Deep Compressed Sensing

Yan Wu, Mihaela Rosca, Tim Lillicrap

DeepMind

Compressed Sensing

A Brief Review

An underdetermined problem: $y = Mx + \eta$

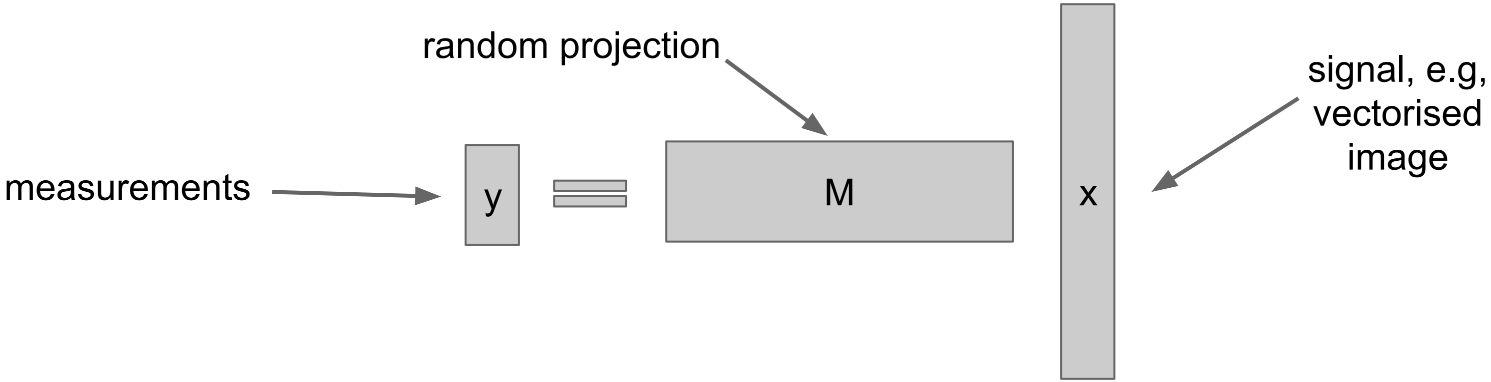

Reconstruction of x is possible when the signal is *sparse* (Candes, Donoho, Romberg, Tao, 2006~):

$$x^* = \text{argmin}_x \, \|y - Mx\|_2^2 + \lambda \, \|x\|_1$$

Why CS works

Restricted Isometry Property (RIP/REC)

- The projection M preserves the Euclidean distance between *k-sparse* signals
- Many random matrices have the RIP with high probability

$$(1-\delta)\left\|x_1-x_2\right\|_2^2 \leq \left\|Mx_1-Mx_2\right\|_2^2 \leq (1+\delta)\left\|x_1-x_2\right\|_2^2$$

Examples: MRI reconstruction, the single pixel camera

RIP can be a trained property

Baseline:

Compressed Sensing using Generative Models

(Bora et al. 2017, almost the same as our model except using separately trained generators)

MNIST

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)	5.3 ± 1.9	3.4 ± 1.2	3

CelebA

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)	23.4 ± 5.8	18.5 ± 4.3	3

Improve GANs by online optimisation

CIFAR, Deep Convolutional GAN

DCGANs, sweeps over 144 hyper-parameters:

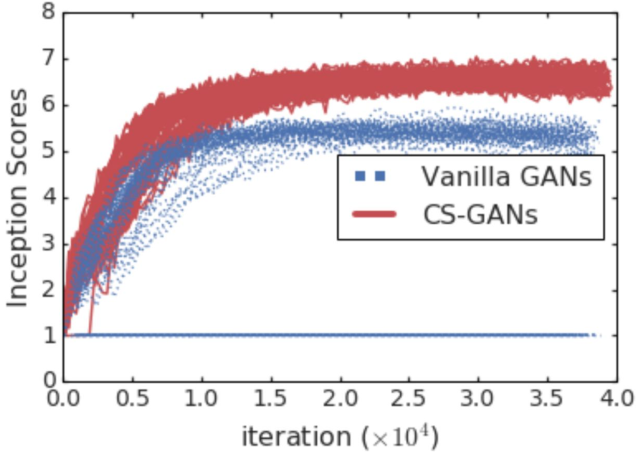

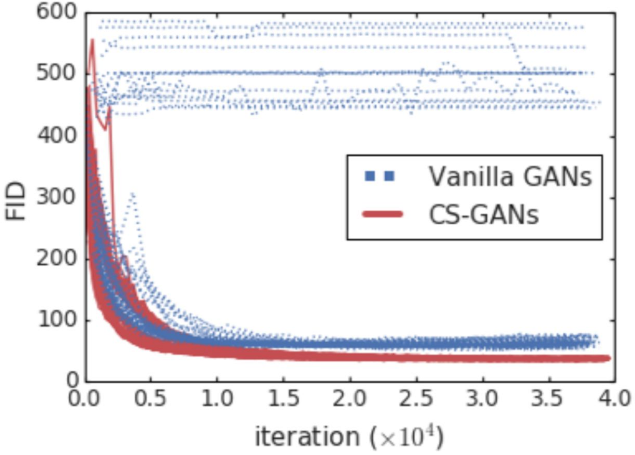

Spectral-normalised GANs

	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	$\mathbf{23.13 \pm 0.50}$

Summary

- A framework based on minimising measurement errors
- Bring RIP to neural networks via training
- Improve GANs: novel use of the discriminator as a measurement function
- A new semi-supervised model

Model	Metric Property
Compressed Sensing	RIP from random projection
Deep Compressed Sensing	Trained RIP
Semi-supervised GANs	Multi-Class Classifier
CS-GANs	Binary Classifier
…	…

Poster #24, Pacific Ballroom