A. Experimental Details

A.1. Fine-tuning

We fine-tuned the MNIST, CIFAR-10, and Glow models at 10 noise levels σ^2 (see Section 3.4) for 50 epochs each on clusters of 4 1080Ti GPU's. This procedure converges rapidly, with no further decrease of the negative log-likelihood after the first 10 epochs. Although Glow models theoretically have full support, the noiseless pre-trained models assign vanishing probability to highly noisy images. In practice, this can cause invertibility assertion failures when fine-tuning directly from the noiseless model. To avoid this we took an iterative approach: first fine-tune the lowest noise level $\sigma = .01$ from the noiseless model, then fine-tune the $\sigma = .016$ model from the $\sigma = .01$ model, etc.

A.2. Scaling and Resources

Scaling Algorithm 1 to richer datasets is constrained primarily by the limited availability of strong, likelihood-based generative models for these datasets. For high resolution images, the running time of Algorithm 1 can also become substantial. Assuming the hyper-parameters T and L discussed in Section 3.4 remain valid at higher resolutions, the computational complexity of BASIS scales linearly with the cost of evaluating gradients of the model (albeit with a large multiplicative constant $T \times L$). Therefore, if a generative model is tractable to train, then it should also be tractable to use for BASIS separation.

In concrete detail, we observe that a batch of 50 BASIS separation results for MNIST or CIFAR-10 using NCSN takes < 5 minutes on a single 1080Ti GPU. Running BASIS with Glow is much slower. We observe that substantial time is spent loading and unloading the noisy models p_σ from memory (in contrast to NCSN, which uses a single noise-conditioned model). A batch of 50 BASIS separation results on MNIST or CIFAR-10 using Glow takes about 30 minutes on a 1080Ti. A batch of 9 BASIS separation result on LSUN using Glow takes 2-3 hours on a 1080Ti.

A.3. Visual Comparisons

When using class-agnostic priors, BASIS separation is symmetric in its output components. To facilitate visual comparisons between original images and separated components, we sort the BASIS separated components to minimize PSNR to the original images. This usually results in the separated components being visually paired with the most similar original components. But due to the deficiencies of PSNR as a comparative metric this is not always the case; the alert reader may have noticed that the yellow and silver car mixture in Figure 1 appears to have been displayed in reverse order. This happens because the separated yellow car component takes the light sky from the original silver car component, and the lightness of the sky dominates the PSNR metric.

For the LSUN separation results, where we use a church model for the first component and a bedroom model for the second, the symmetry is broken. For these results, components naturally sort themselves into church and bedroom components, which can be compared directly to the original images.

B. Intermediate Samples During the Annealing Process

Figure 9. Intermediate CIFAR-10 separation results taken at noise levels σ during the annealing process of BASIS separation.

C. MNIST Separation Results Under Different Models and Sampling Procedures

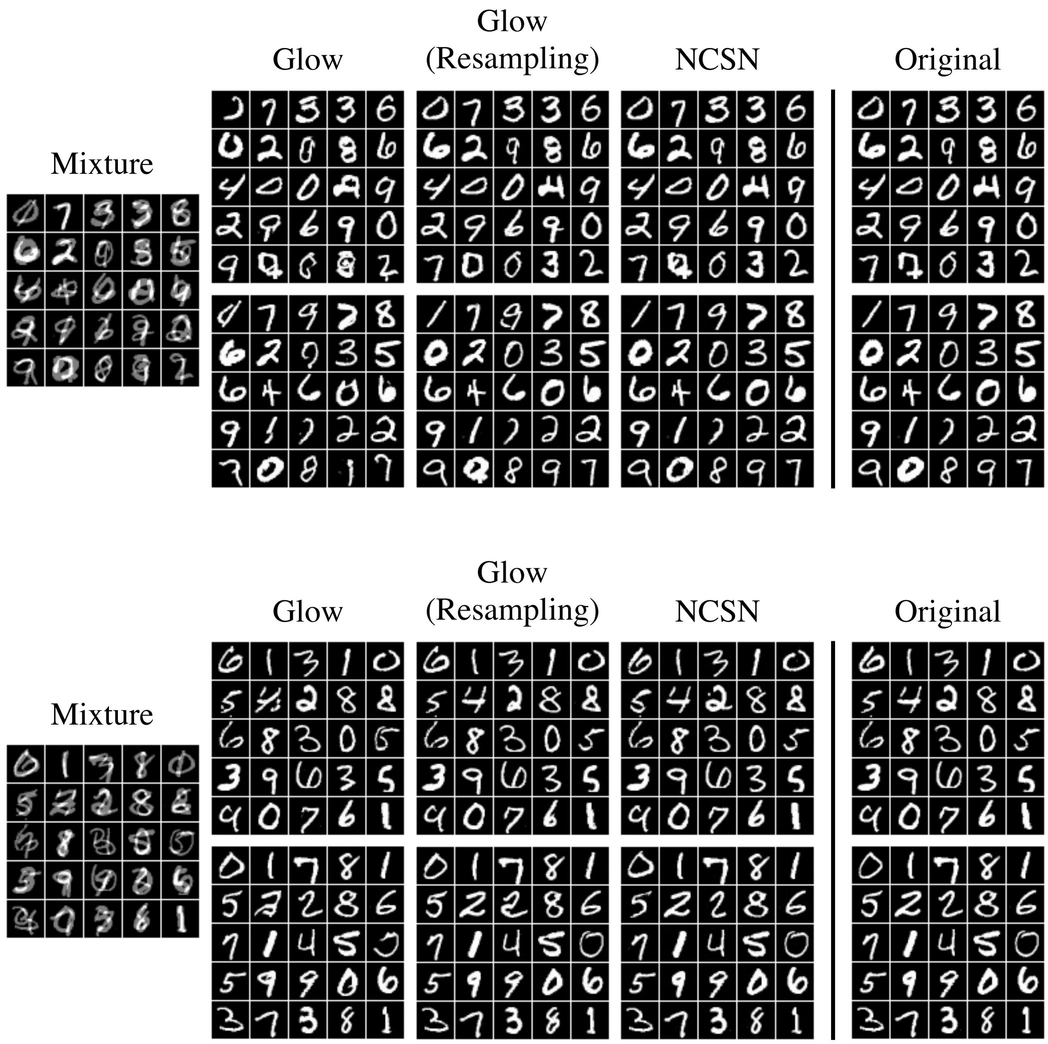

Figure 10. Uncurated class-agnostic separation results using: (1) samples from the posterior with Glow as a prior (2) an approximate MAP estimate using the maximum over 10 samples from the posterior with Glow as a prior (3) samples from the posterior with NCSN as a prior.

D. Extended CIFAR-10 Separation Results

D.1. NCSN Prior

Figure 11. Uncurated class-agnostic CIFAR-10 separation results using NCSN as a prior.

D.2. Glow Prior

Figure 12. Uncurated class-agnostic CIFAR-10 separation results using Glow as a prior.

E. Extended CIFAR-10 Colorization Results

E.1. NCSN Prior

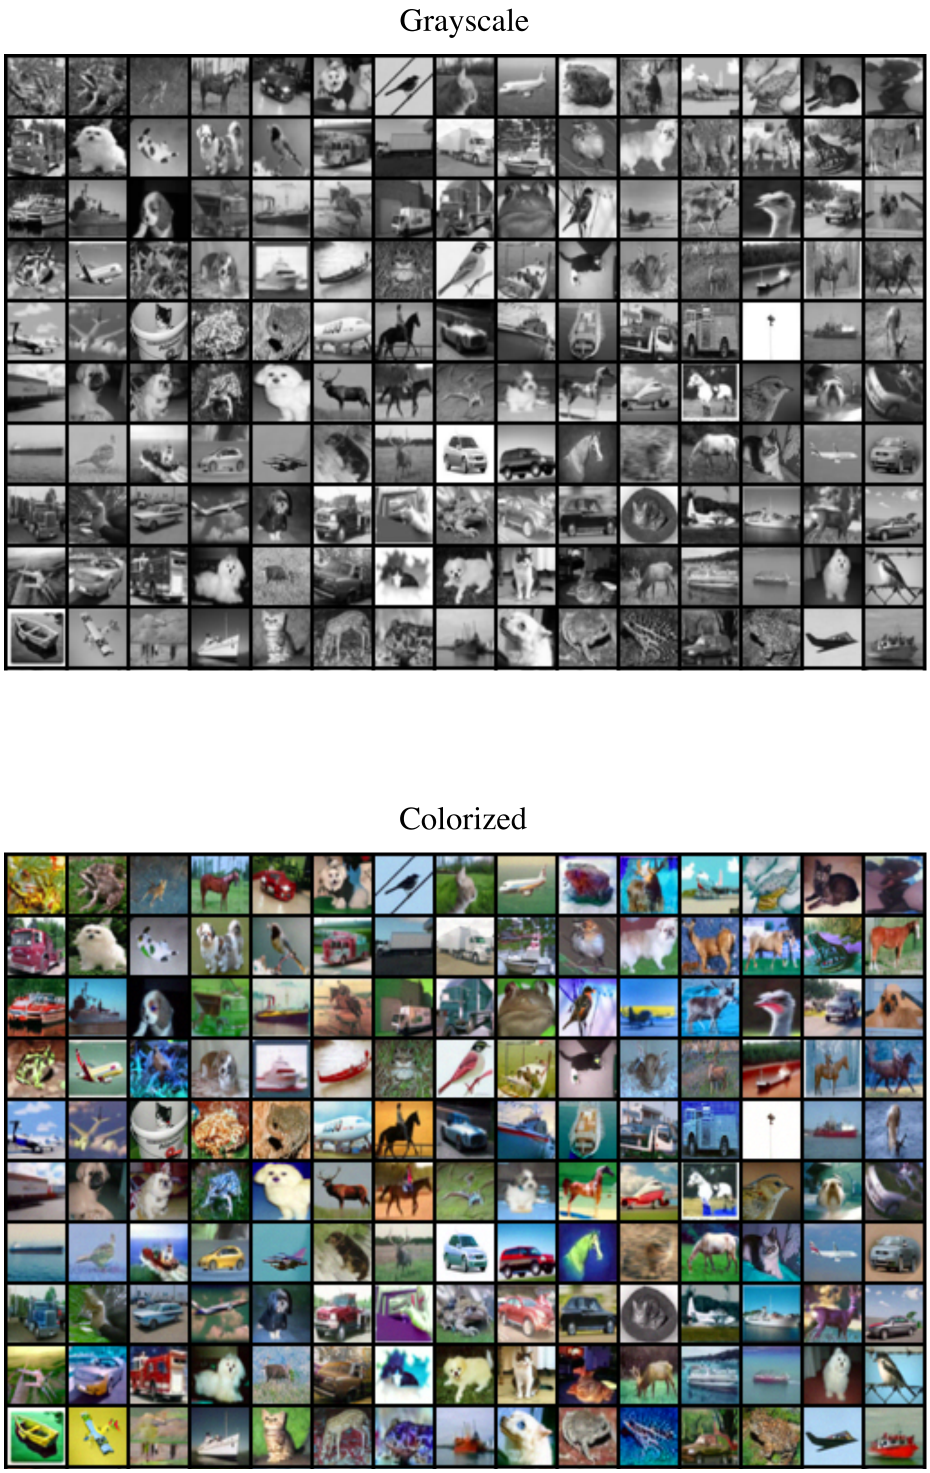

Figure 13. Uncurated CIFAR-10 colorization results using NCSN as a prior.

E.2. Glow Prior

Figure 14. Uncurated CIFAR-10 colorization results using Glow as a prior.

F. Extended LSUN Separation Results

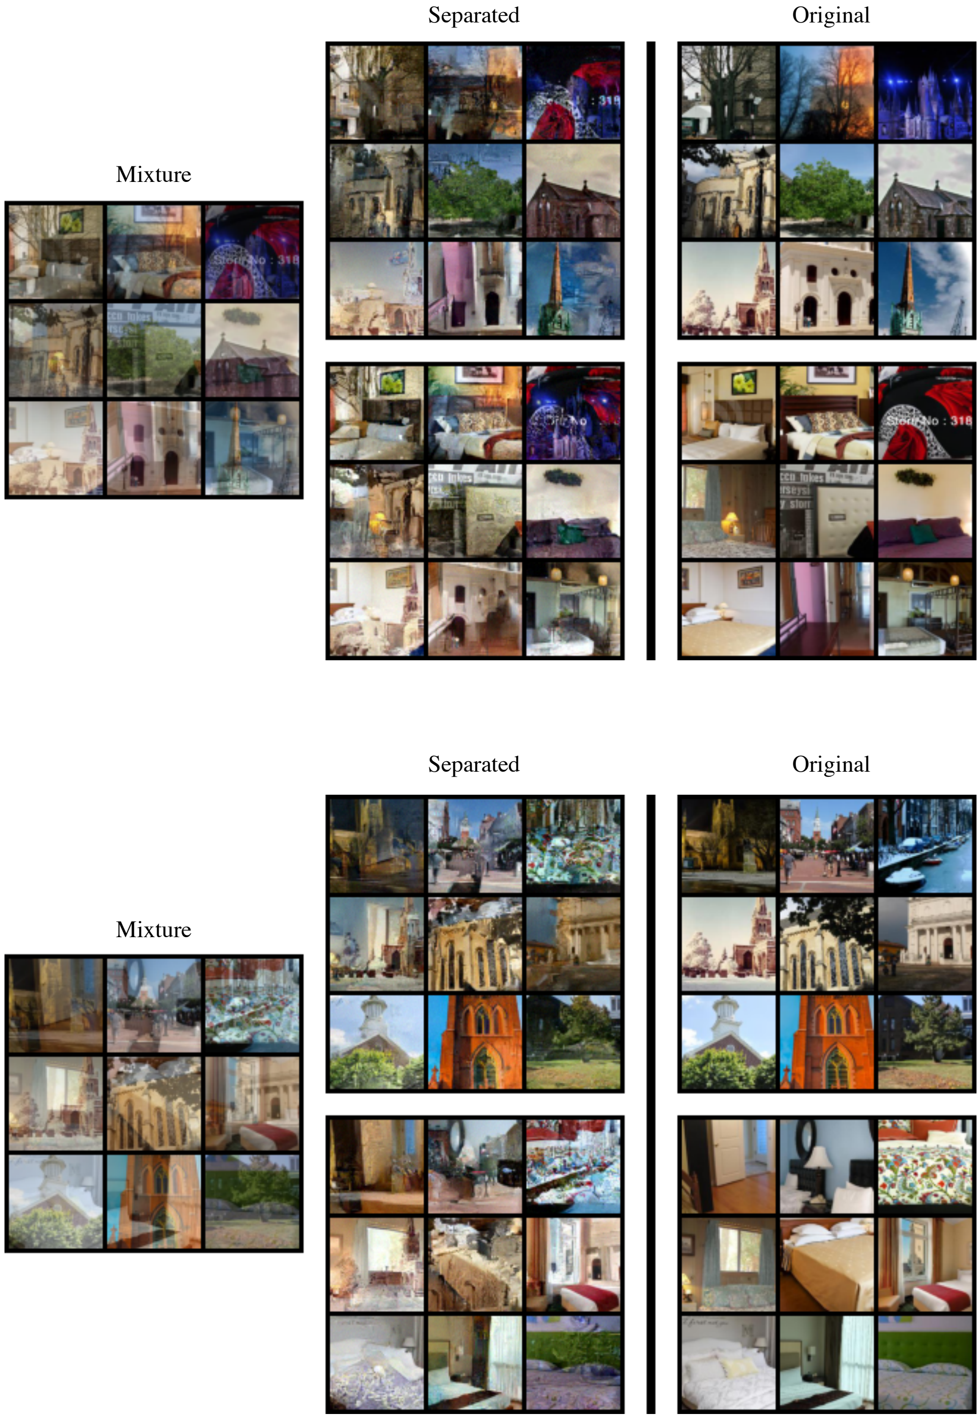

Figure 15. Uncurated church/bedroom LSUN separation results using Glow as a prior.