

Generator Surgery for Compressed Sensing

Jung Yeon Park^{1*}, Niklas Smedemark-Margulies^{1*}, Max Daniels¹, Rose Yu², Jan-Willem Van De Meent¹, Paul Hand¹

*Equal Contribution

Code: <https://github.com/nik-sm/generator-surgery>



*We show that **cutting initial layers** of a **generative signal prior** at test time **improves image recovery performance**.*

Problem Statement

- We study **compressed sensing of images** under generative priors.
 - Goal: recover images from noisy undersampled measurements.
 - Unique recovery is possible under prior assumptions on image.
 - CS with Generative Models (CSGM): assume images are sampled from a generative model (e.g. GAN, VAE, etc).

- Procedure:** project noisy measurements onto range of model.
 - Given measurement matrix A , noise η , generator $G(z)$:

$$y = Ax^* + \eta \quad A_{ij} \sim \mathcal{N}(0, 1/m) \quad \eta \sim \mathcal{N}(0, \sigma^2)$$

$$\hat{z} = \arg \min_z \|y - AG(z)\|_2^2 \quad \hat{x} = G(\hat{z})$$


- Intuition:** of all images consistent with measurements, use a generative model to decide which looks most natural.

- Key Problem:** typical generative models cannot represent all natural images, inducing **representation error**.

$$\text{Rep-Err}(x^*) = \min_z \|x^* - G(z)\|_2^2$$

We introduce **Generator Surgery (GS)**, which mitigates representation error and improves CSGM recovery performance.

Method




- Consider the generator as a composition of blocks. Each block may contain convolutions, upsample, batch norm, etc.

$$G_0(z_0) = [B_{d-1} \circ \dots \circ B_1 \circ B_0](z_0)$$

- Generator Surgery (GS):** Remove the first c blocks of the generator. Recover images by projection onto range of the new model.


$$G_c(z_c) = [B_{d-1} \circ \dots \circ B_c](z_c)$$

Main Results




- CelebA (in-training-distribution): GS substantially improves recovery performance for all considered architectures (DCGAN, BEGAN, and VAE).

- COCO (out-of-training-distribution): performance increase over no surgery is similar to CelebA images.



- Baselines: IAGAN, mGANprior, Deep Decoder (with similar number of optimized parameters).
- GS generally beats mGANprior and performs similarly to DD. IAGAN is highly overparameterized.


Original
mGANprior
DeepDecoder
Control
GS (ours)




- Qualitative comparison of GS to other baseline methods.
- Note: all models trained on CelebA.

Supplemental Results


- We select the cut index c based on 100 CelebA validation images. Recovery quality is roughly a concave function of c .



- The generator's trained weights contribute to recovery quality: reinitializing weights after surgery is harmful to performance.



- Samples from the uncut generator have zero representation error. Gradient descent converges with negligible error in this case.



Discussion

- We find promising results using test-time modifications for image recovery.
- Surprisingly, GS allows models trained on CelebA to recover images from COCO with high quality.
- Our method trades the model's generative sampling procedure for increased recovery quality.