in [?], except we use different levels of jitter dependent on the depth of the layer (we still use half the initial jitter value for the second jitter transform). This is done due to the unreliability high jitter has at producing a high activation on the center neuron once optimization has completed. We sacrificed some interpretability that jitter offers for more accurate activation values.

We omitted jitter for the entire 1.1 and 2.0 blocks, we used jitter=4 for 2.1 Pre and the entire 3.0 block, then used the default jitter=16 value for the remainder of the blocks.

1.2. Visualization grids for remaining scale invariant channels

We present the grids of all criteria-passing channels throughout ResNet18 in Fig. 1 (blocks 1.1, 2.0, and 2.1) and Fig. 2 (block 3.1).





2.0: channel 50



2.1: channel 13



36



Figure 1. Grids of maximally exciting images for all remaining scale invariant criteria-passing channels in blocks 1.1, 2.0, and 2.1. All channels are zero-indexed.



Figure 2. Grids of maximally exciting images for all remaining scale invariant criteria-passing channels in block 3.1. All channels are zero-indexed.