Deep Compressed Sensing

Yan Wu, Mihaela Rosca, Tim Lillicrap

DeepMind

Introduction

Compressed sensing (CS) provides an elegant framework for recovering sparse signals from compressed measurements. It is highly flexible and data efficient, but its application has been restricted by the strong assumption of sparsity and costly optimisation process. Here we propose a novel framework that significantly improves both the performance and speed of signal recovery by jointly training a generator and the optimisation process for reconstruction via meta-learning. We explore training the measurements with different objectives, and derive a family of models based on minimising measurement errors. We show that Generative Adversarial Nets (GANs) can be viewed as a special case in this family of models. Borrowing insights from the CS perspective, we develop a novel way of stabilising GAN training using gradient information from the discriminator.

Contributions

- We incorporate generic neural networks with CS-style online optimisation
- The meta-learned reconstruction process is more accurate and orders of magnitudes faster compared with previous models (Bora et al. 2017)
- We improve GAN performance by online minimising generator loss as a measurement error.
- We derive a novel conditional generative models based on latent optimisation, which results in semantically meaningful latent spaces.

Generalisation

Deep Compressed Sensing

A real-valued m preserves image Euclidean distance

CS-GANs

A binary *m* that preserves image validity

$$\mathcal{L}_F = t(\boldsymbol{x}) \, \ln \left[D_\phi(\boldsymbol{x})\right] + (1 - t(\boldsymbol{x})) \, \ln \left[1 - D_\phi(\boldsymbol{x})\right]$$

Semi-supervised GANs

A categorical *m* that preserves image class

$$\mathcal{L}_F = \sum_{k=1}^{K+1} t^k(\boldsymbol{x}) \, \ln \left[C_\phi^k(\boldsymbol{x})\right]$$

Adversarially trained classifier

Framework

x: signal
z: latent code
F: measurement function
G: generator
m: measurement of signal
E: measurement error

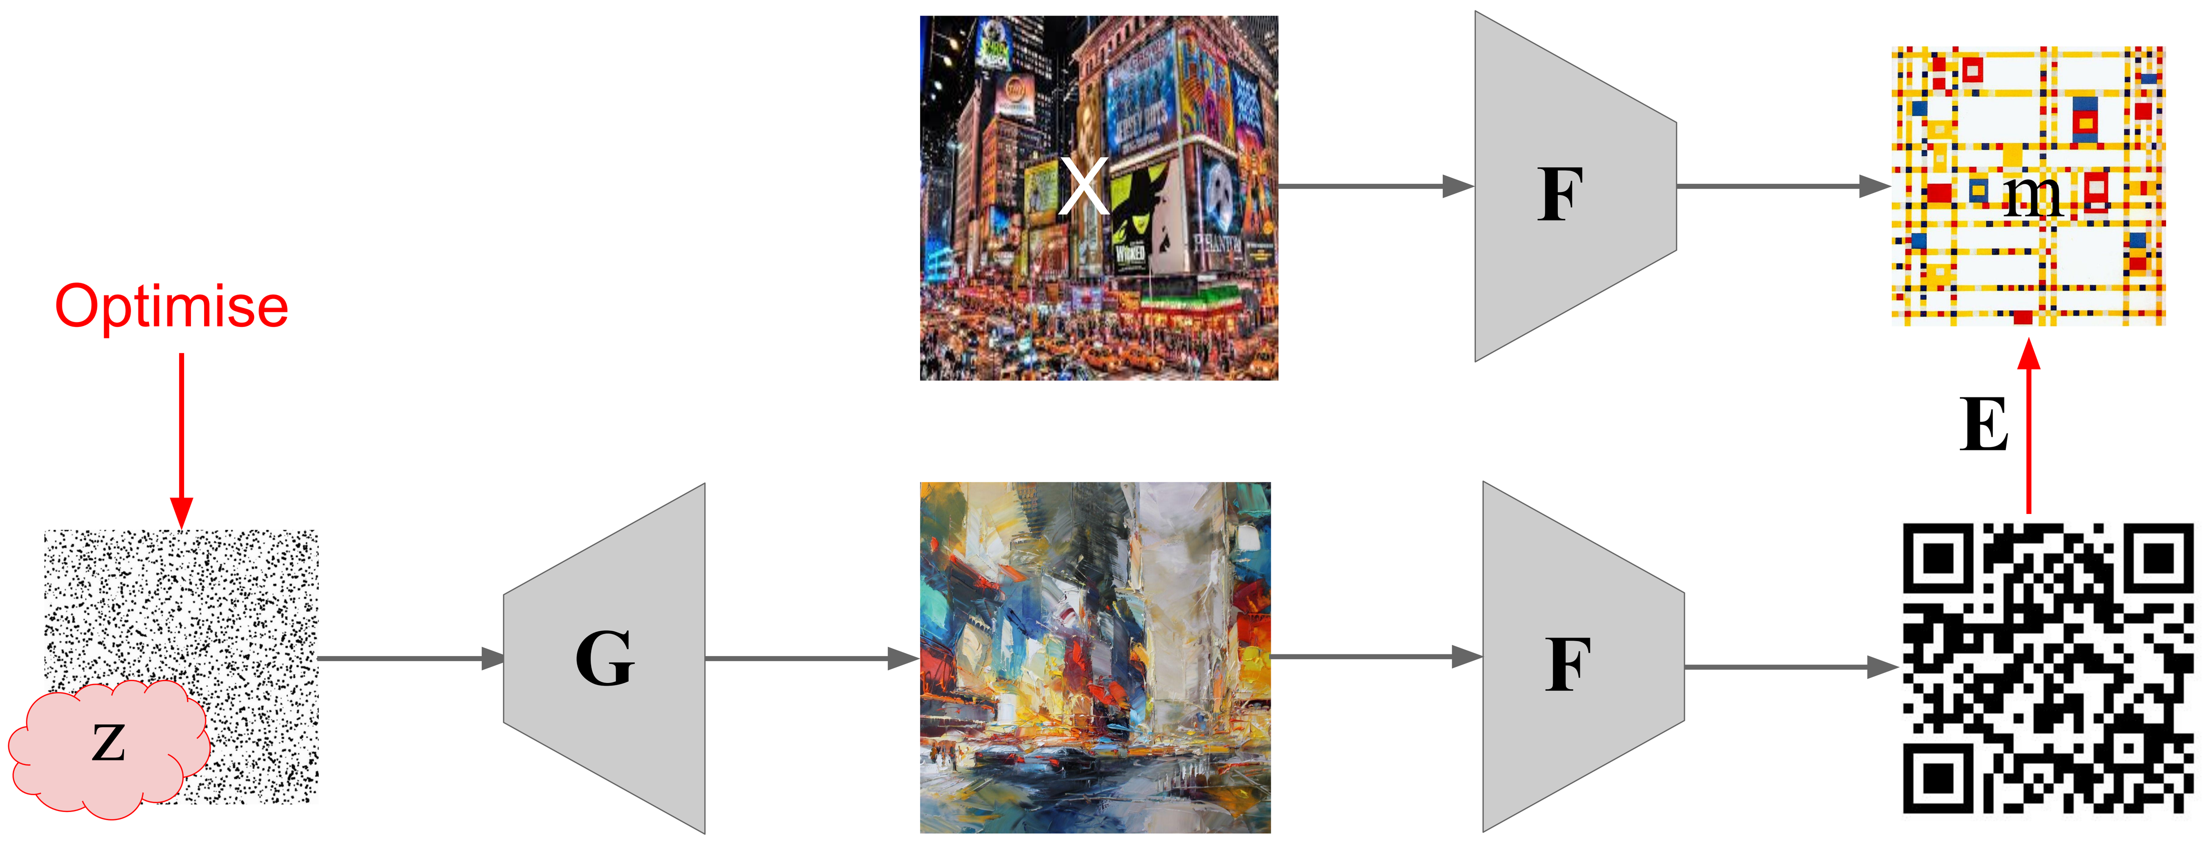

Formulation

Compressed Sensing

Linear random metric:

$$\boldsymbol{m} = \mathbf{F}\,\boldsymbol{x} + \eta$$

Directly optimising signal x (subject to sparsity constraint):

$$\hat{\boldsymbol{x}} = \arg\min_{\boldsymbol{x}} \|\boldsymbol{m} - \mathbf{F}\,\boldsymbol{x}\|_2^2$$

Restricted isometry property (RIP, i.e., distance preserving) comes with random matrix:

$$(1-\delta)\ \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2 \le \|\mathbf{F}\,(\boldsymbol{x}_1 - \boldsymbol{x}_2)\|_2^2 \le (1+\delta)\ \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2^2$$

Deep Compressed Sensing

Nonlinear parametrised metric and generator:

$$m \leftarrow F_\phi(x) \qquad \boldsymbol{x} = G_\theta(\boldsymbol{z})$$

Indirectly optimising latent (grad-descent on z):

$$\hat{\boldsymbol{z}} = \arg\min_{\boldsymbol{z}} E_\theta(\boldsymbol{m}, \boldsymbol{z}) \quad E_\theta(\boldsymbol{m}, \boldsymbol{z}) = \|\boldsymbol{m} - F_\phi\left(G_\theta(\boldsymbol{z})\right)\|_2^2$$

$$\hat{\boldsymbol{z}} \leftarrow \hat{\boldsymbol{z}} - \alpha \left.\frac{\partial E_\theta(\boldsymbol{m}, \boldsymbol{z})}{\partial \boldsymbol{z}}\right|_{\boldsymbol{z}=\hat{\boldsymbol{z}}}$$

Train the metric for RIP:

$$\mathcal{L}_F = \mathbb{E}_{\boldsymbol{x}_1, \boldsymbol{x}_2}\left[\left(\|F_\phi(\boldsymbol{x}_1 - \boldsymbol{x}_2)\|_2 - \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2\right)^2\right]$$

Train the generator to minimise measurement error:

$$\min_\theta \mathcal{L}_G, \ \text{for}\ \mathcal{L}_G = \mathbb{E}_{\boldsymbol{x}_i \sim p_{\text{data}}(\boldsymbol{x})}\left[E_\theta(\boldsymbol{m}_i, \hat{\boldsymbol{z}}_i)\right]$$

Unlike autoencoders, no reconstruction error is involved in reconstruction (likelihood-free)

Compressed Sensing Results

Model	10	25 measurements	steps
Baseline	54.8	17.2	10×1000
Linear	10.8 ± 3.8	6.9 ± 2.7	3
NN	12.5 ± 2.2	10.2 ± 1.7	3
Linear(L)	6.5 ± 2.1	$4. \pm 1.4$	3
NN(L)	$\mathbf{5.3 \pm 1.9}$	$\mathbf{3.4 \pm 1.2}$	3

Model	20	50 measurements	steps
Baseline	156.8	82.3	2×500
Linear	34.7 ± 7.9	27.1 ± 6.1	3
NN	46.1 ± 8.9	41.2 ± 8.3	3
Linear(L)	26.2 ± 5.9	20.5 ± 4.3	3
NN(L)	$\mathbf{23.4 \pm 5.8}$	$\mathbf{18.5 \pm 4.3}$	3

DC-GAN CIFAR

Hyper-parameter sweep over 144 configs: generator learning rates {1 × 10−4 , 2 × 10−4 , 3 × 10−4 }, discriminator learning rates {1×10−4, 2×10−4, 3×10−4}, latent variable sizes {100, 200}, mini-batch sizes {32, 64}, and 2 instances for each combination to account for the effect of random seeds.

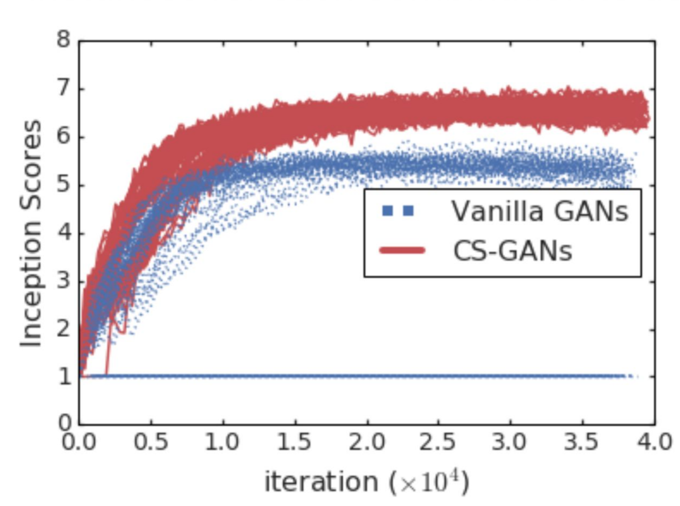

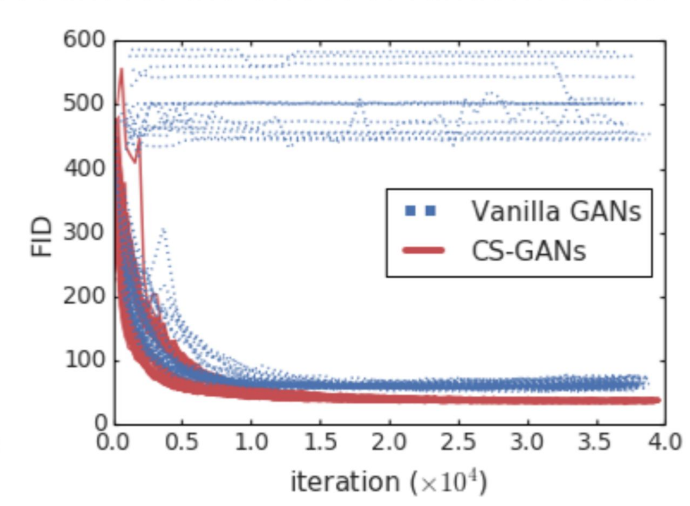

	SN-GAN	SN-GAN (ours)	CS+SN-GAN
IS	7.42 ± 0.08	7.34 ± 0.07	$\mathbf{7.80 \pm 0.05}$
FID	29.3	29.53 ± 0.36	$\mathbf{23.13 \pm 0.50}$

Compressed Sensing Examples

random linear

learned linear

learned non-linear

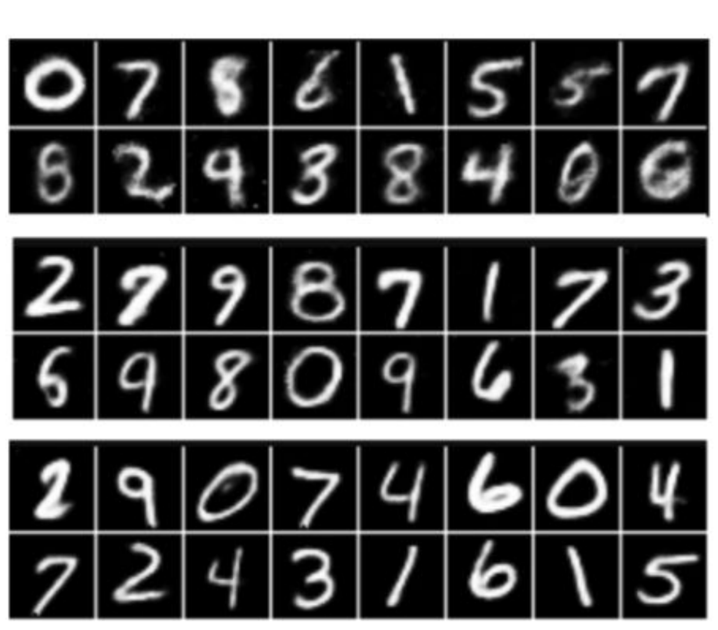

GAN Toy Examples

0 iteration | 3 iteration | 5 iterations

Generative Classifier

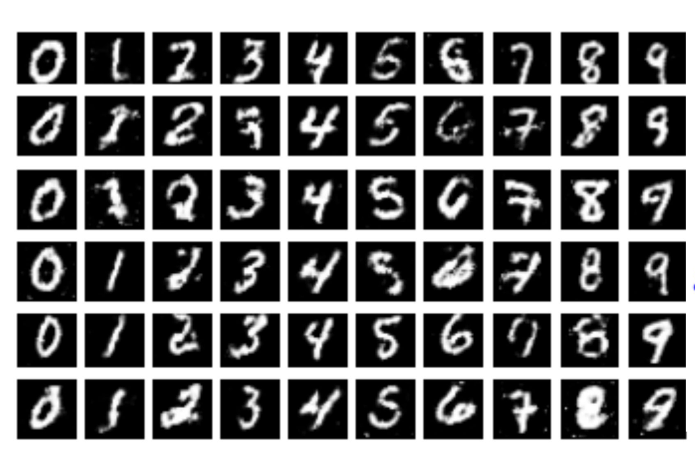

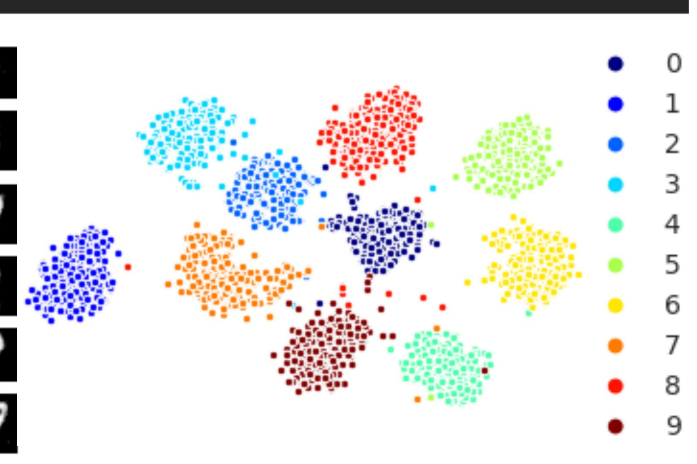

All the conditioning information is in the latent variable.

Bora, A., Jalal, A., Price, E., and Dimakis, A. G. Compressed sensing using generative models.