


Generative models as data-driven priors: how to learn them efficiently?


Vincent Schellekens & Laurent Jacques

UCLouvain




Motivation: inverse problems

$$\boldsymbol{y} = \Phi(\boldsymbol{x}^*) + \boldsymbol{w} \quad \text{Forward model}$$




Motivation: inverse problems

$$\mathbf{y} = \Phi(\mathbf{x}^*) + \mathbf{w} \quad \text{Forward model}$$



Motivation: inverse problems

$$\mathbf{y} = \Phi(\mathbf{x}^*) + \mathbf{w} \quad \text{Forward model}$$




Data-fidelity loss

- Euclidean $l(\mathbf{x}; \mathbf{y}) = \|\mathbf{y} - \Phi(\mathbf{x})\|_2^2$

- ...

Motivation: inverse problems

$$\mathbf{y} = \Phi(\mathbf{x}^*) + \mathbf{w} \quad \text{Forward model}$$



Data-fidelity loss


- Euclidean $l(\mathbf{x}; \mathbf{y}) = \|\mathbf{y} - \Phi(\mathbf{x})\|_2^2$
- ...

Regularization (= prior)

- Tikhonov $r(\mathbf{x}) = \|\mathbf{x}\|_2^2$
- TV-norm $r(\mathbf{x}) = \|\mathbf{x}\|_{TV}$
- Sparsity (in wavelets) $r(\mathbf{x}) = \|\Psi \mathbf{x}\|_1$
- Sparsity (in dictionary) $r(\mathbf{x}) = \|D \mathbf{x}\|_1$
- Generative priors

Motivation: inverse problems

$$\mathbf{y} = \Phi(\mathbf{x}^*) + \mathbf{w} \quad \text{Forward model}$$



Data-fidelity loss

- Euclidean $l(\mathbf{x}; \mathbf{y}) = \|\mathbf{y} - \Phi(\mathbf{x})\|_2^2$
- ...


Regularization (= prior)

- Tikhonov $r(\mathbf{x}) = \|\mathbf{x}\|_2^2$
- TV-norm $r(\mathbf{x}) = \|\mathbf{x}\|_{TV}$
- Sparsity (in wavelets) $r(\mathbf{x}) = \|\Psi \mathbf{x}\|_1$
- Sparsity (in dictionary) $r(\mathbf{x}) = \|D \mathbf{x}\|_1$
- Generative priors

From model-based
to data-driven priors!

Motivation: inverse problems

$$\mathbf{y} = \Phi(\mathbf{x}^*) + \mathbf{w} \quad \text{Forward model}$$



Data-fidelity loss

- Euclidean $l(\mathbf{x}; \mathbf{y}) = \|\mathbf{y} - \Phi(\mathbf{x})\|_2^2$
- ...

Regularization (= prior)


- Tikhonov $r(\mathbf{x}) = \|\mathbf{x}\|_2^2$
- TV-norm $r(\mathbf{x}) = \|\mathbf{x}\|_{TV}$
- Sparsity (in wavelets) $r(\mathbf{x}) = \|\Psi \mathbf{x}\|_1$
- Sparsity (in dictionary) $r(\mathbf{x}) = \|D \mathbf{x}\|_1$
- Generative priors

From model-based
to data-driven priors!

Idea: approximate true
prior distribution from
learning examples

Generative priors in inverse problems (1)

1) Learn a generative network \mathcal{G}_θ

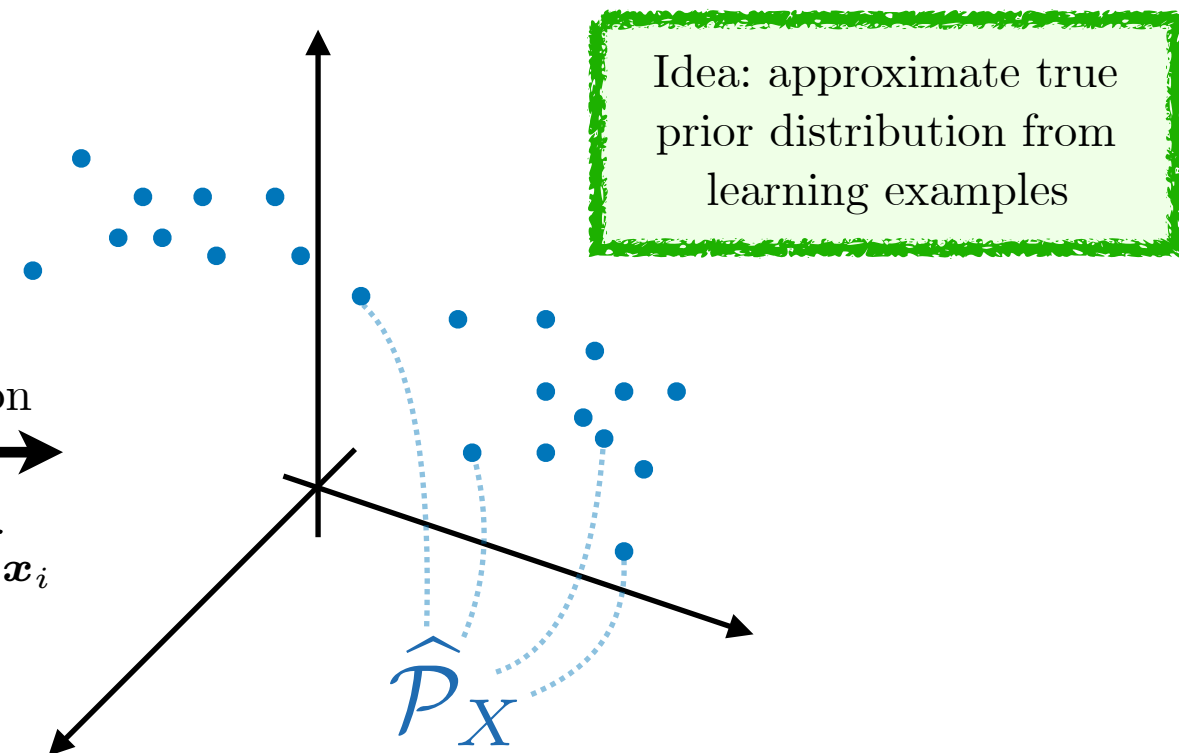


Dataset

We have **samples** (signals)


Empirical distribution

$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{\mathbf{x}_i \in X} \delta_{\mathbf{x}_i}$$



Generative priors in inverse problems (1)


1) Learn a generative network \mathcal{G}_θ



Dataset

Empirical distribution


$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{\mathbf{x}_i \in X} \delta_{\mathbf{x}_i}$$



We have **samples** (signals) from a high-dimensional **distribution**...

Generative priors in inverse problems (1)


1) Learn a generative network \mathcal{G}_θ



Dataset

Empirical distribution



$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{x_i \in X} \delta_{x_i}$$



We have **samples** (signals) from a high-dimensional **distribution**...

... and a way to generate “artificial” samples...

$$\mathcal{G}_\theta(z) = \rho(\Theta_L \cdot \rho(\Theta_{L-1} \cdots \rho(\Theta_2 \cdot \rho(\Theta_1 \cdot z))))$$




Generative network

Latent space

Random “noise”

Generative priors in inverse problems (1)


1) Learn a generative network \mathcal{G}_θ



We have **samples** (signals) from a high-dimensional **distribution**...


... and a way to generate “artificial” samples...

$$\mathcal{G}_\theta(\mathbf{z}) = \rho(\Theta_L \cdot \rho(\Theta_{L-1} \cdots \rho(\Theta_2 \cdot \rho(\Theta_1 \cdot \mathbf{z}))))$$



Generative priors in inverse problems (1)


1) Learn a generative network \mathcal{G}_θ



Dataset

Empirical distribution



$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{x_i \in X} \delta_{x_i}$$




We have **samples** (signals) from a high-dimensional **distribution**...

... and a way to generate “artificial” samples...

$$\mathcal{G}_\theta(z) = \rho(\Theta_L \cdot \rho(\Theta_{L-1} \cdots \rho(\Theta_2 \cdot \rho(\Theta_1 \cdot z))))$$



Generative network



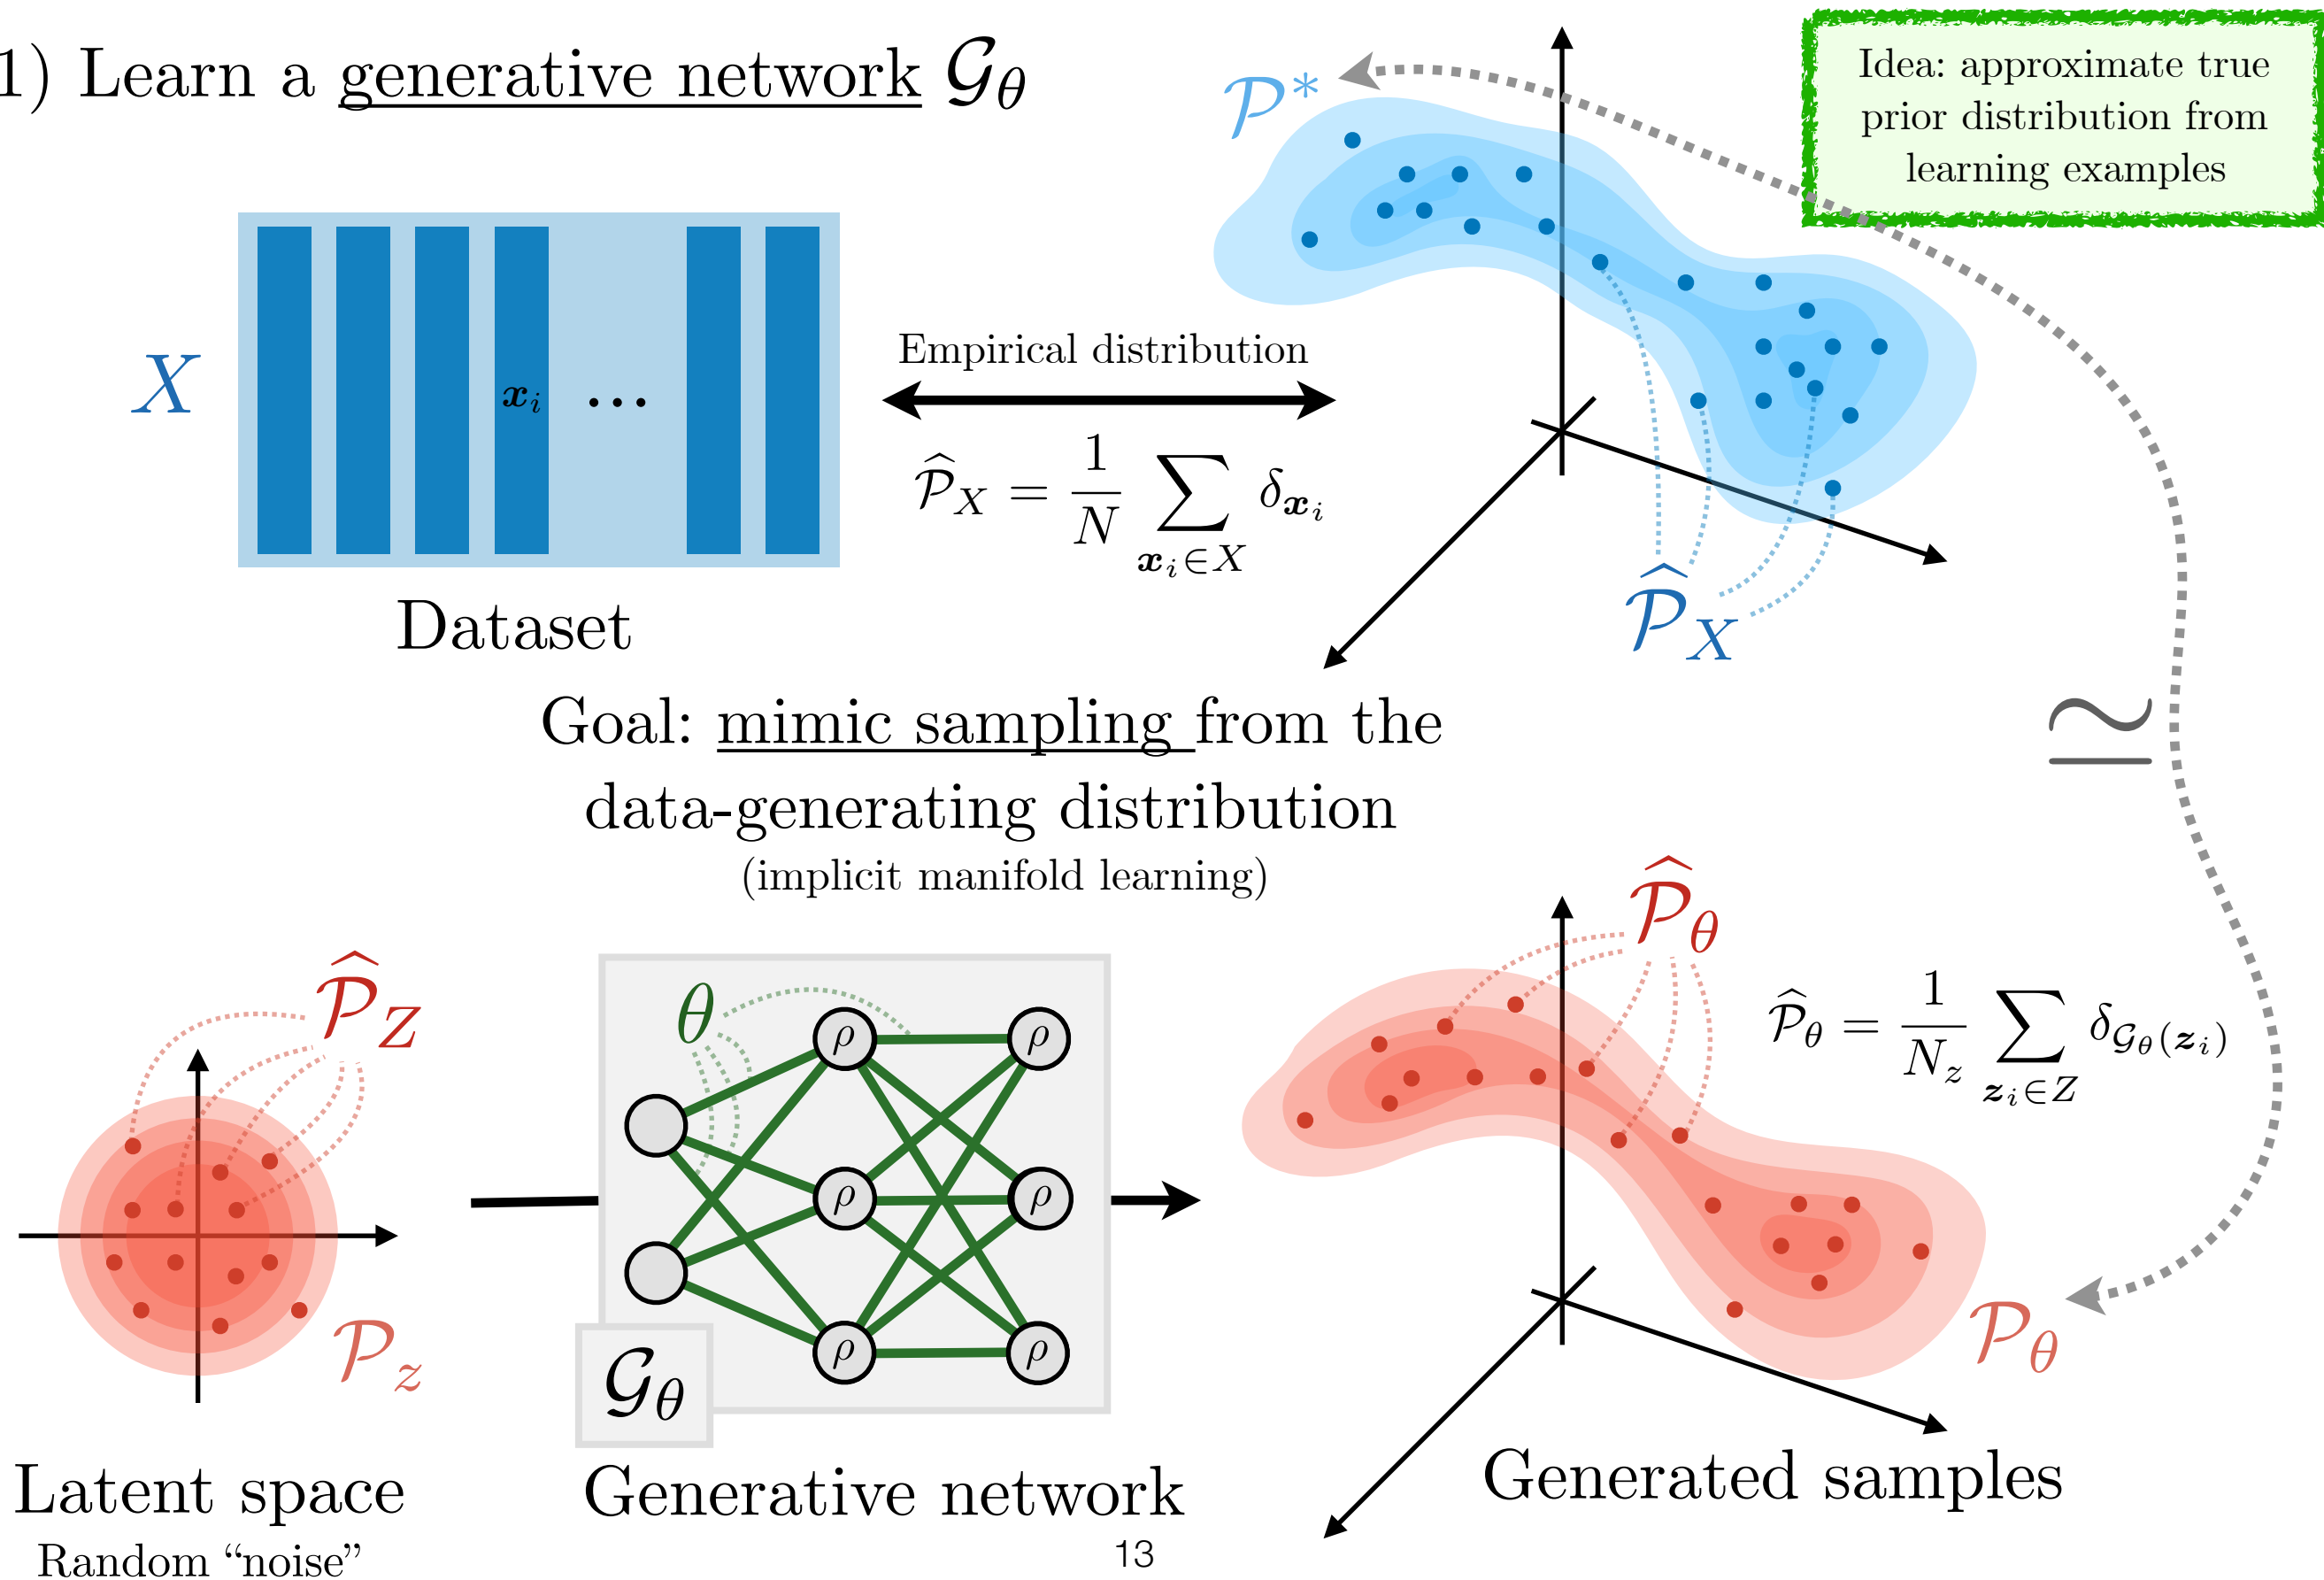
Generated samples

Latent space

Random “noise”

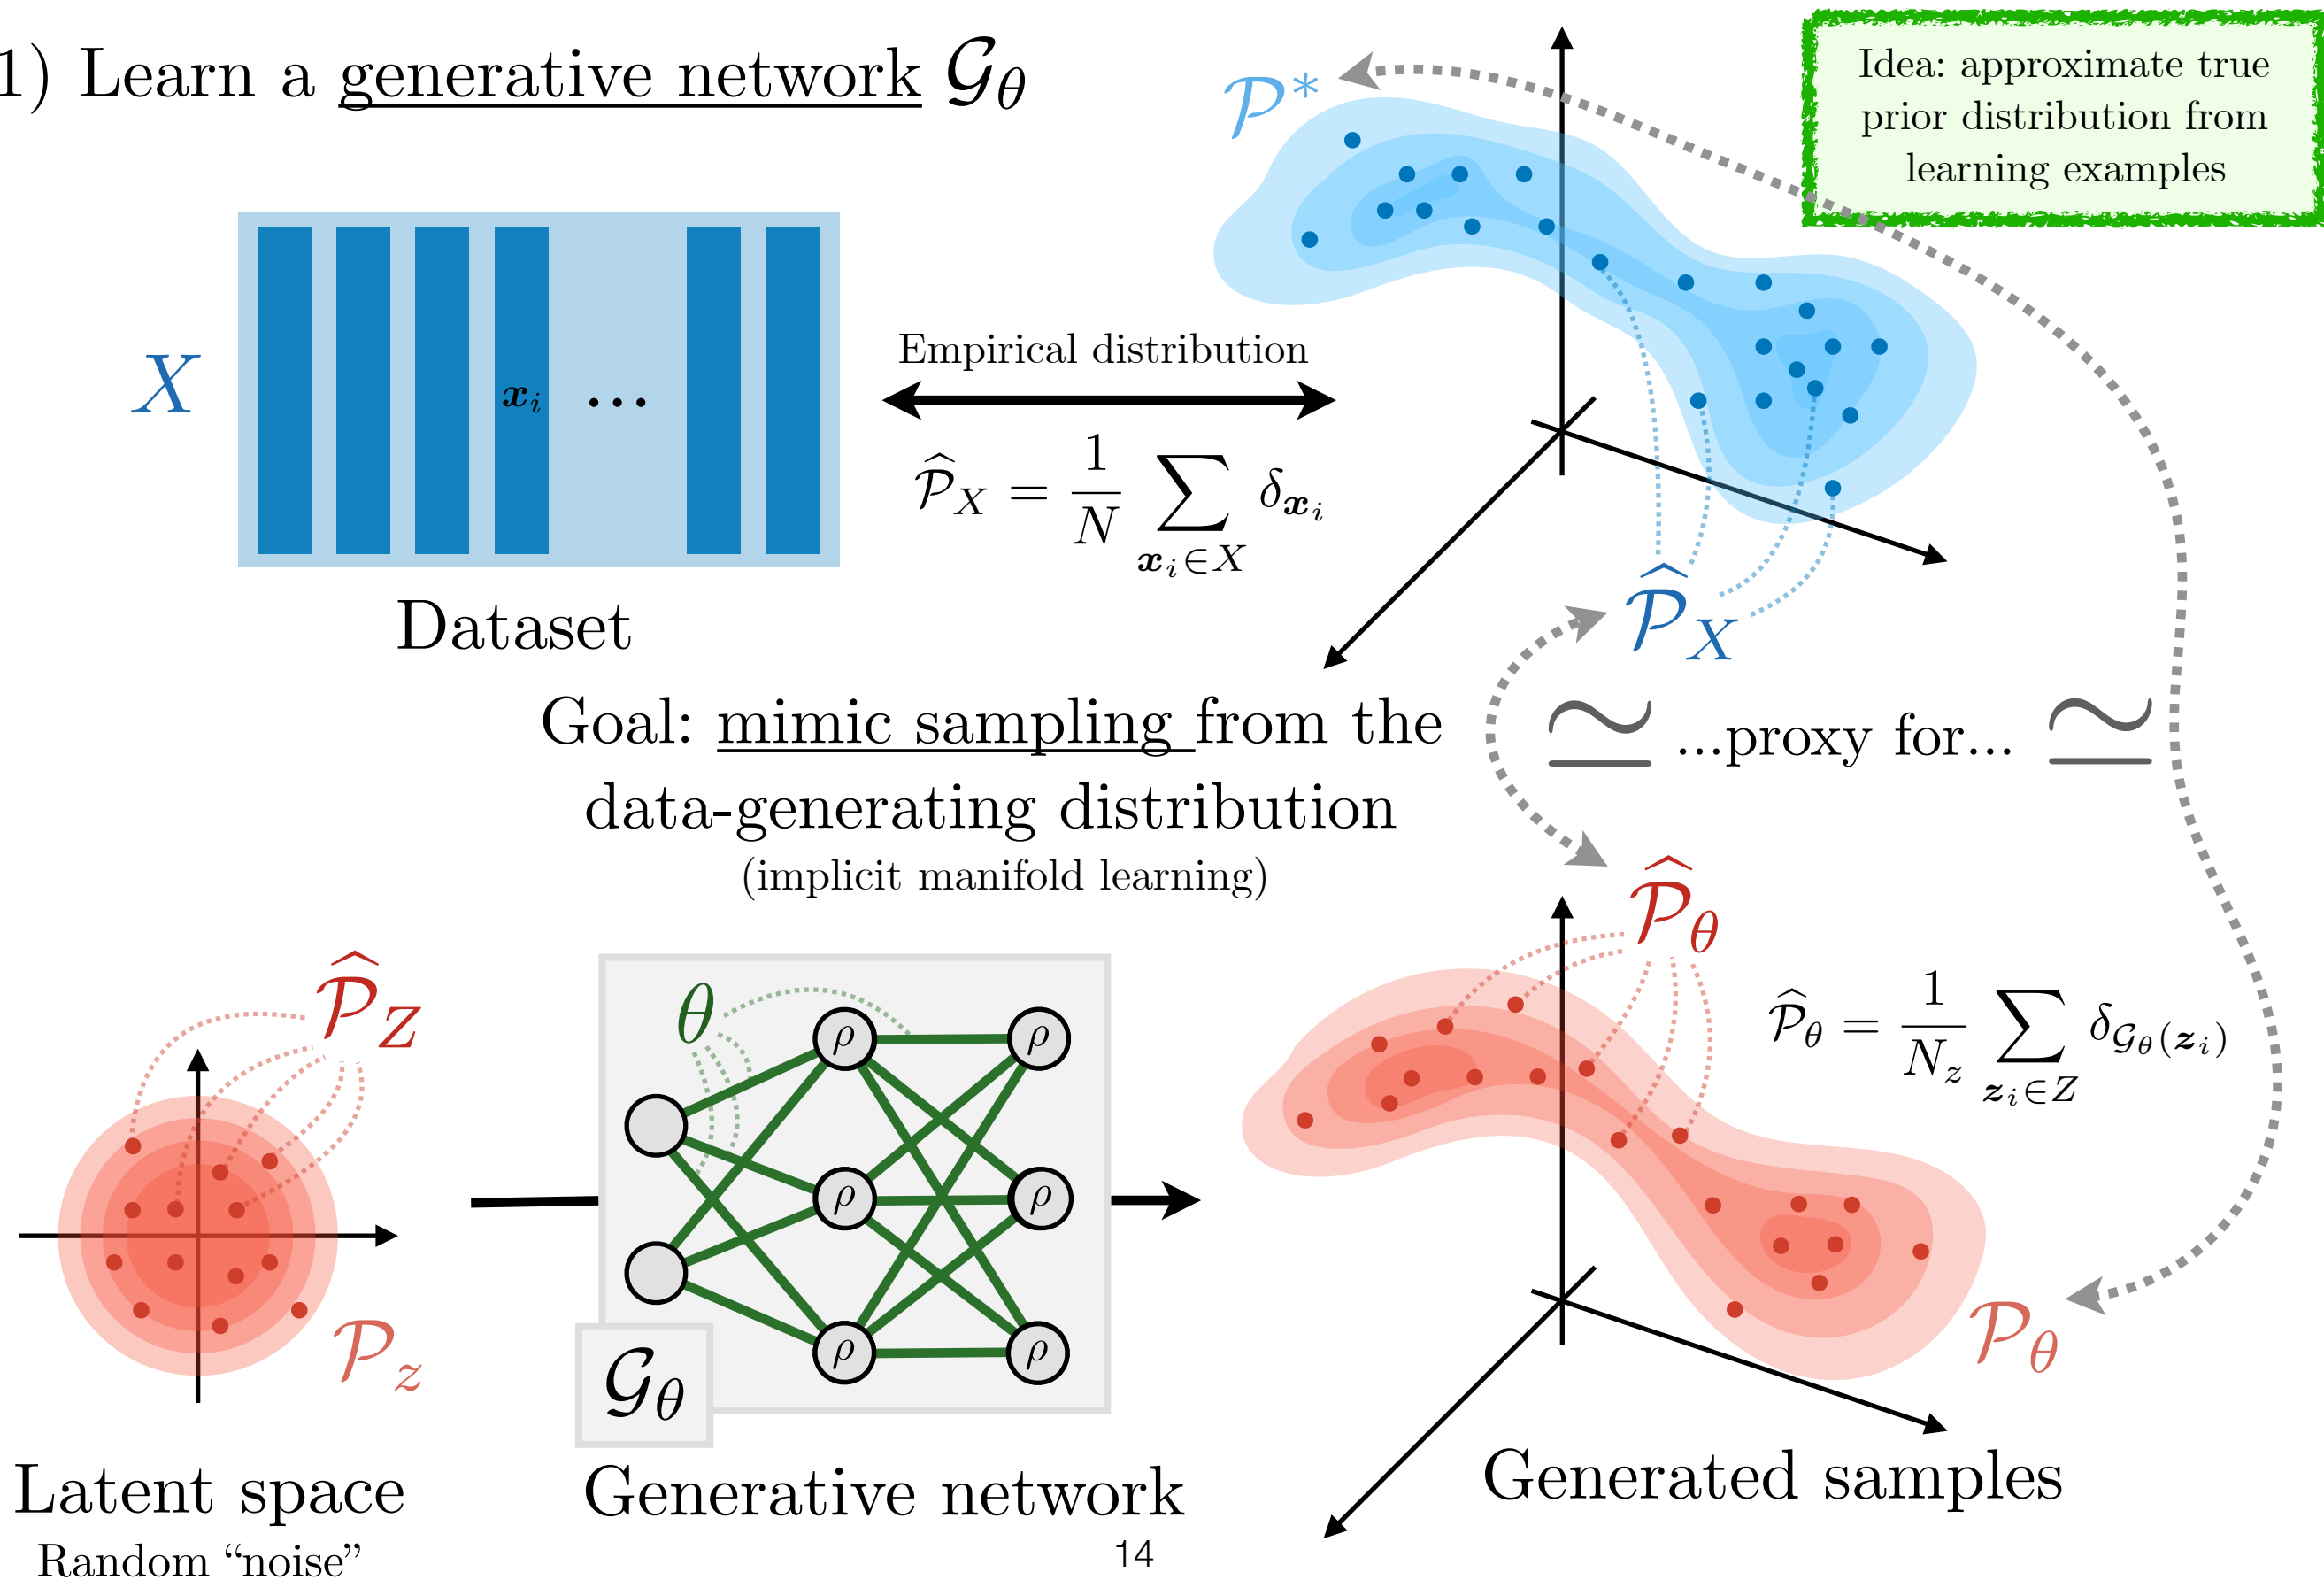
Generative priors in inverse problems (1)

1) Learn a generative network \mathcal{G}_θ



Generative priors in inverse problems (1)

1) Learn a generative network \mathcal{G}_θ



Generative priors in inverse problems (2)

- 1) Learn a generative network \mathcal{G}_θ
- 2) Use the generative network in the inverse problem


Generative priors in inverse problems (2)

1) Learn a generative network \mathcal{G}_θ

2) Use the generative network in the inverse problem

e.g. reconstructed signal in the range of the network

$$\hat{\mathbf{x}} = \mathcal{G}_\theta(\hat{\mathbf{z}}) \quad \text{with} \quad \hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \|\mathbf{y} - \Phi(\mathcal{G}_\theta(\mathbf{z}))\|_2^2$$



Generative priors in inverse problems

Some examples...

Compressed Sensing using Generative Models

Ashish Bora* Ajil Jalal† Eric Price‡ Alexandros G. Dimakis§

Abstract

The goal of compressed sensing is to estimate a vector from an underdetermined system of noisy linear measurements, by making use of prior knowledge on the structure of vectors in the relevant domain. For almost all results in this literature, the structure is represented by sparsity in a well-chosen basis. We show how to achieve guarantees similar to those achieved by sparsity, but using generative models. Our algorithm can use 5-10x fewer measurements than Lasso, and achieves near the rate of compressed sensing, $O(k \log L)$.




Figure 3: Reconstruction results on celebA with $m = 500$ measurements (of $n = 12288$ dimensional vector). We show original images (top row), and reconstructions by Lasso with DCT basis (second row), Lasso with wavelet basis (third row), and our algorithm (last row).

Generative priors in inverse problems

Some examples...

Deep Generative Adversarial Networks for Compressed Sensing (GANCS) Automates MRI

*Morteza Mardani^{1,3}, Enhao Gong¹, Joseph Y. Cheng^{1,2}, Shreyas Vasanawala²,
Greg Zaharchuk², Marcus Alley², Neil Thakur², Song Han⁴, William Dally⁴,
John M. Pauly¹, and Lei Xing^{1,3*}*

Magnet
inverse
tially tra
compres
To cope
benefits
manifold
mixture




Fig. 2: Representative coronal (1st row) and axial (3rd row) images for a test patient retrieved by ZF (1st), CS-WV (2nd), ℓ_2 -net (3th), ℓ_1 -net (4th), GAN (5th), GANCS (6th), and gold-standard (7th).

Generative priors in inverse problems

Some examples...

SUNLayer: Stable denoising with generative networks

Dustin G. Mixon* Soledad Villar†

Abstract

It has been experimentally established that deep neural networks are good for real world data. It has also been established that such generative models can be used to solve inverse problems like compressed sensing and super resolution. In this paper, we propose a theoretical setting that uses the properties of the activation functions will allow signal denoising.

1 Introduction

Deep neural networks, in particular generative adversarial networks, have been used to produce generative models for real world data that can be used for natural images (see for instance [Nguyen et al., 2016]). They can also efficiently solve classical inverse problems in signal processing, such as compressed sensing ([Bora et al., 2017]). The latter numerically solve the compressed sensing problem with ten times fewer measurements than required. Follow-up work by [Hand and Voroninski, 2017] (reducing empirical risk minimization) in the compressed sensing task by using a network with random weights and ReLU activation functions.




Figure 1: Denoising with generative priors

(**First line**) Digits from the MNIST test set ([LeCun, 1998]). (**Second line**) random noise is added to the digits. (**Third line**) Denoising of images by shrinkage in wavelet domain ([Donoho and Johnstone, 1994]). (**Fourth line**) Denoising by minimizing total variation ([Rudin et al., 1992]). (**Fifth line**) We train a GAN using the training set of MNIST to obtain a generative model G . We denoise by finding the closest element in the image of G using stochastic gradient descent.

Generative priors in inverse problems

Some examples...

Adversarial Regularizers in Inverse Problems




Sebastian Lunz
DAMTP
University of Cambridge
Cambridge CB3 0
lunz@math.cam.ac.uk

Ozan Öktem
Department of Mathematics
KIT - Karlsruhe Institute of Technology

Carola-Bibiane Schönlieb
DAMTP
University of Cambridge

Inverse Problems using purely data-driven regularization functions trained, the variational algorithm

(a) High noise			(b) Low noise		
Method	PSNR (dB)	SSIM	Method	PSNR (dB)	SSIM
MODEL-BASED			MODEL-BASED		
Filtered Backprojection	14.9	.227	Filtered Backprojection	23.3	.604
Total Variation [18]	27.7	.890	Total Variation [18]	30.0	.924
SUPERVISED			SUPERVISED		
Post-Processing [15]	31.2	.936	Post-Processing [15]	33.6	.955
RED [21]	29.9	.904	RED [21]	32.8	.947
UNSUPERVISED			UNSUPERVISED		
Adversarial Reg. (ours)	30.5	.927	Adversarial Reg. (ours)	32.5	.946



(a) Ground Truth

(b) FBP

(c) TV

(d) Post-Processing (e) Adversarial Reg.

Figure 2: Reconstruction from simulated CT measurements on the LIDC dataset

Generative priors in inverse problems

Some examples...

Blind Image Deconvolution using Deep Generative Priors

Muhammad Asim^{*}, Fahad Shamshad^{*}, and Ali Ahmed

Abstract—This paper proposes a novel approach to regularize the *ill-posed* and *non-linear* blind image deconvolution (blind deblurring) using deep generative networks as priors. We employ two separate generative models — one trained to produce sharp images while the other trained to generate blur kernels from lower-dimensional parameters. To deblur, we propose an alternating gradient descent scheme operating in the latent lower-dimensional space of each of the pretrained generative models. Our experiments show promising deblurring results on images even under large blurs, and heavy noise. To address the shortcomings of generative models such as mode collapse, we augment our generative priors with classical image priors and report improved performance on complex image datasets. The deblurring performance depends on how well the range of the generator spans the image class. Interestingly, our experiments show that even an untrained structured (convolutional) generative networks acts as an image prior in the image deblurring context allowing us to extend our results to more diverse natural image datasets.


Index Terms—Blind image deblurring, generative adversarial networks, variational autoencoders, deep image prior.

I. INTRODUCTION

BLIND image deblurring aims to recover a true image i and a blur kernel k from blurry and possibly noisy observation y . For a uniform and spatially invariant blur, it can be mathematically formulated as

$$y = i \otimes k + n, \quad (1)$$

where \otimes is a convolution operator and n is an additive Gaussian noise. In its full generality, the inverse problem (1) is severely ill-posed as many different instances of i , and k fit the observation y . See [1], [2] for a thorough discussion on




(a) Blurry (b) Ours (c) Original

Fig. 1: Blind image deblurring using deep generative priors.

It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ





It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ


1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$




On twitter...

It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator




Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$

2) Maximum Mean Discrepancy

$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_{\theta})$$

Easier to train (no balancing)...




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)


$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$

2) Maximum Mean Discrepancy

$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_{\theta})$$

$$\sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{x}_j \in X}} \kappa(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{z}_j \in Z}} \kappa(\mathbf{x}_i, \mathcal{G}_{\theta}(\mathbf{z}_j)) + \sum_{\substack{\mathbf{z}_i \in Z \\ \mathbf{z}_j \in Z}} \kappa(\mathcal{G}_{\theta}(\mathbf{z}_i), \mathcal{G}_{\theta}(\mathbf{z}_j))$$

Easier to train (no balancing)...




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$


2) Maximum Mean Discrepancy

$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_{\theta})$$

$$\sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{x}_j \in X}} \kappa(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{z}_j \in Z}} \kappa(\mathbf{x}_i, \mathcal{G}_{\theta}(\mathbf{z}_j)) + \sum_{\substack{\mathbf{z}_i \in Z \\ \mathbf{z}_j \in Z}} \kappa(\mathcal{G}_{\theta}(\mathbf{z}_i), \mathcal{G}_{\theta}(\mathbf{z}_j))$$

Similarity between samples

Easier to train (no balancing)...




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$

2) Maximum Mean Discrepancy

$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_{\theta})$$

$$\sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{x}_j \in X}} \kappa(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{z}_j \in Z}} \kappa(\mathbf{x}_i, \mathcal{G}_{\theta}(\mathbf{z}_j)) + \sum_{\substack{\mathbf{z}_i \in Z \\ \mathbf{z}_j \in Z}} \kappa(\mathcal{G}_{\theta}(\mathbf{z}_i), \mathcal{G}_{\theta}(\mathbf{z}_j))$$


Similarity between samples

equivalent to (for later)

$$\mathbb{E}_{\omega \sim \Lambda} \left| \sum_{\mathbf{x}_i \in X} e^{i\omega^T \mathbf{x}_i} - \sum_{\mathbf{z}_i \in Z} e^{i\omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right|^2$$

Easier to train (no balancing)...

With $\Lambda = \mathcal{F}\kappa$




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$

2) Maximum Mean Discrepancy

$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_{\theta})$$

equivalent to (for later)


$$\sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{x}_j \in X}} \kappa(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{z}_j \in Z}} \kappa(\mathbf{x}_i, \mathcal{G}_{\theta}(\mathbf{z}_j)) + \sum_{\substack{\mathbf{z}_i \in Z \\ \mathbf{z}_j \in Z}} \kappa(\mathcal{G}_{\theta}(\mathbf{z}_i), \mathcal{G}_{\theta}(\mathbf{z}_j))$$

Similarity between samples

$$\mathbb{E}_{\omega \sim \Lambda} \left| \sum_{\mathbf{x}_i \in X} e^{i\omega^T \mathbf{x}_i} - \sum_{\mathbf{z}_i \in Z} e^{i\omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right|^2$$

Easier to train (no balancing) but quadratic complexity...

With $\Lambda = \mathcal{F}_{\kappa}$




It's nice! ...where's the catch?

How to learn the generative network \mathcal{G}_θ

1) Golden standard: Generative Adversarial Networks

Learn a second “discriminator” network that classifies real/fake at the same time as the generator



Very difficult to train (due to balancing of training discriminator/generator)

$$\min_{\theta} \max_{\phi} \mathcal{L}(\theta, \phi)$$

2) Maximum Mean Discrepancy


$$\min_{\theta} \text{MMD}_{\kappa}(\hat{\mathcal{P}}_X, \hat{\mathcal{P}}_\theta)$$

equivalent to (for later)

$$\sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{x}_j \in X}} \kappa(\mathbf{x}_i, \mathbf{x}_j) - 2 \sum_{\substack{\mathbf{x}_i \in X \\ \mathbf{z}_j \in Z}} \kappa(\mathbf{x}_i, \mathcal{G}_\theta(\mathbf{z}_j)) + \sum_{\substack{\mathbf{z}_i \in Z \\ \mathbf{z}_j \in Z}} \kappa(\mathcal{G}_\theta(\mathbf{z}_i), \mathcal{G}_\theta(\mathbf{z}_j))$$

Similarity between samples

Easier to train (no balancing) but quadratic complexity...





Training generative networks typically requires massive amounts of data!

$$\left| \sum_{\mathbf{z}_i \in Z} e^{i\omega^T \mathcal{G}_\theta(\mathbf{z}_i)} \right|^2$$

Compressive Learning to the rescue!



Compressive Learning: principle

Usual machine learning




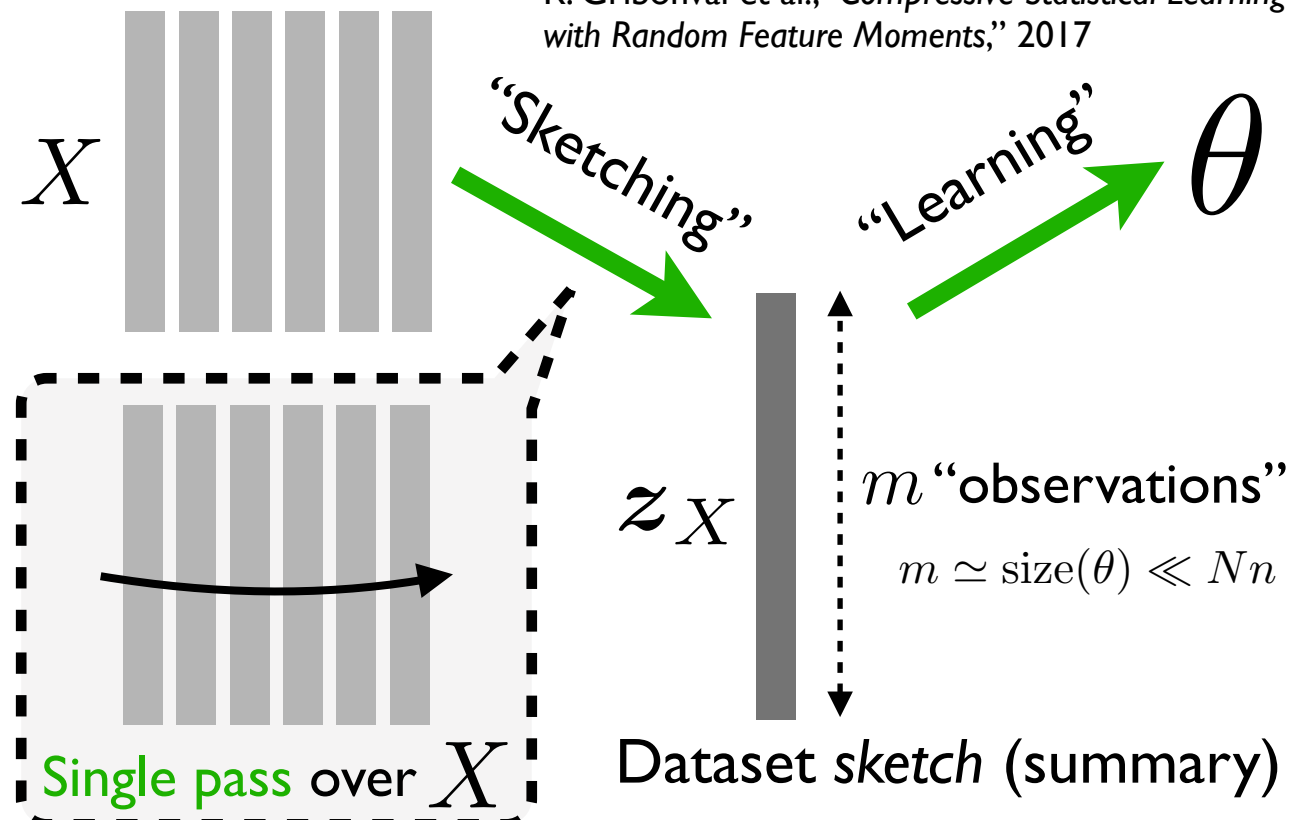
Compressive Learning: principle

Usual machine learning




Compressive Learning

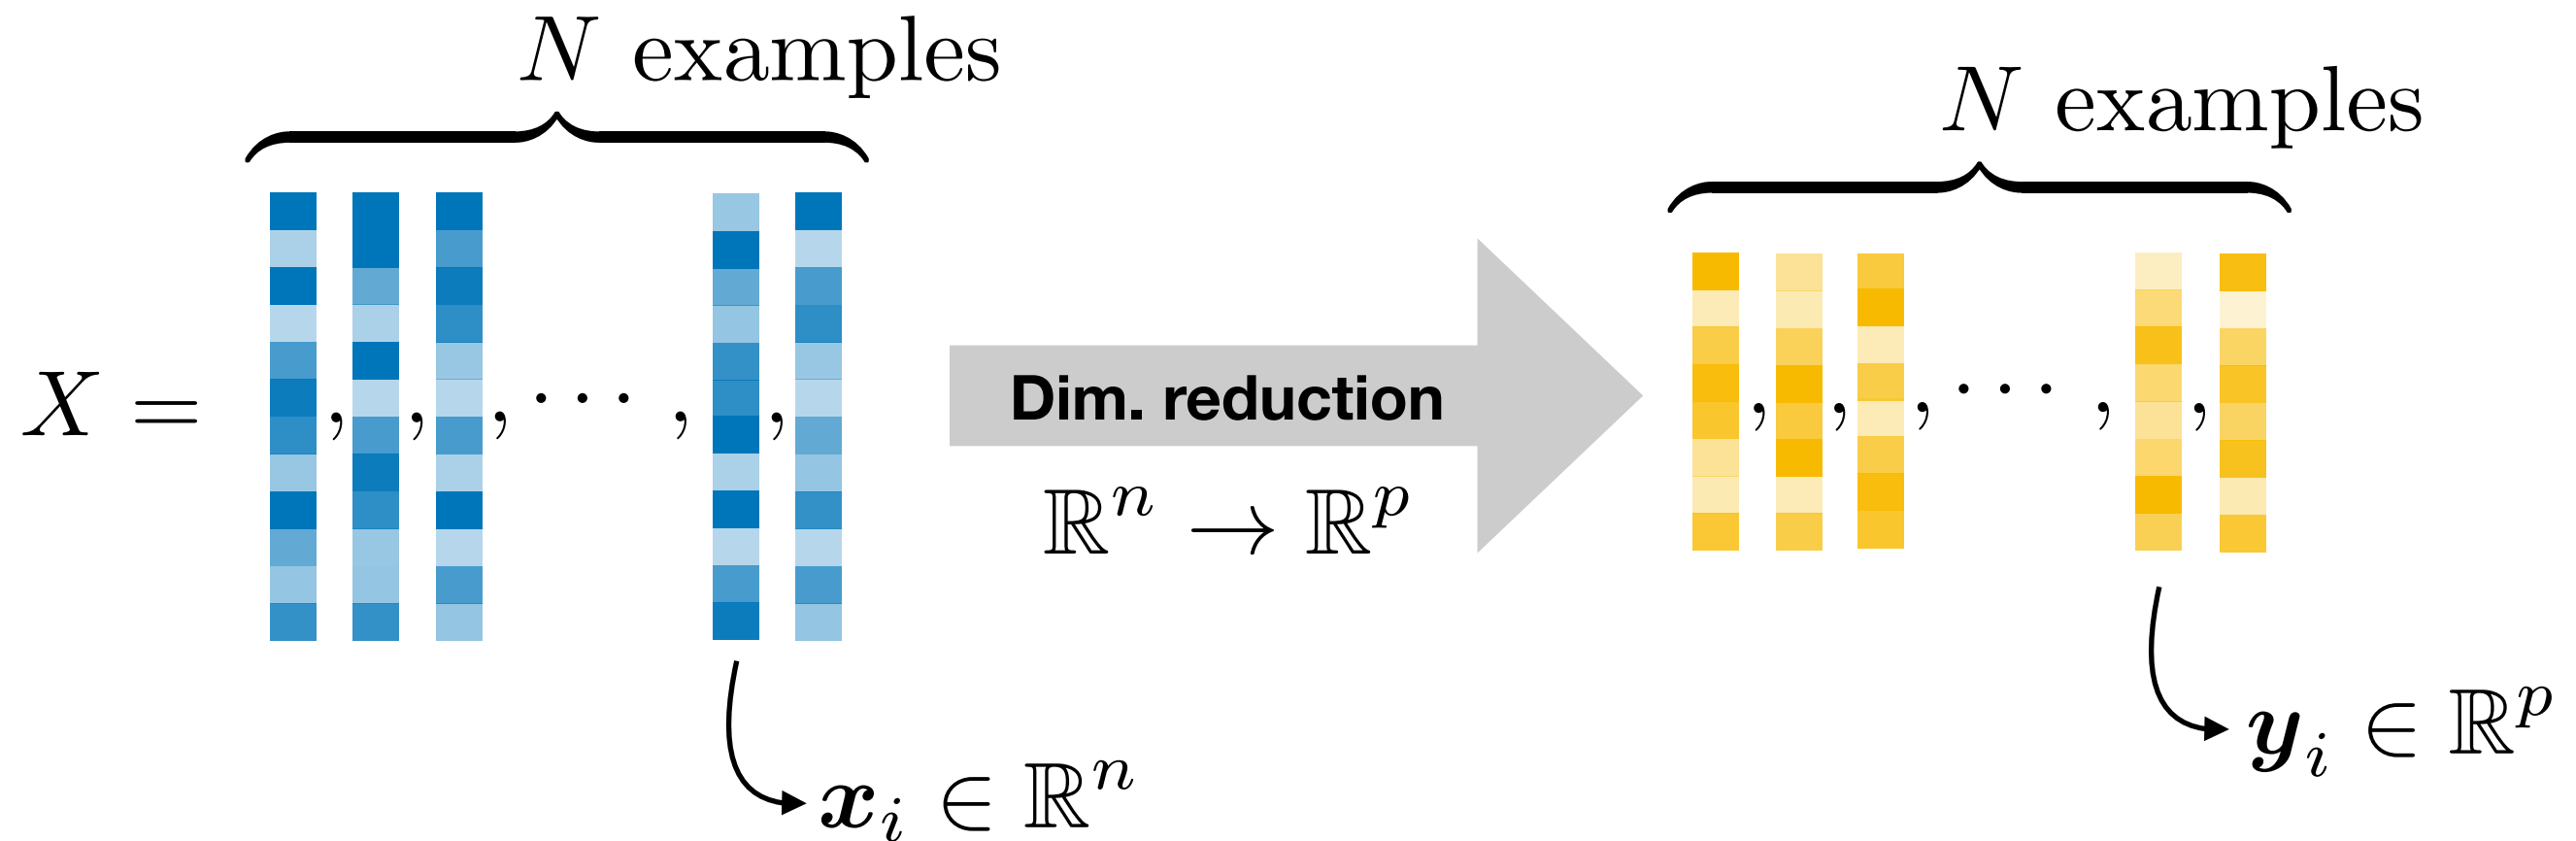
R. Gribonval et al., "Compressive Statistical Learning with Random Feature Moments," 2017



Sketching a dataset

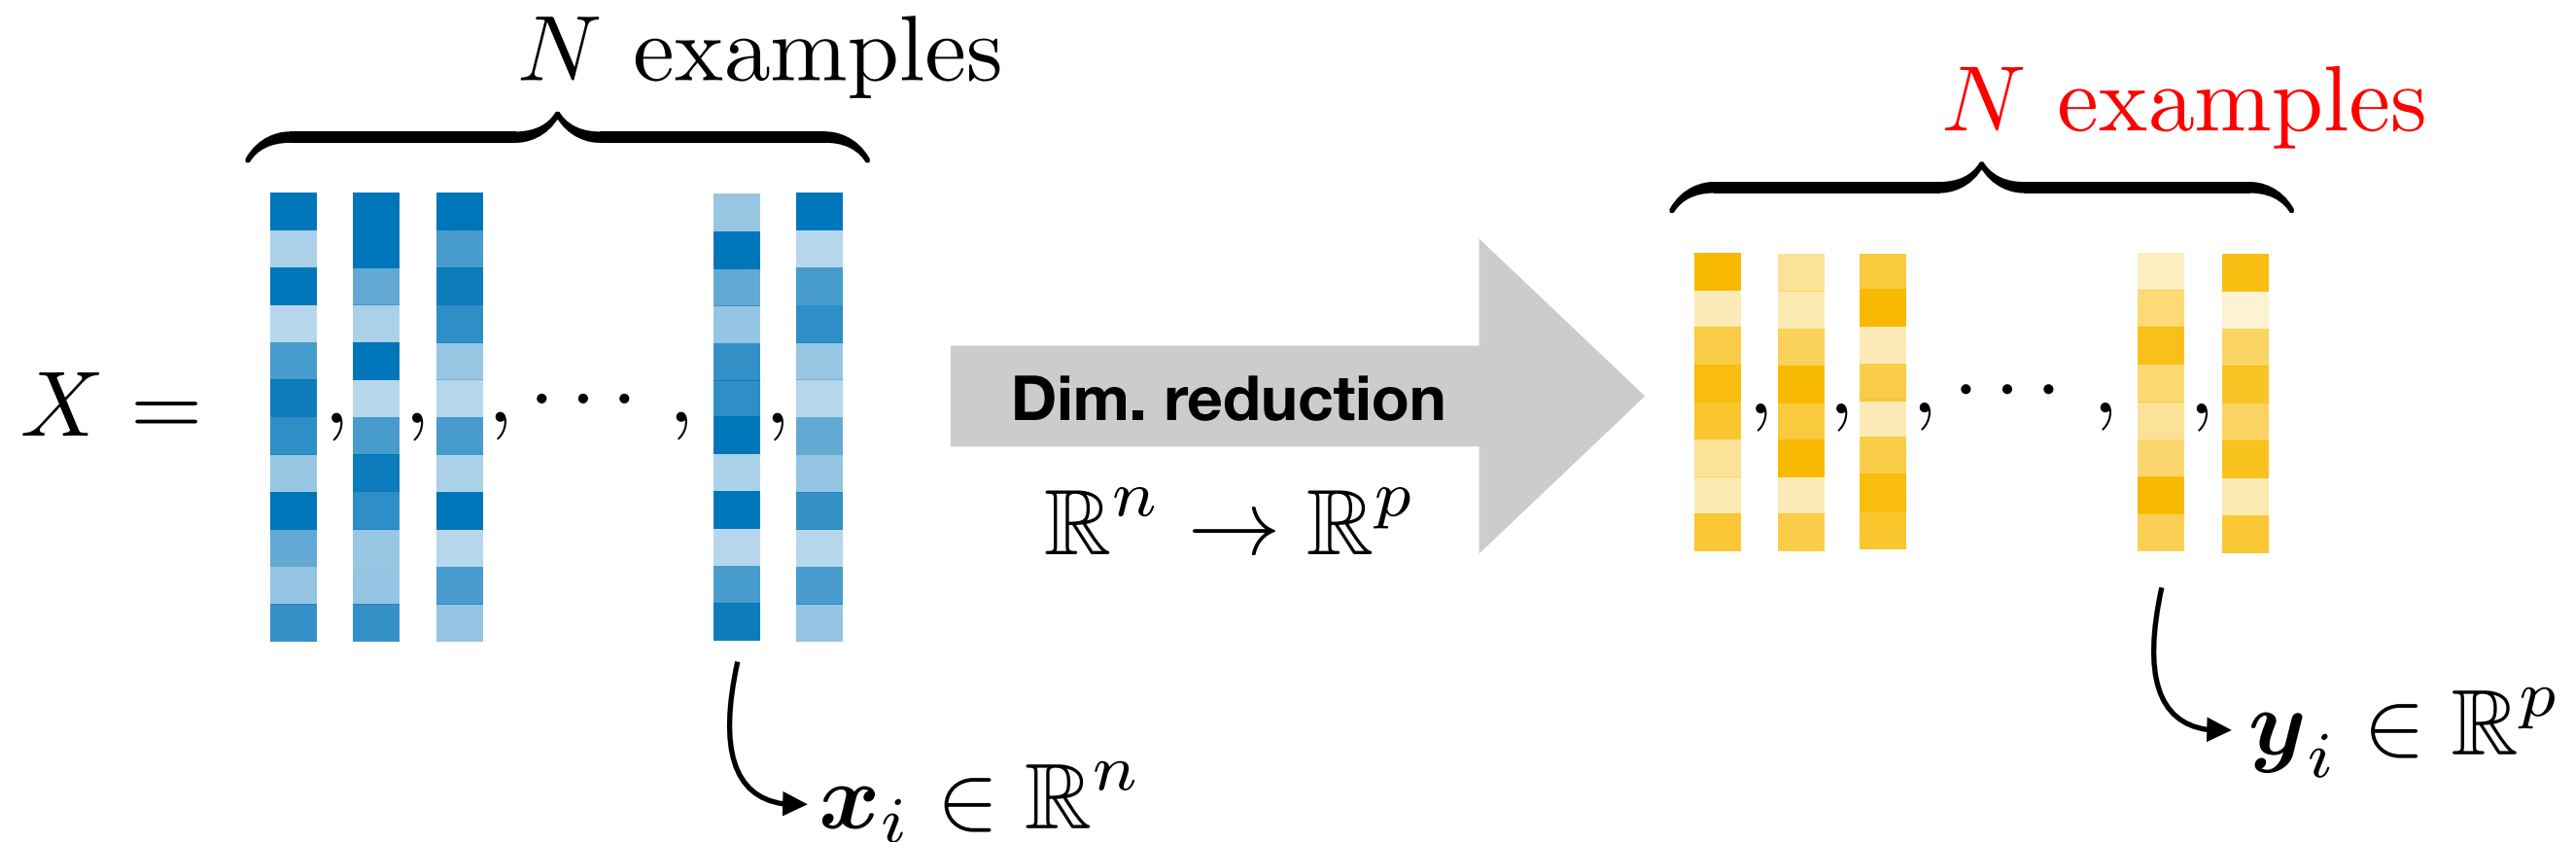


Sketching a dataset



- Compressed representation ✓
- Preserves relevant information ✓


Sketching a dataset



- Compressed representation ✓
- Preserves relevant information ✓
- **Constant number of examples** ✗


N can be **VERY** large (“big data”)!

Sketching a dataset



- **Compressed representation** ✓
- **Preserves relevant information** ✓
- **Dataset summary = single vector** ✓


Sketching a dataset



- **Compressed representation** ✓
- **Preserves relevant information** ✓
- **Dataset summary = single vector** ✓


Sketching a dataset

z_X



=

Ω^T



x_i



.


$\omega_j \sim \Lambda$



1. Project on m (random) vectors


e.g.

Λ



Controls the *cluster scale*

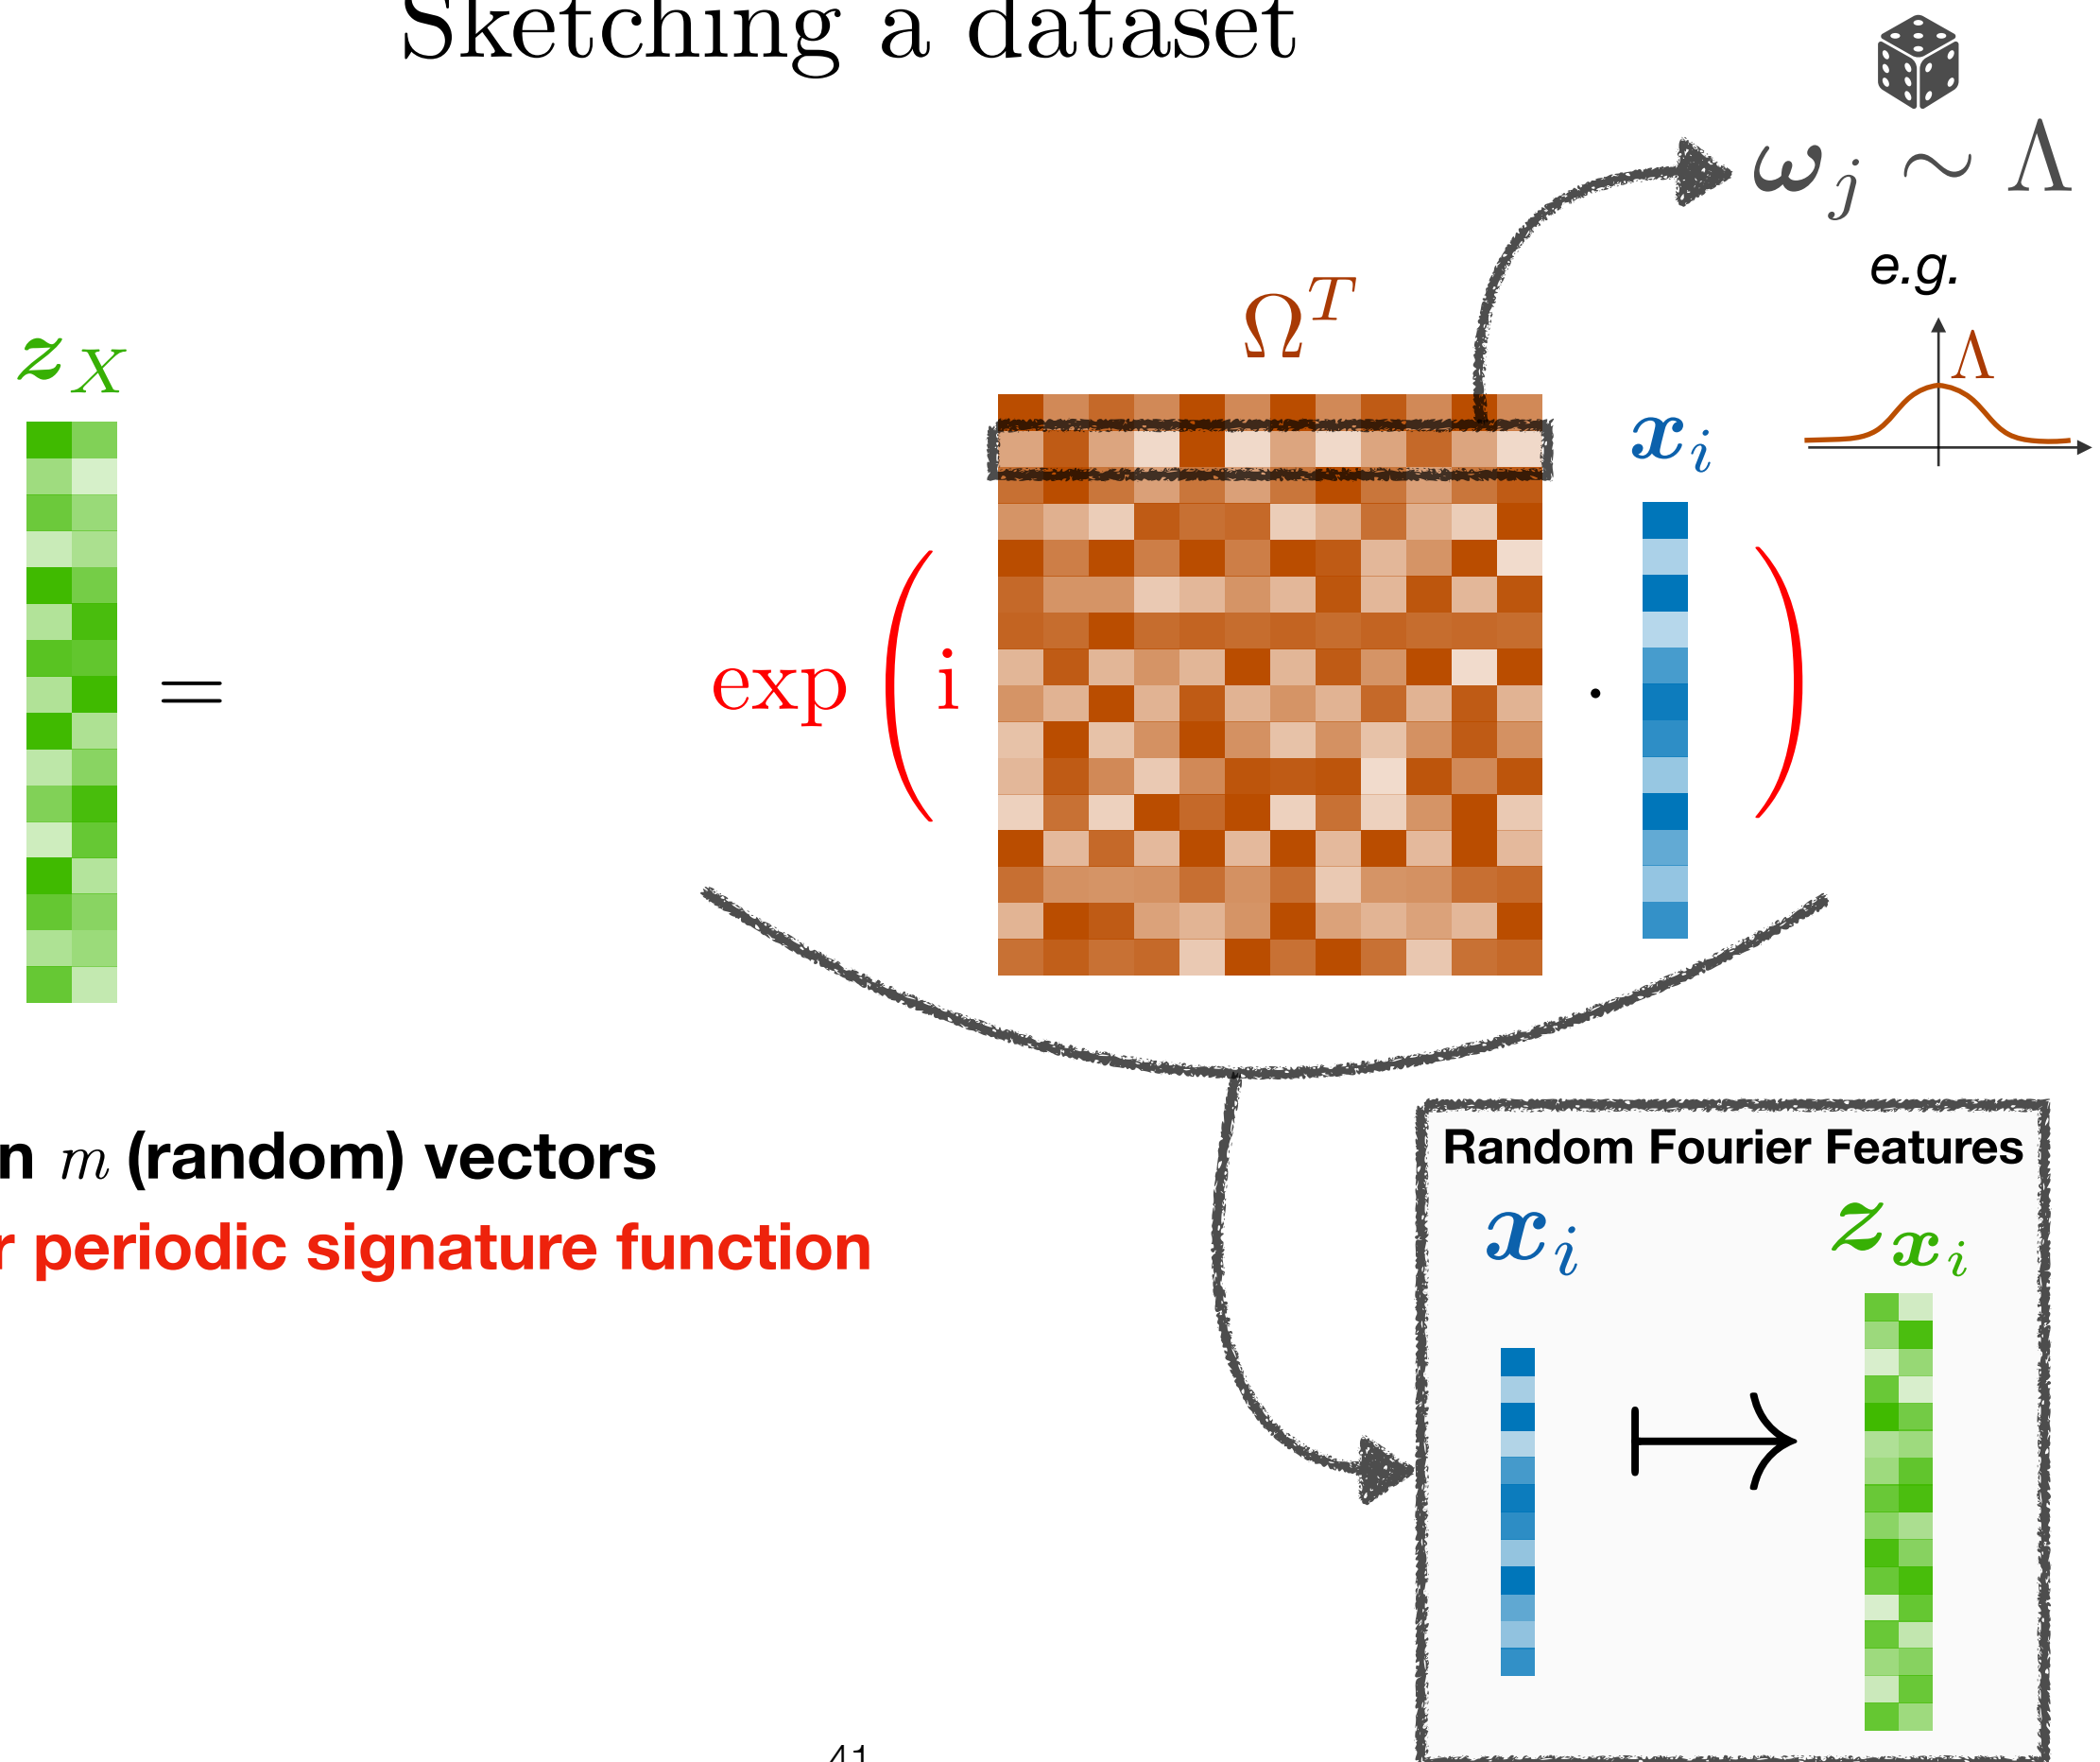
Sketching a dataset



1. Project on m (random) vectors

2. Nonlinear periodic signature function


Sketching a dataset




1. Project on m (random) vectors

2. Nonlinear periodic signature function


Sketching a dataset



1. Project on m (random) vectors
2. Nonlinear periodic signature function
3. Pooling (average)



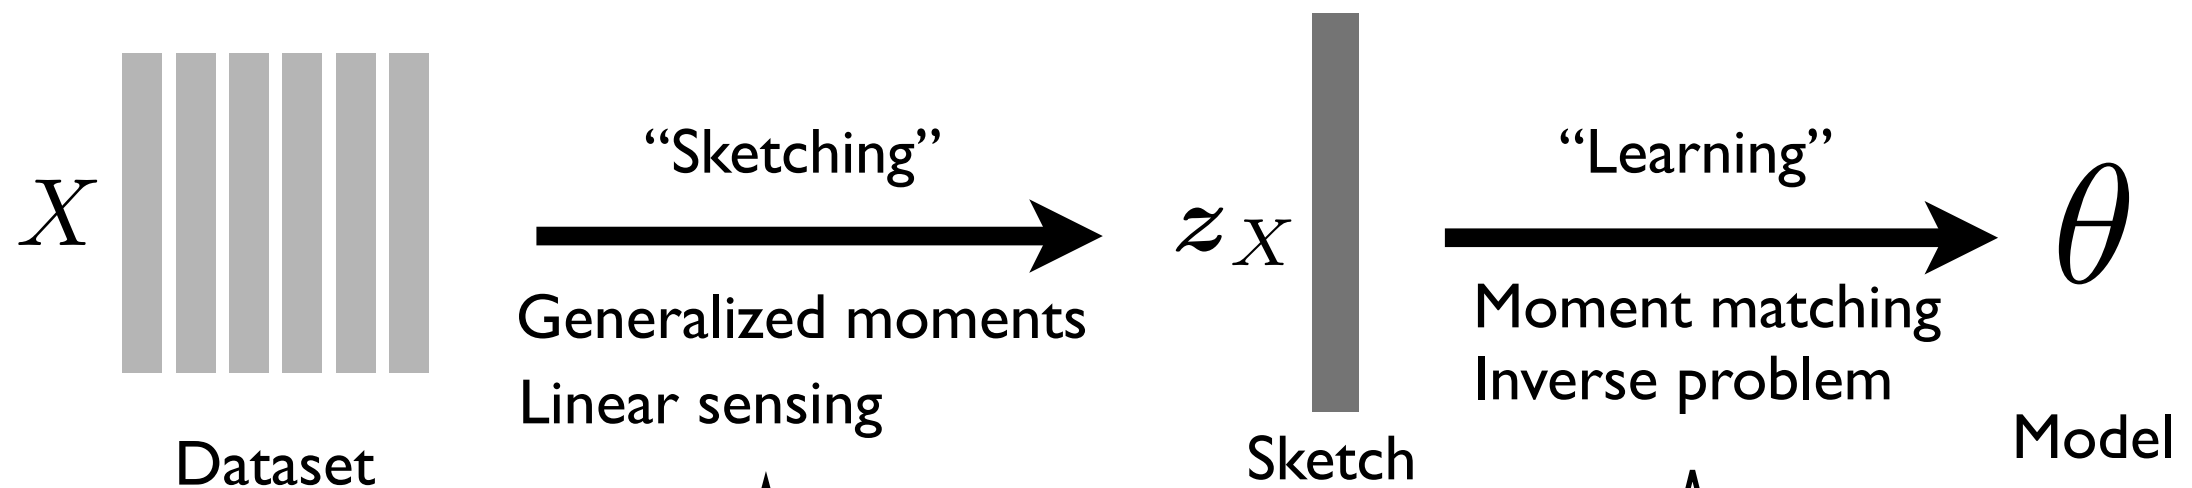
Sketching a dataset



1. Project on m (random) vectors
2. Nonlinear periodic signature function
3. Pooling (average)

$$z_X = \left[\frac{1}{N} \sum_{x_i \in X} e^{i\omega_j^T x_i} \right]_{j=1}^m \in \mathbb{C}^m$$

Compressive Learning: SoA in one slide



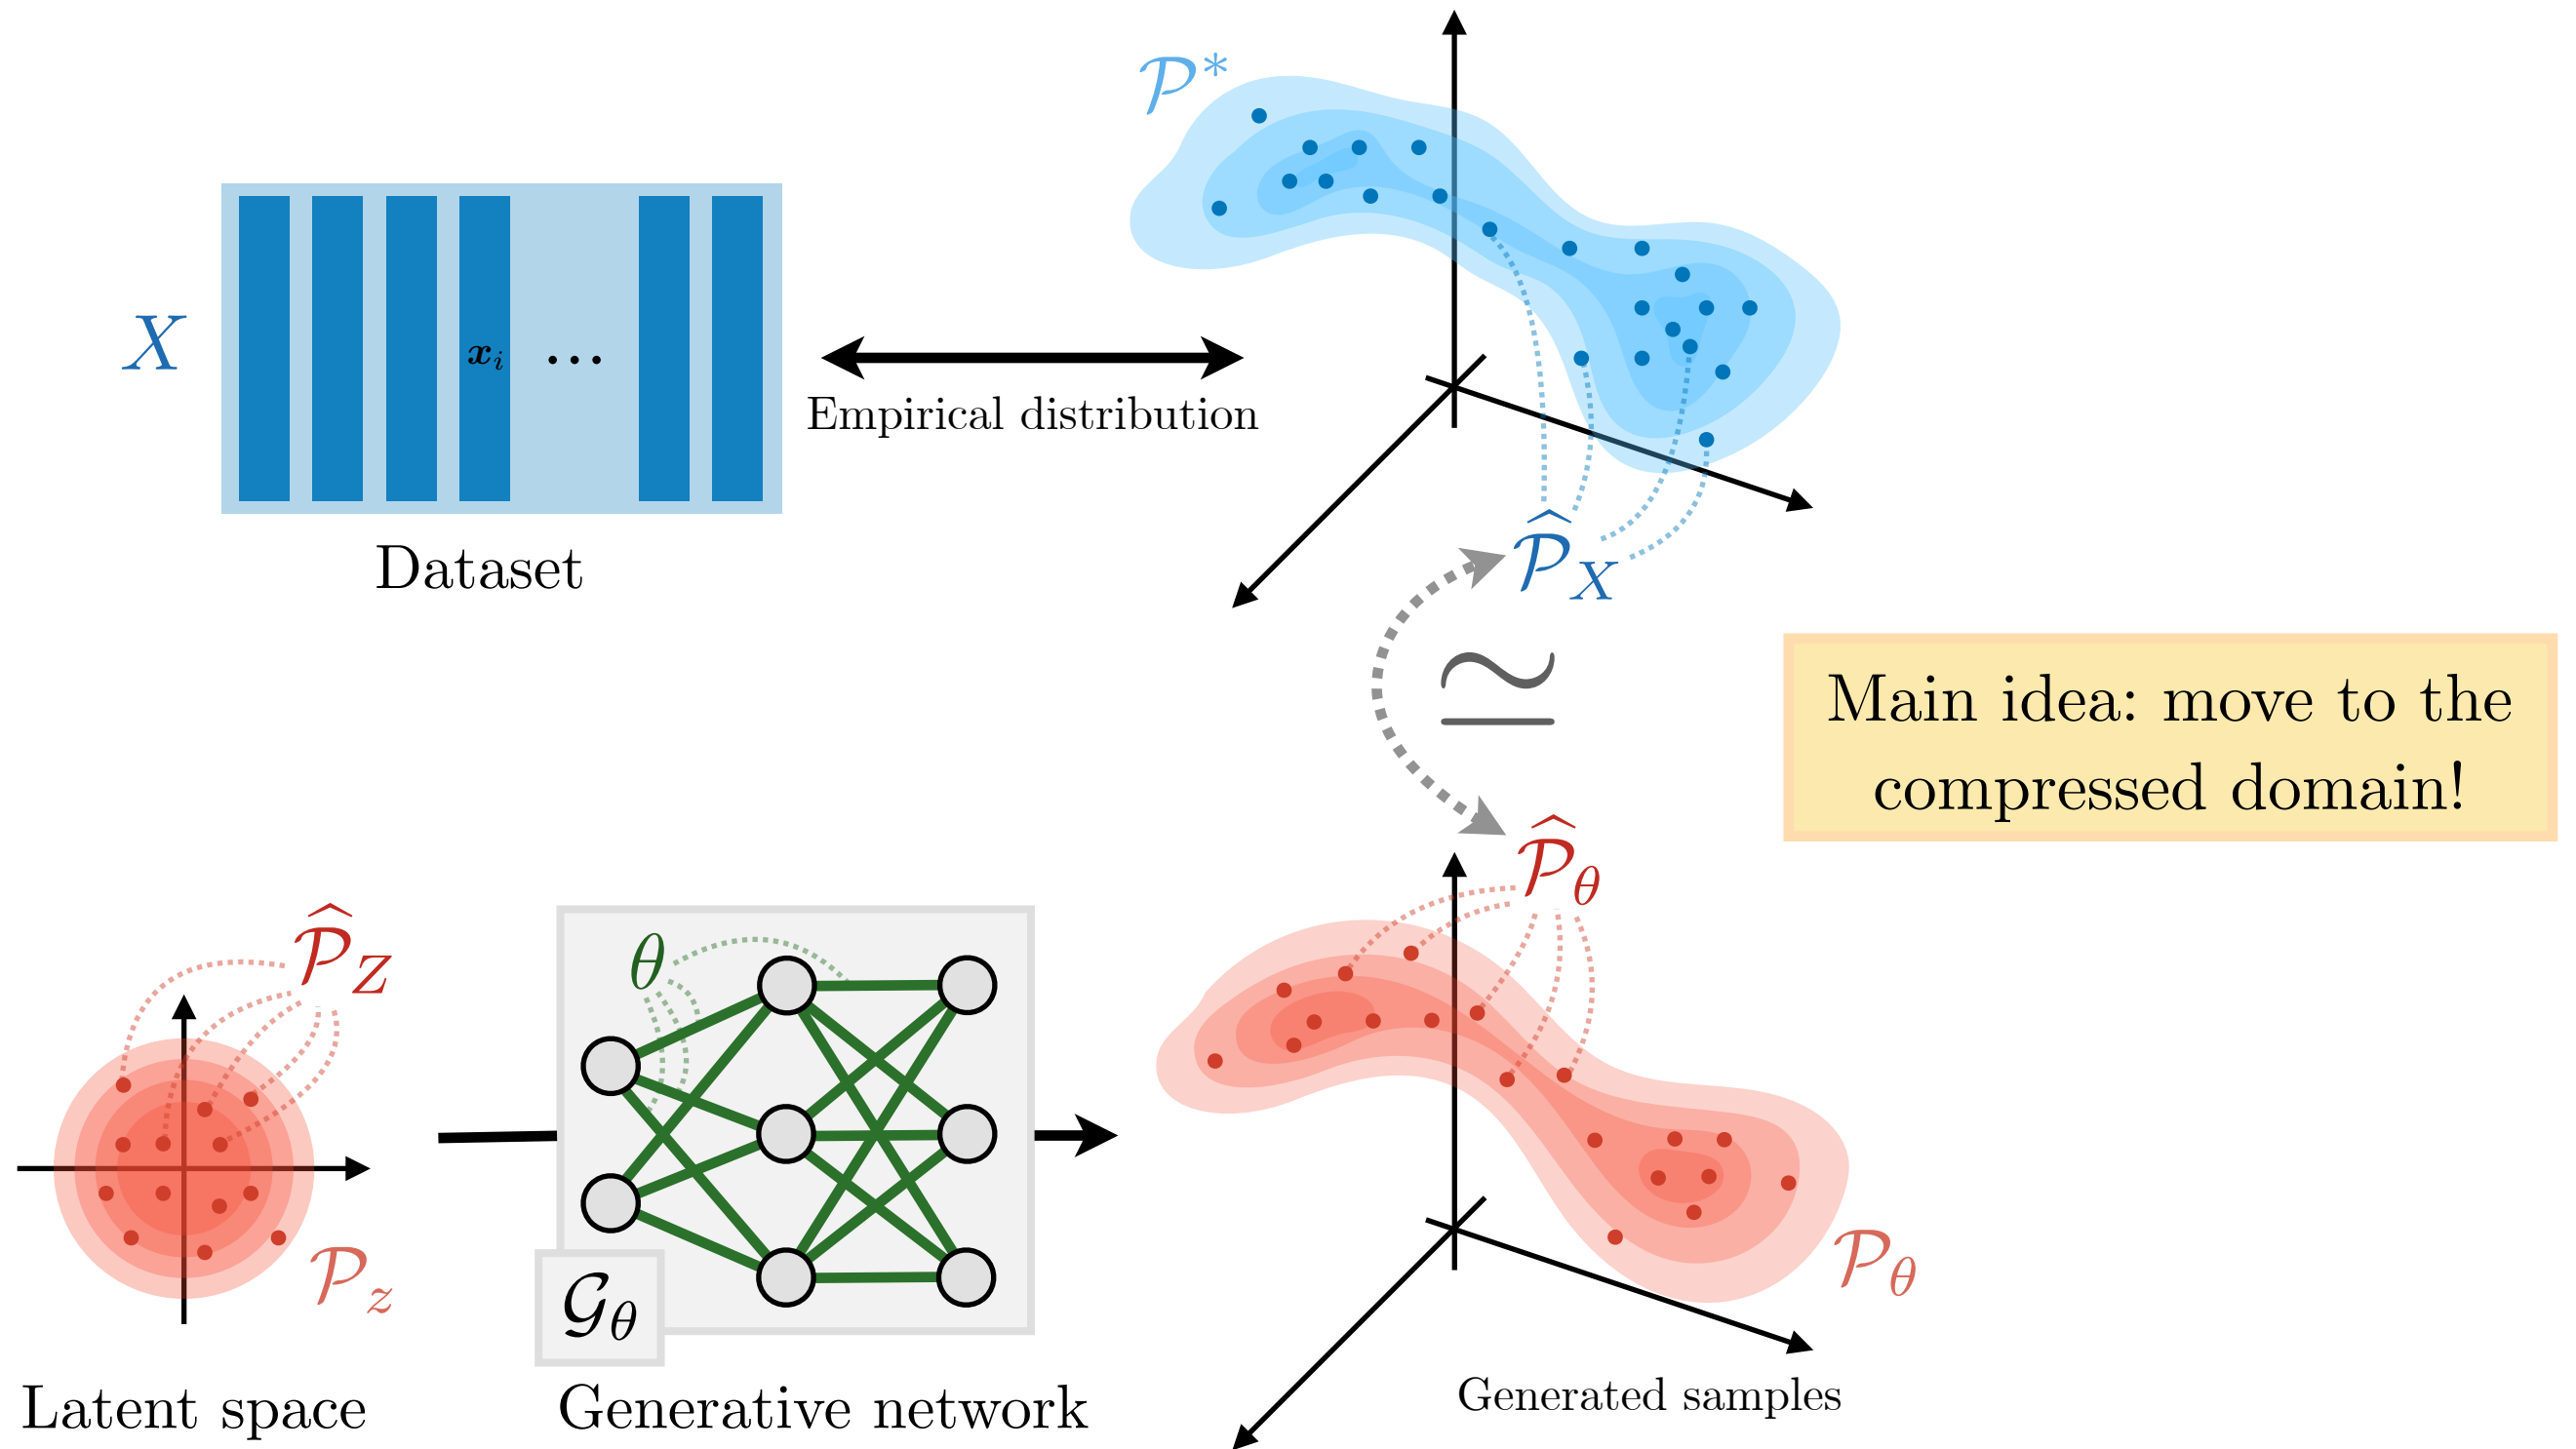
Advantages:

- One single pass on dataset (parallelizable)
- Harsh compression (ideal for large-scale datasets)
- Handy for privacy preservation

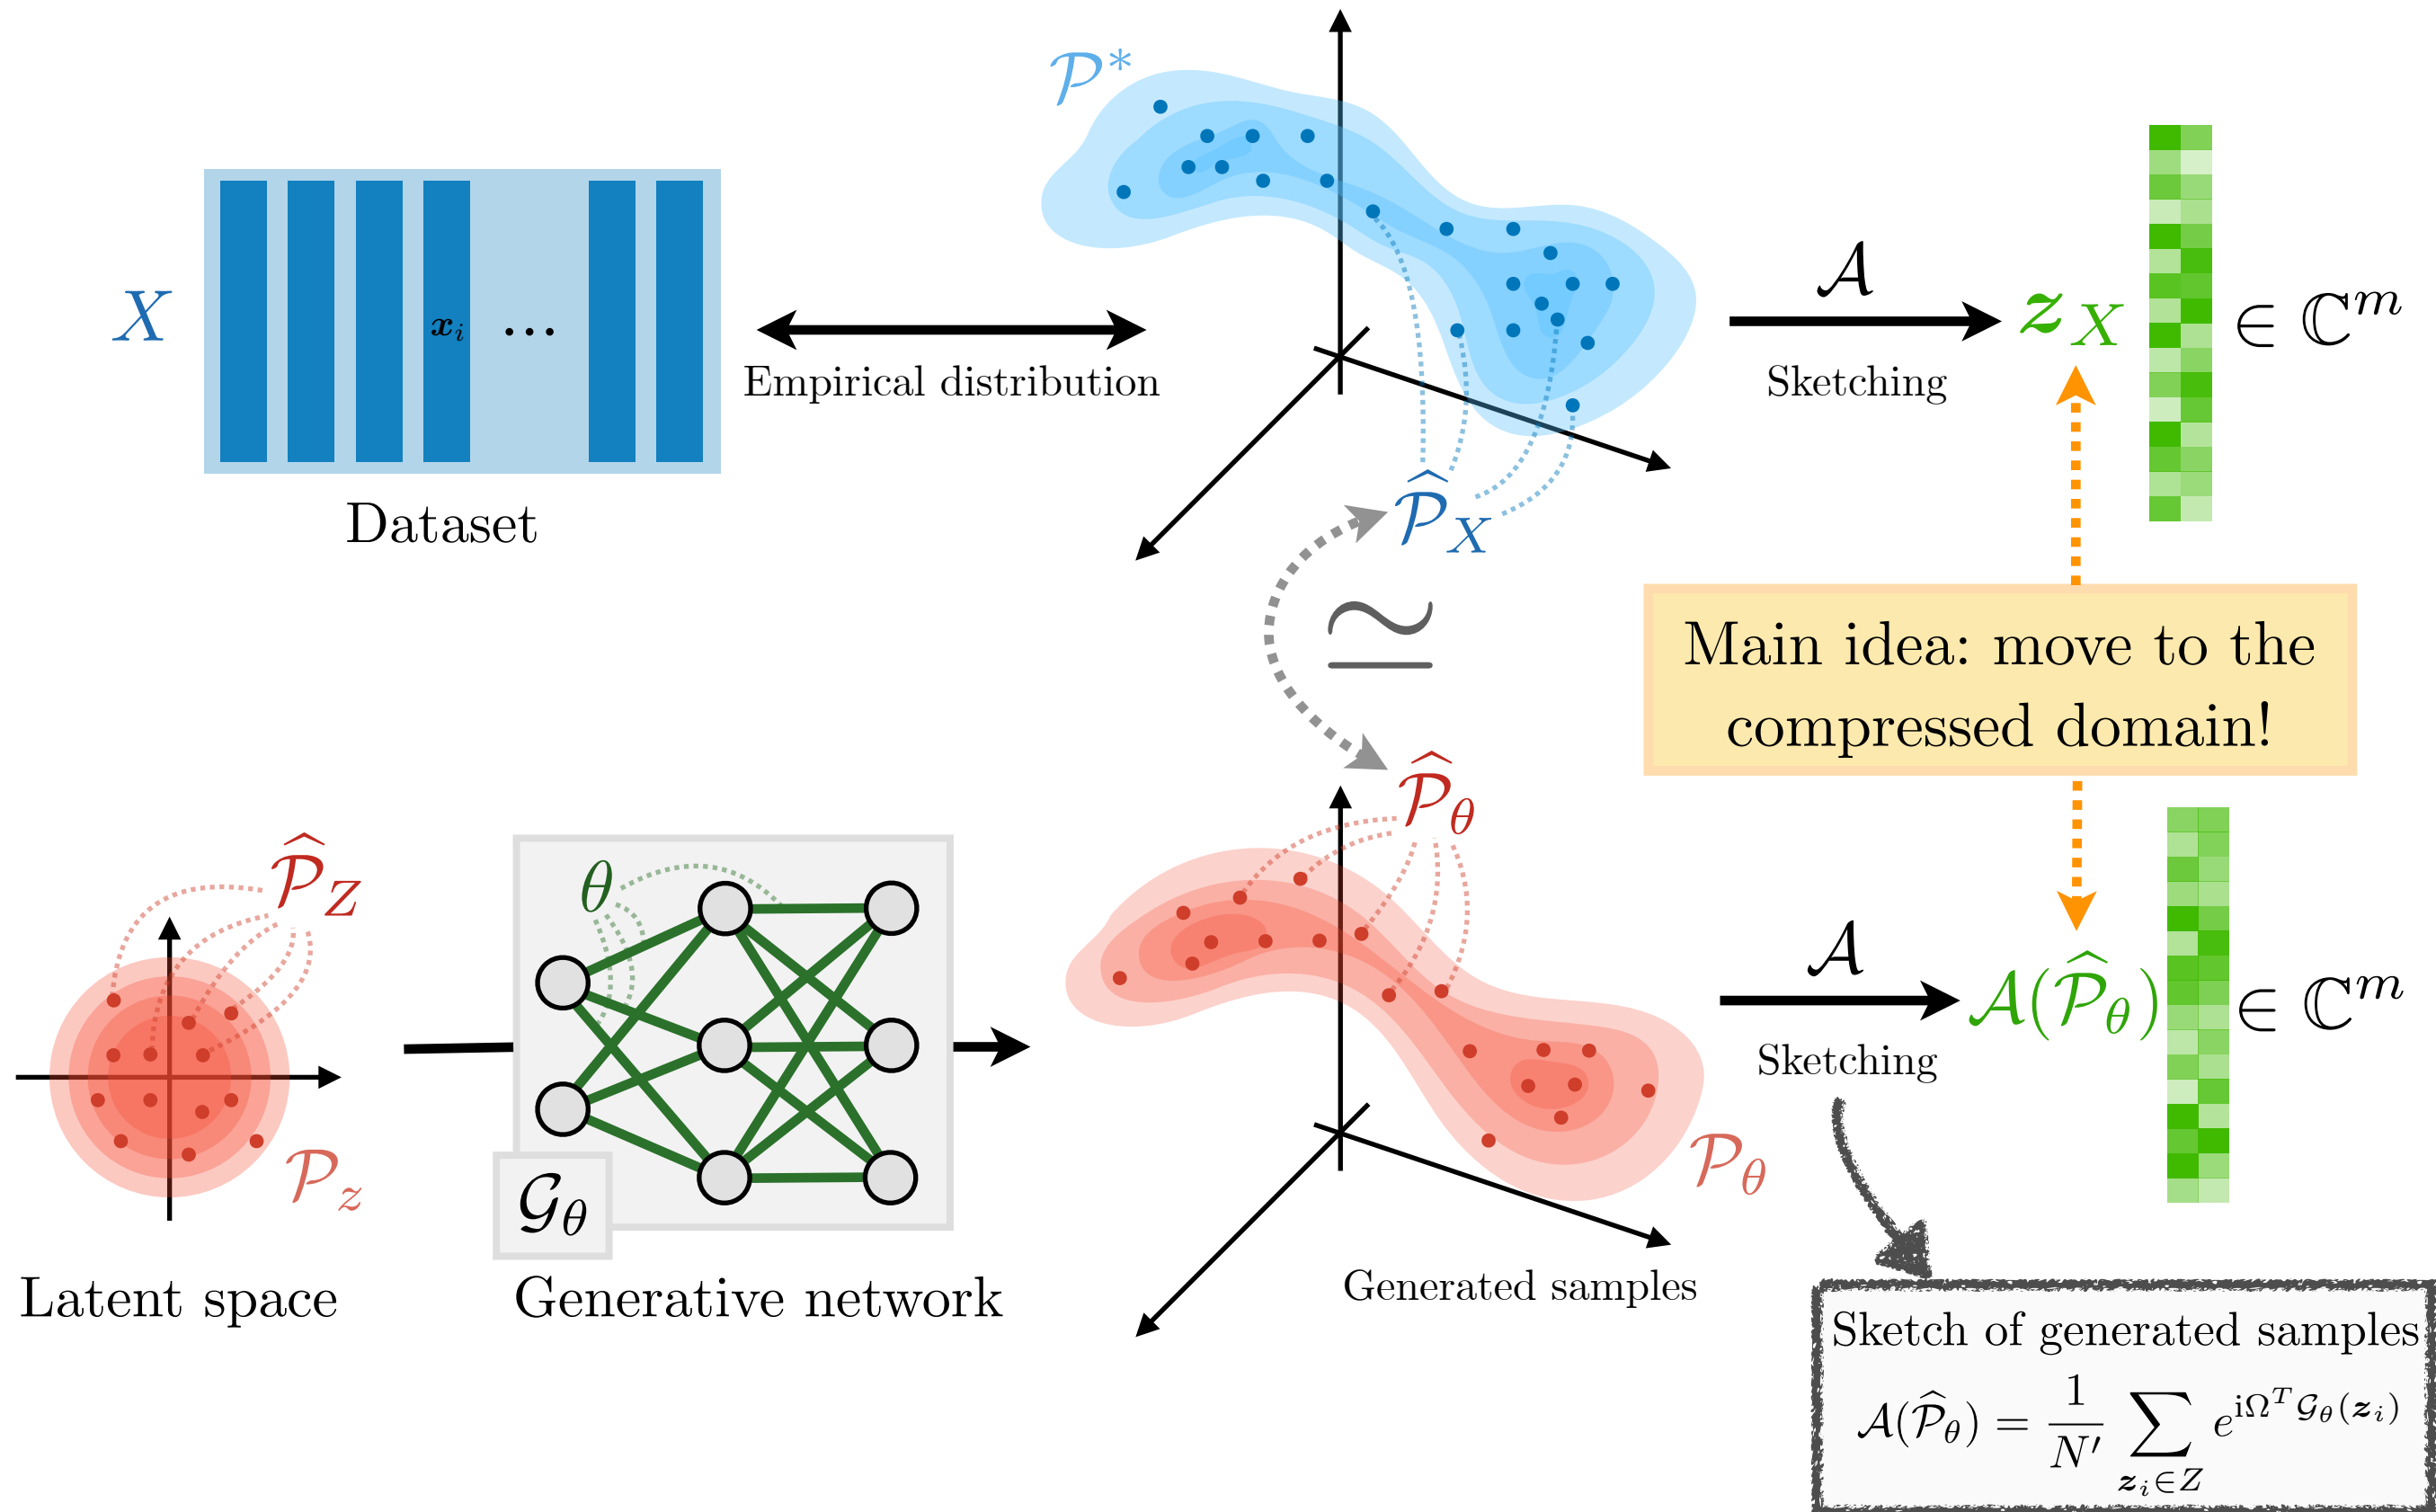
Existing (unsupervised) tasks:

- clustering: k-means, subspace clustering
- mixture model estimation: GMM, alpha-stable distributions
- Principal Component Analysis
- Independent Component Analysis
- **Generative networks (this work)**

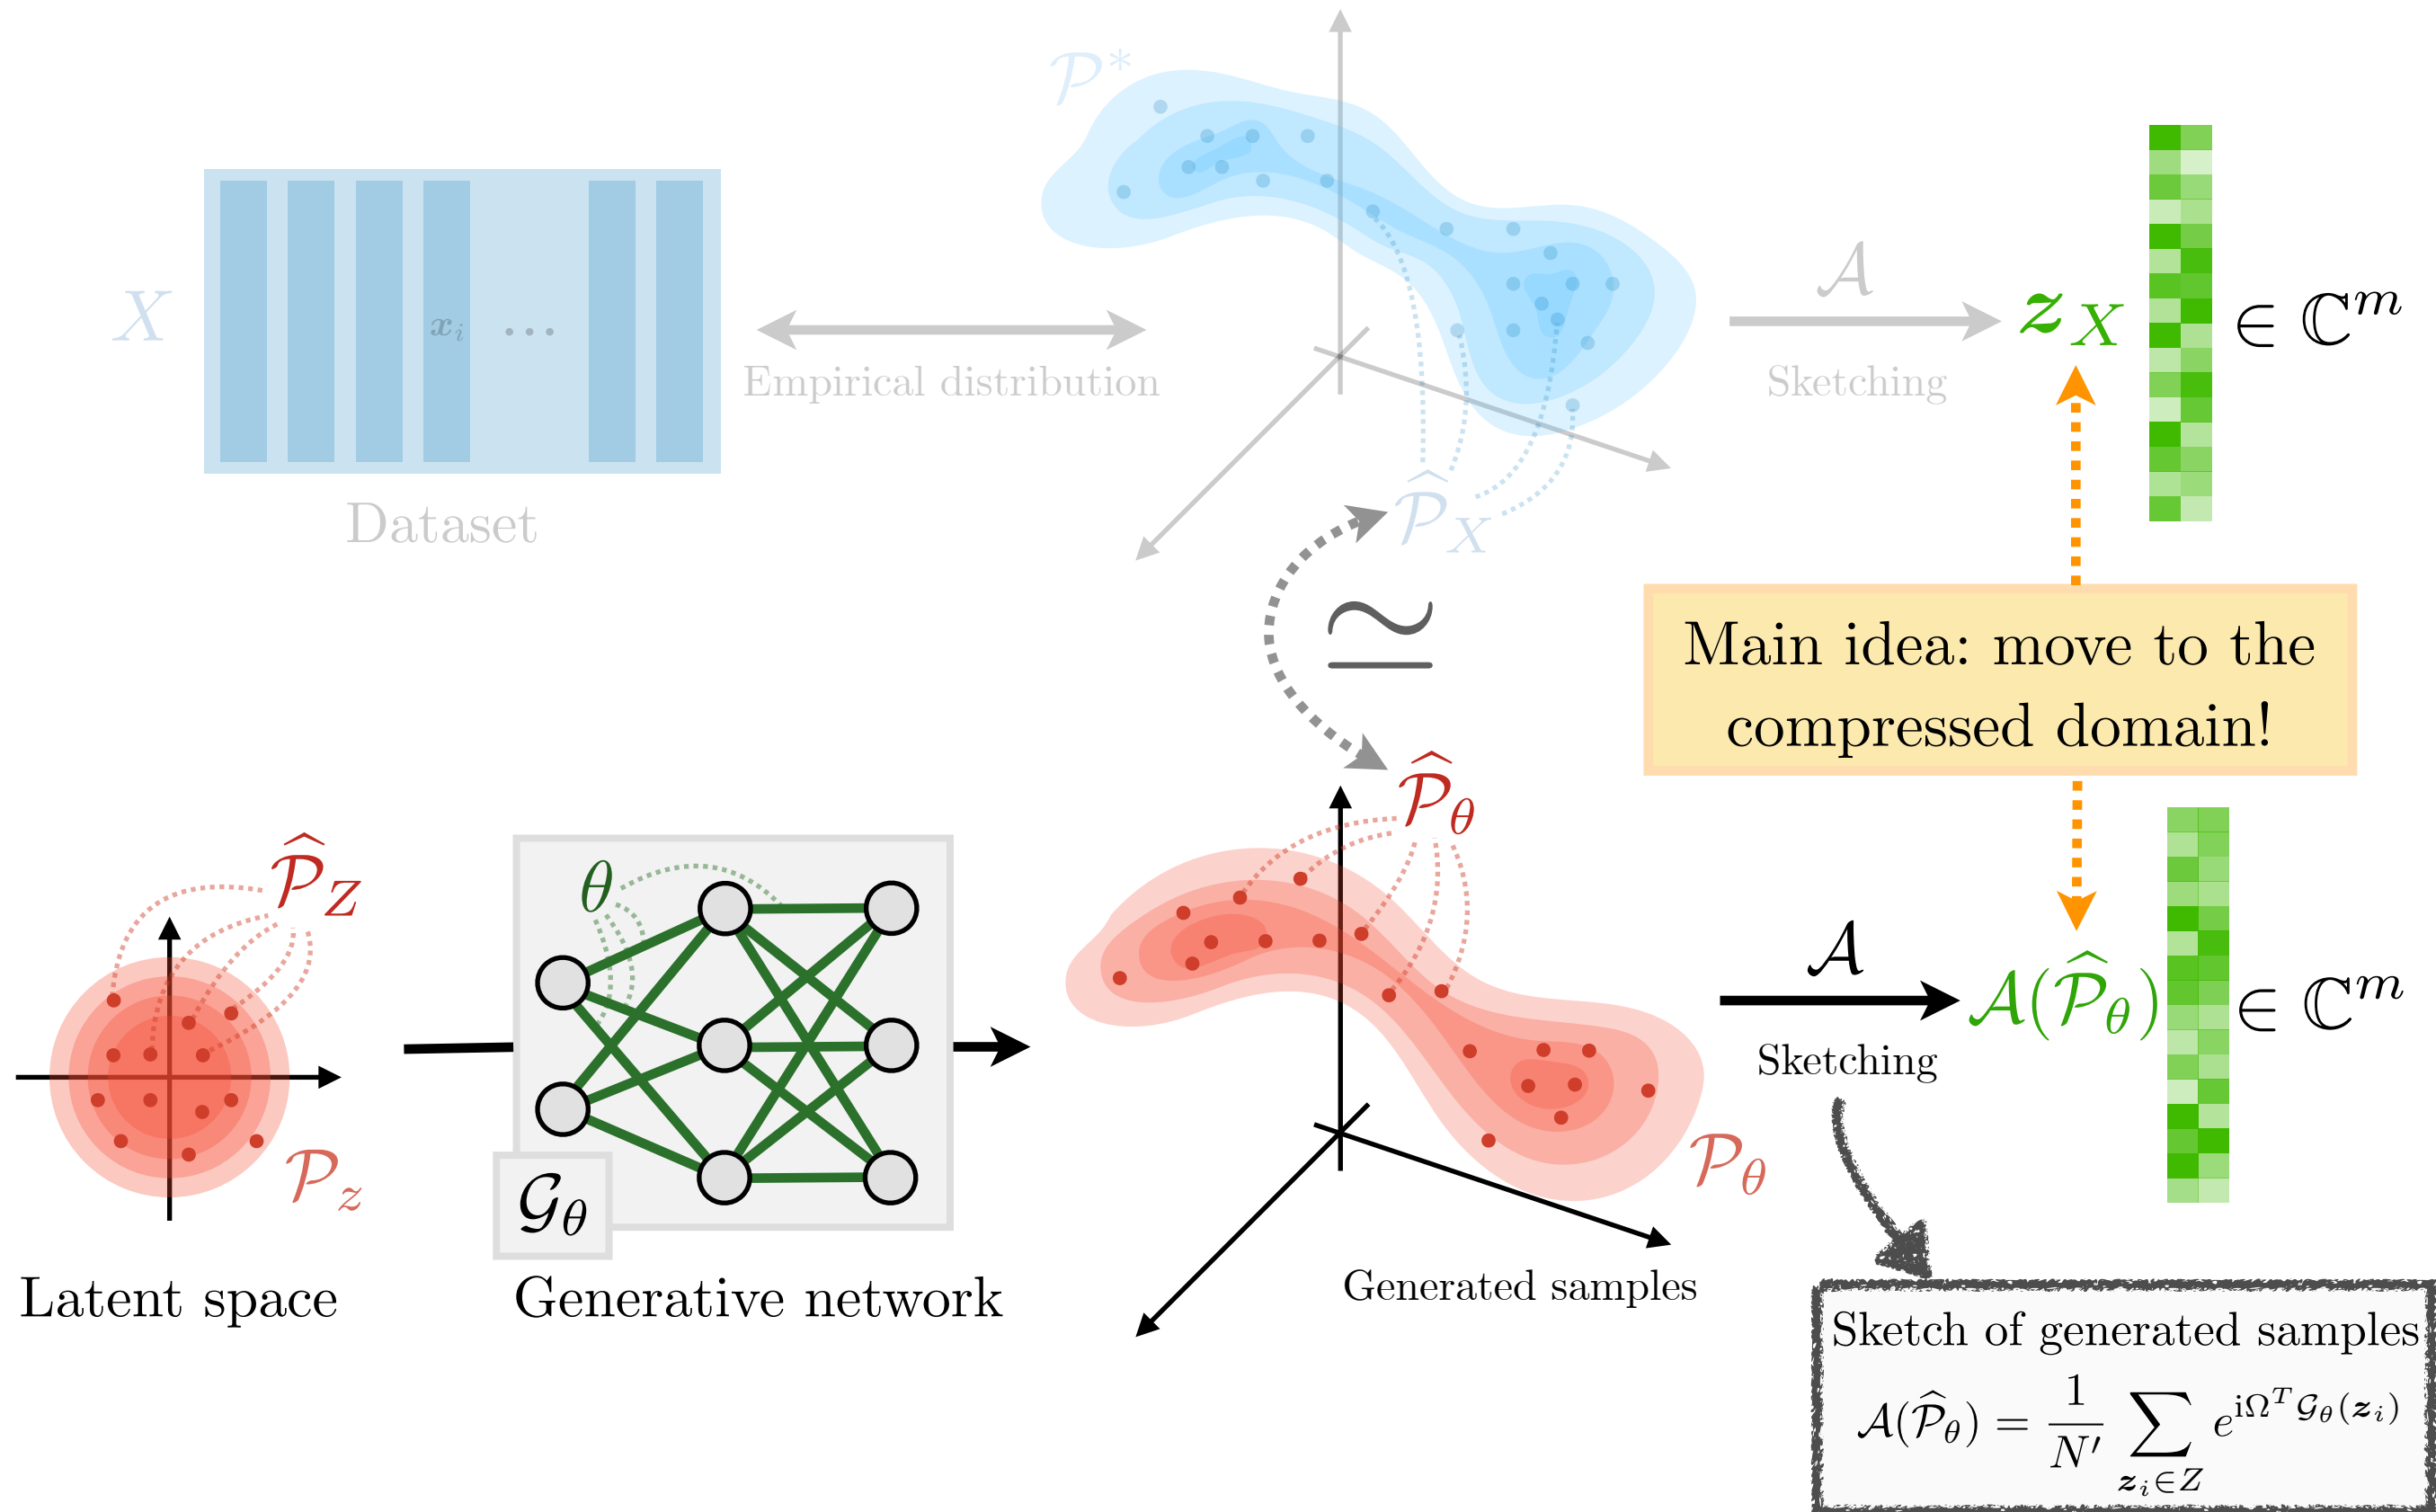
Compressively learning generative networks



Compressively learning generative networks



Compressively learning generative networks



Compressively learning generative networks

Proposed approach: match the sketches of real and generated data

$$\min_{\theta} \left\| \mathbf{z}_X - \frac{1}{N'} \sum_{\mathbf{z}_i \in Z} e^{i\Omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right\|_2^2$$

Compressively learning generative networks

Proposed approach: match the sketches of real and generated data

$$\min_{\theta} \left\| \mathbf{z}_X - \frac{1}{N'} \sum_{\mathbf{z}_i \in Z} e^{i\Omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right\|_2^2$$

Sampled (Monte Carlo)
estimation to the MMD!

$$\omega_j \sim \Lambda$$

$$\min_{\theta} \mathbb{E}_{\omega \sim \Lambda} \left| \frac{1}{N} \sum_{\mathbf{x}_i \in X} e^{i\omega^T \mathbf{x}_i} - \frac{1}{N'} \sum_{\mathbf{z}_i \in Z} e^{i\omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right|^2$$

...but where dataset is accessed only once 🍀

Compressively learning generative networks

Proposed approach: match the sketches of real and generated data

$$\min_{\theta} \left\| \mathbf{z}_X - \frac{1}{N'} \sum_{\mathbf{z}_i \in Z} e^{i\Omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right\|_2^2$$

Sampled (Monte Carlo)
estimation to the MMD!

$$\omega_j \sim \Lambda$$

Practical learning algorithm?

$$\min_{\theta} \mathbb{E}_{\omega \sim \Lambda} \left| \frac{1}{N} \sum_{\mathbf{x}_i \in X} e^{i\omega^T \mathbf{x}_i} - \frac{1}{N'} \sum_{\mathbf{z}_i \in Z} e^{i\omega^T \mathcal{G}_{\theta}(\mathbf{z}_i)} \right|^2$$

...but where dataset is accessed only once 🍀




Differentiable by chain rule (feat. backprop) 🍀

$$\nabla_{\theta} \hat{\mathcal{L}}(\theta; \mathbf{z}_X) = -2 \cdot \frac{1}{n'} \sum_{i=1}^{n'} \Re \left[\mathbf{r}^H \left(\frac{\partial \Phi(\mathbf{u})}{\partial \mathbf{u}} \Big|_{\mathbf{u}=\mathcal{G}_{\theta}(\mathbf{z}_i)} \cdot \frac{\partial \mathcal{G}_{\theta}(\mathbf{z}_i)}{\partial \theta} \right) \right]$$

Preliminary results

Simple 2d signals...

...it works (in principle)!



To conclude: open challenges

