

## Update 8/24/21

In our initial explorations, we first trained image classifiers on various input modalities: grayscale, edges, superpixels, etc. We then ensembled various combinations of these models and evaluated against IID data and OOD image corruptions. Our preliminary experiments on small datasets showed improvements to corruption robustness by a significant margin over equivalently sized ensembles of standard models. However, these gains disappeared after scaling up to larger datasets.

As a result, we decided to put our work here on hold and pursue new directions. Our next project is **Better Visual Representations through Language Supervision**.