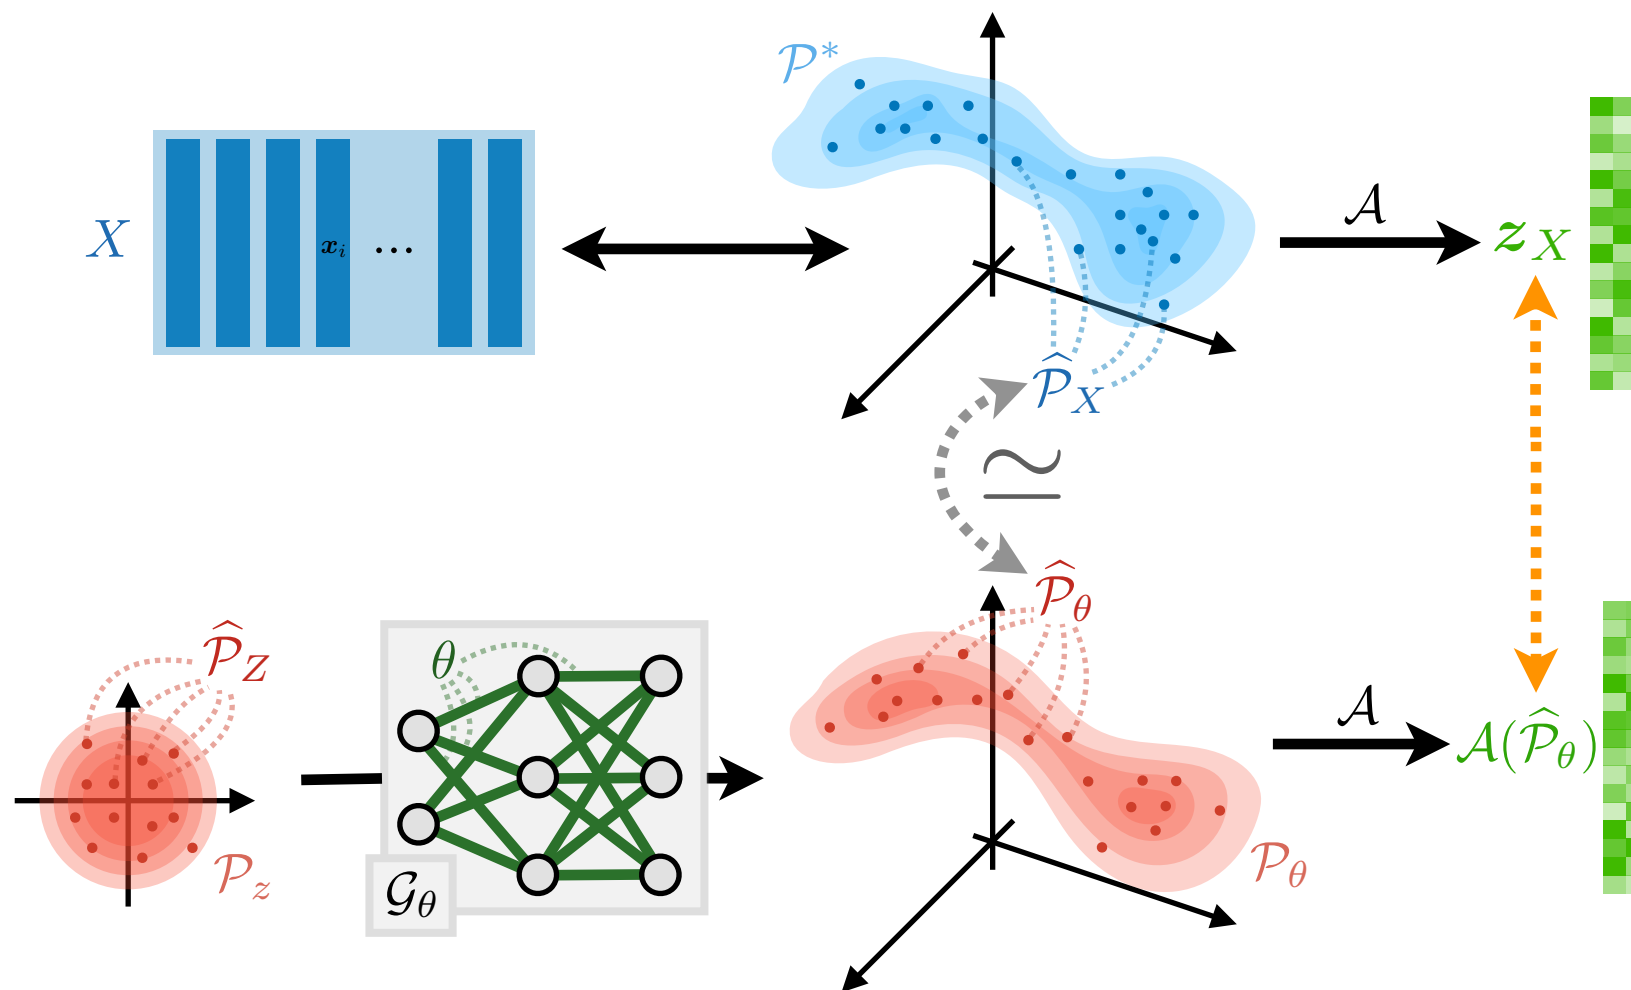


Compressive Learning of Generative Networks

Vincent Schellekens & Laurent Jacques

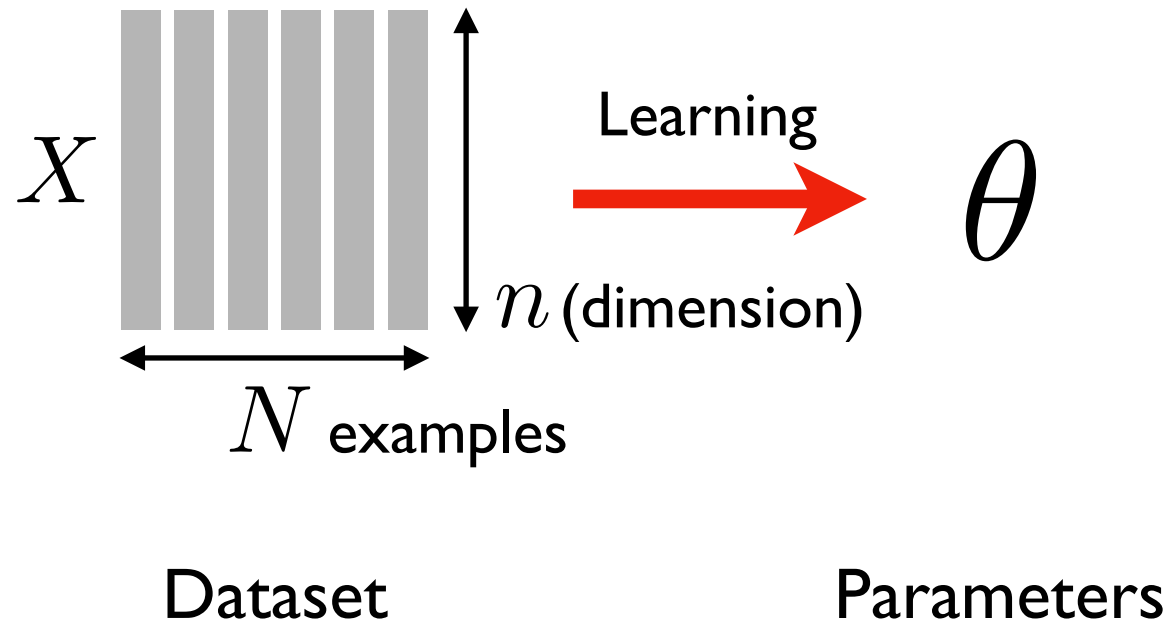
UCLouvain



Compressive Learning of Generative Networks

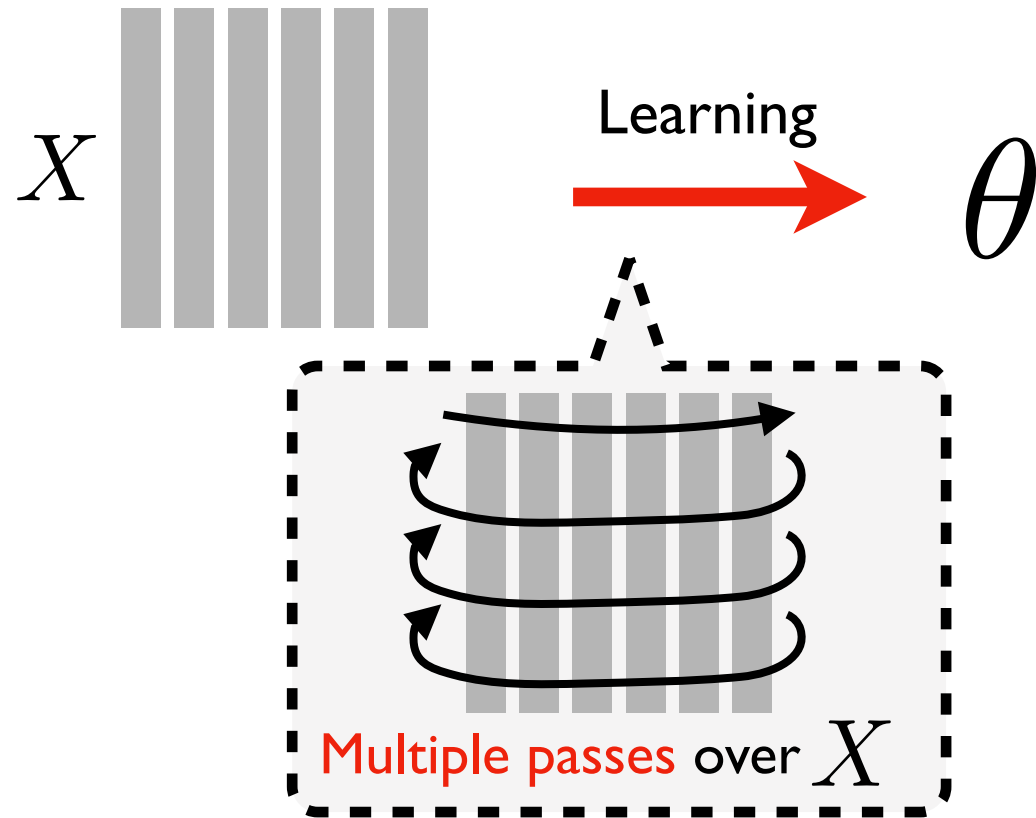
Compressive Learning: Why?

Usual machine learning



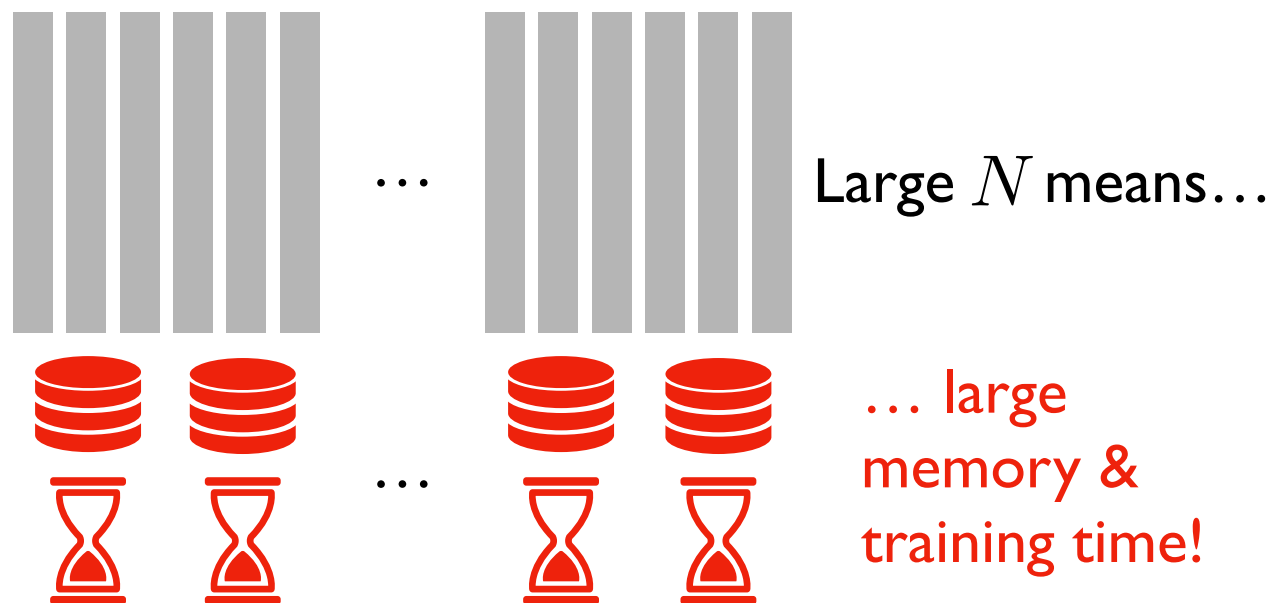
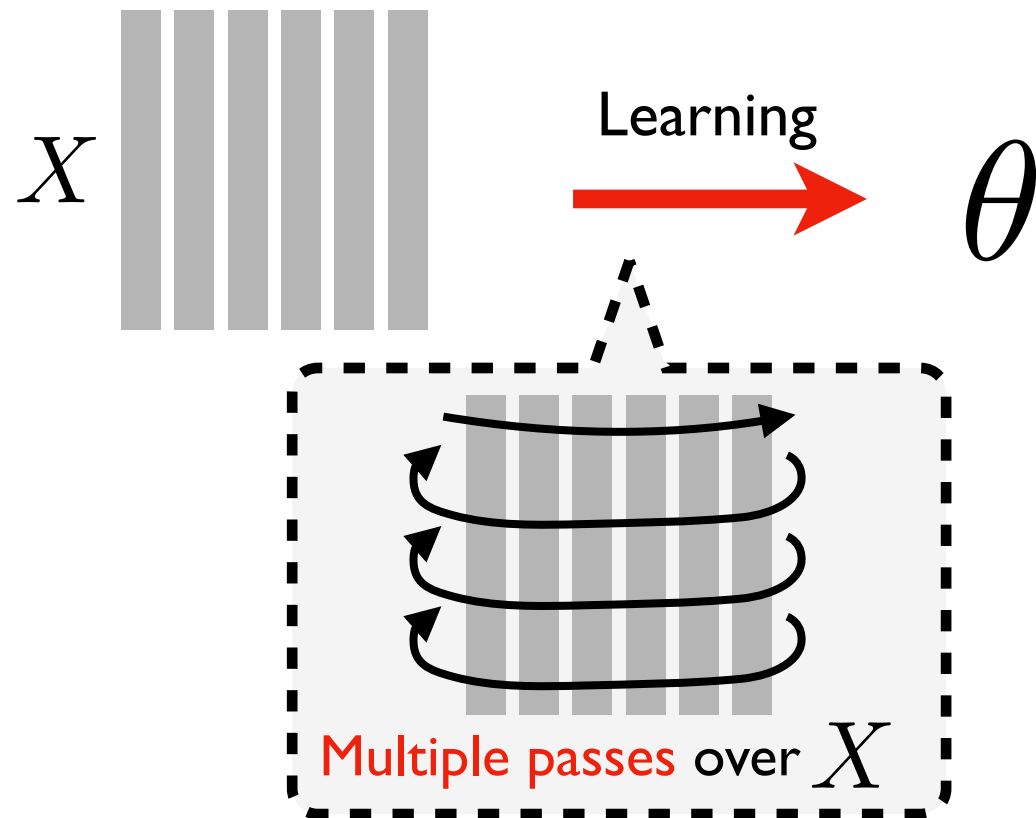
Compressive Learning: Why?

Usual machine learning



Compressive Learning: Why?

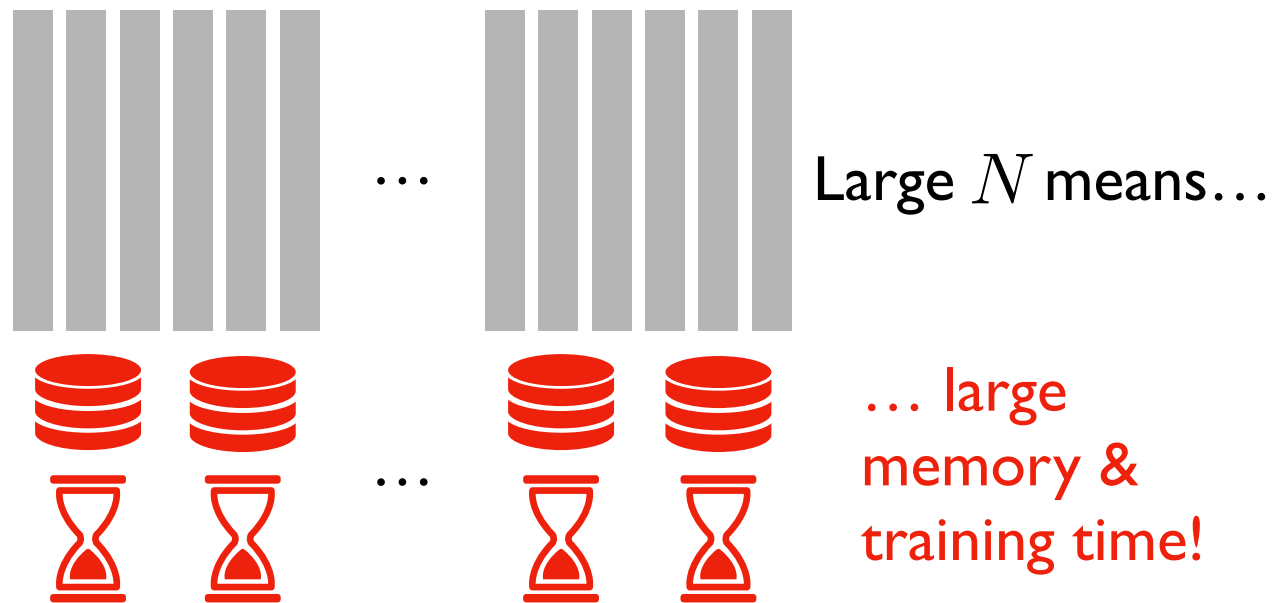
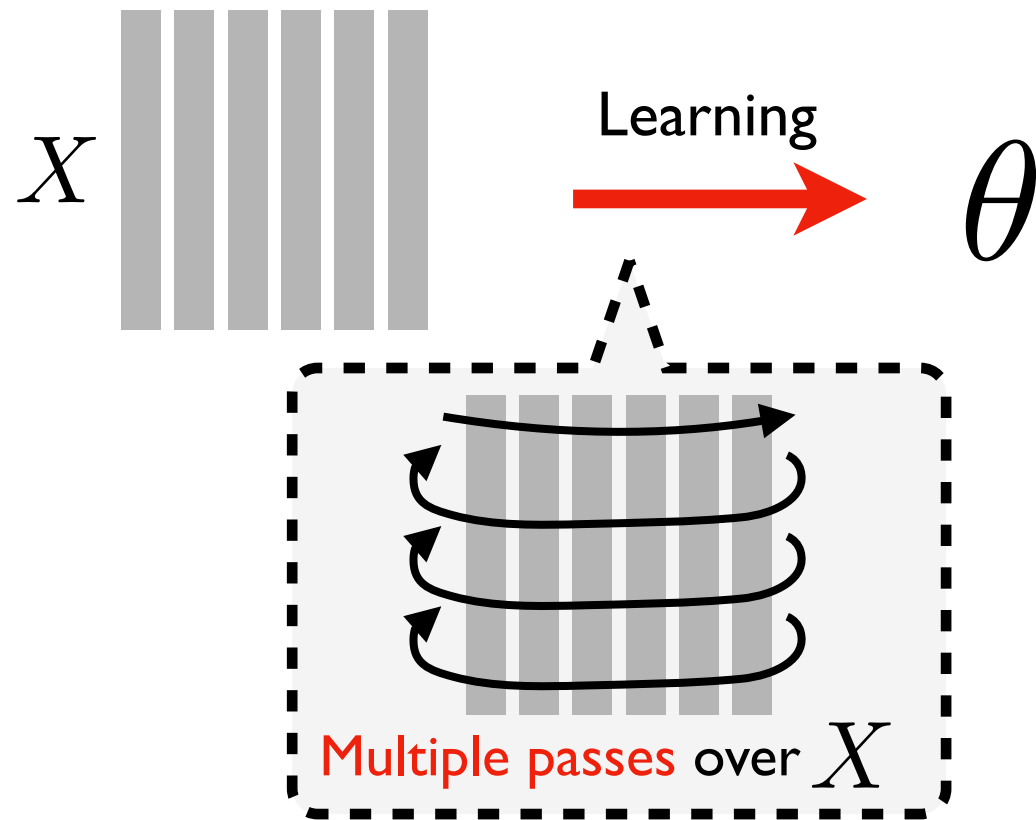
Usual machine learning



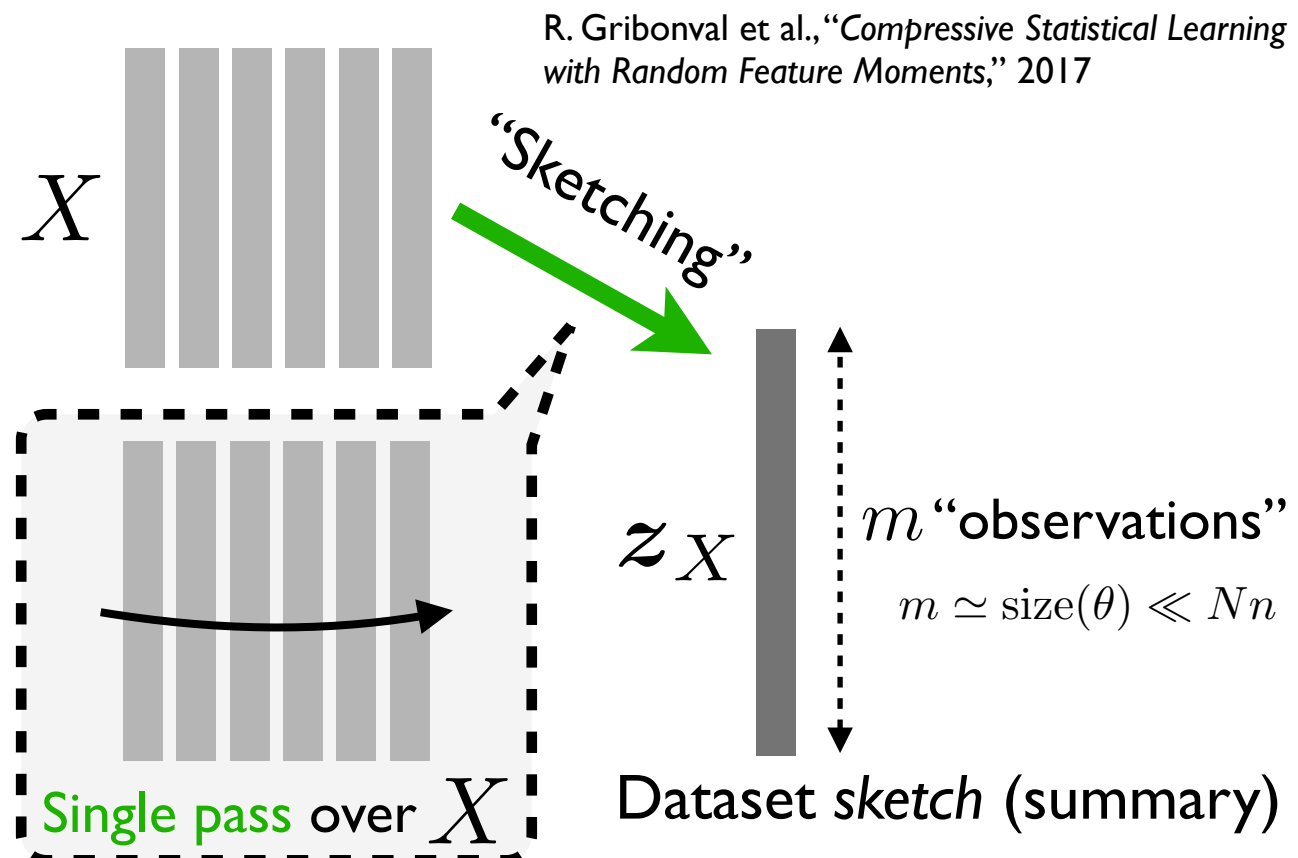
Context: large-scale learning

Compressive Learning: What?

Usual machine learning

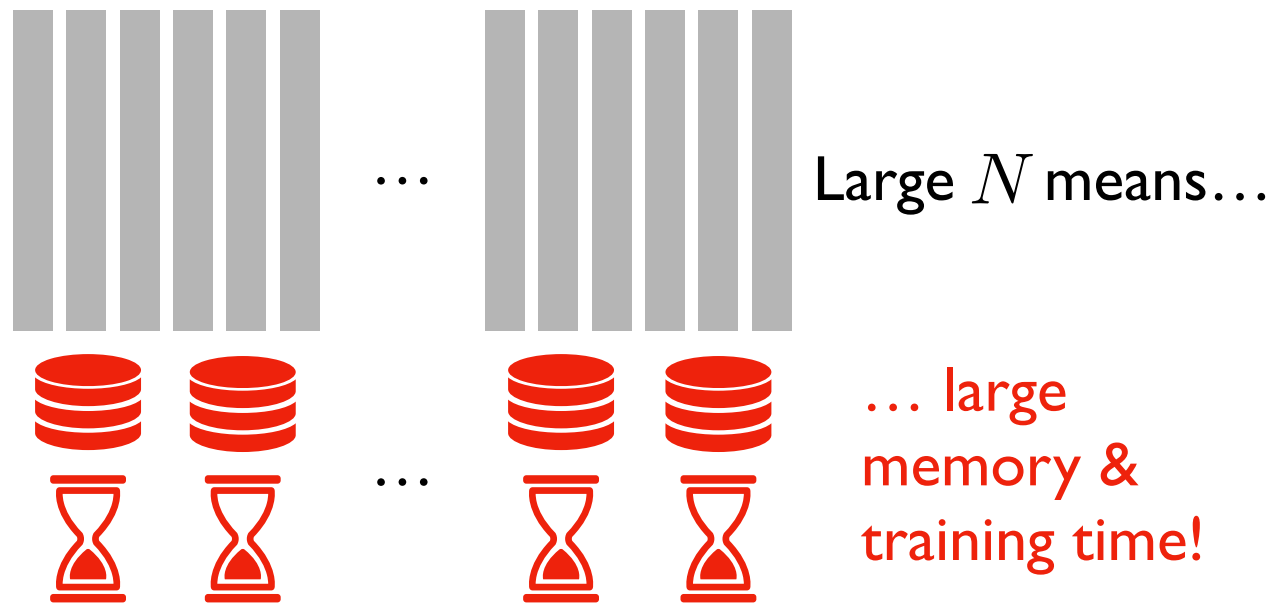
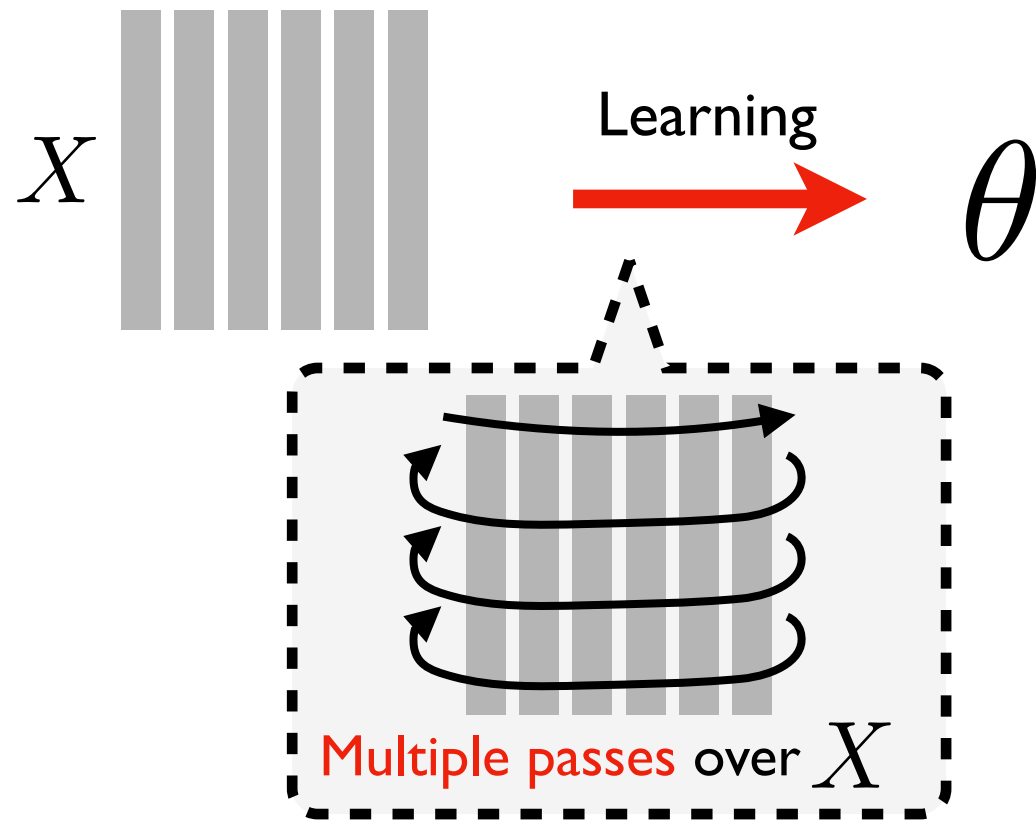


Compressive Learning



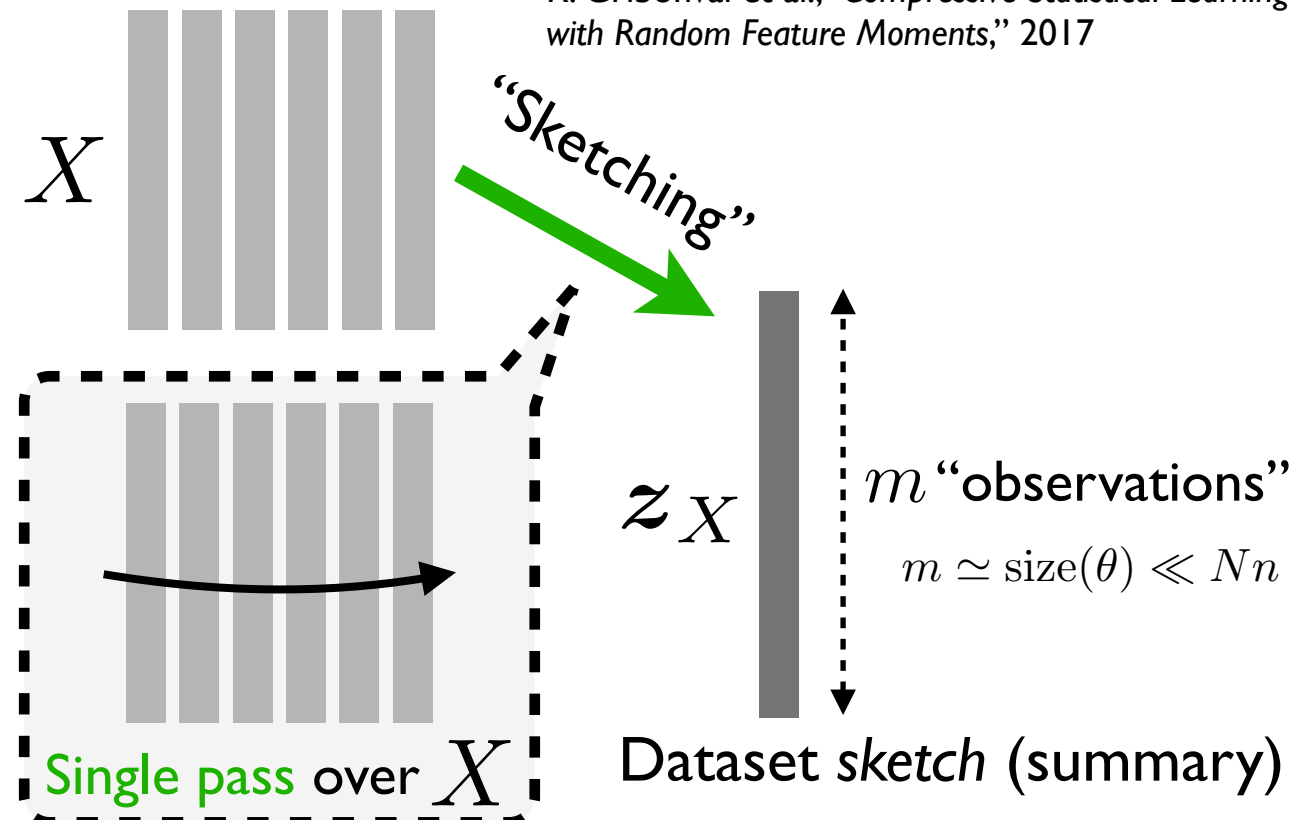
Compressive Learning: What?

Usual machine learning



Compressive Learning

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

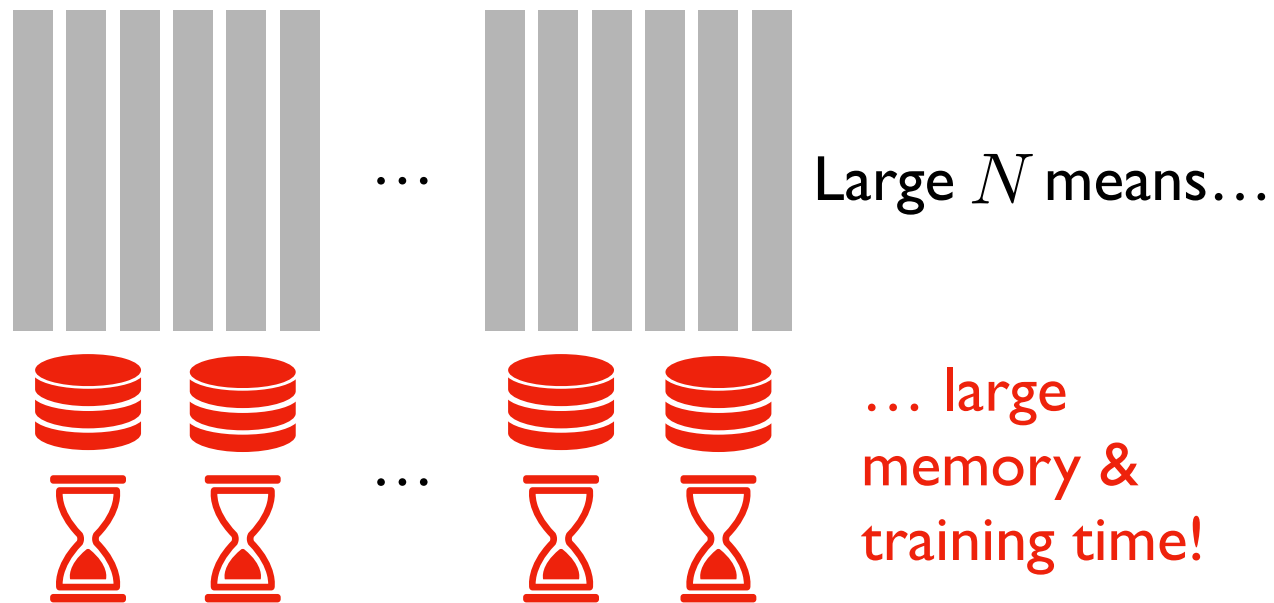
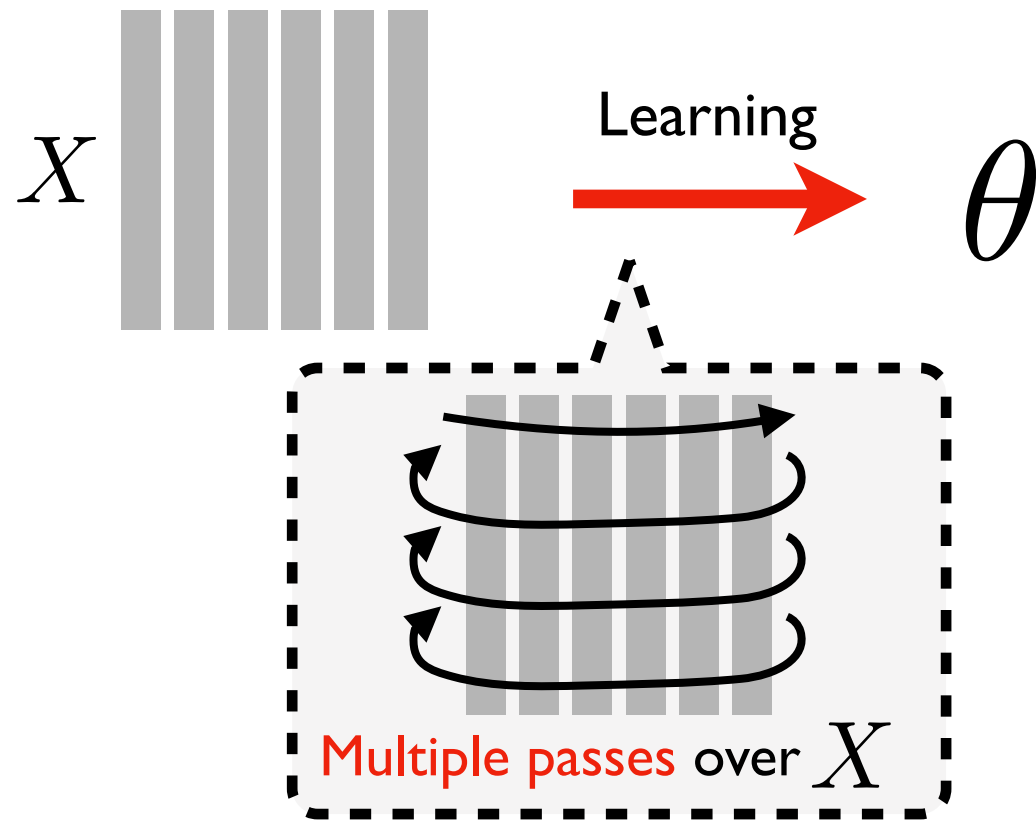


$$z_X = \frac{1}{N} \sum_{i=1}^N \Phi(\mathbf{x}_i)$$

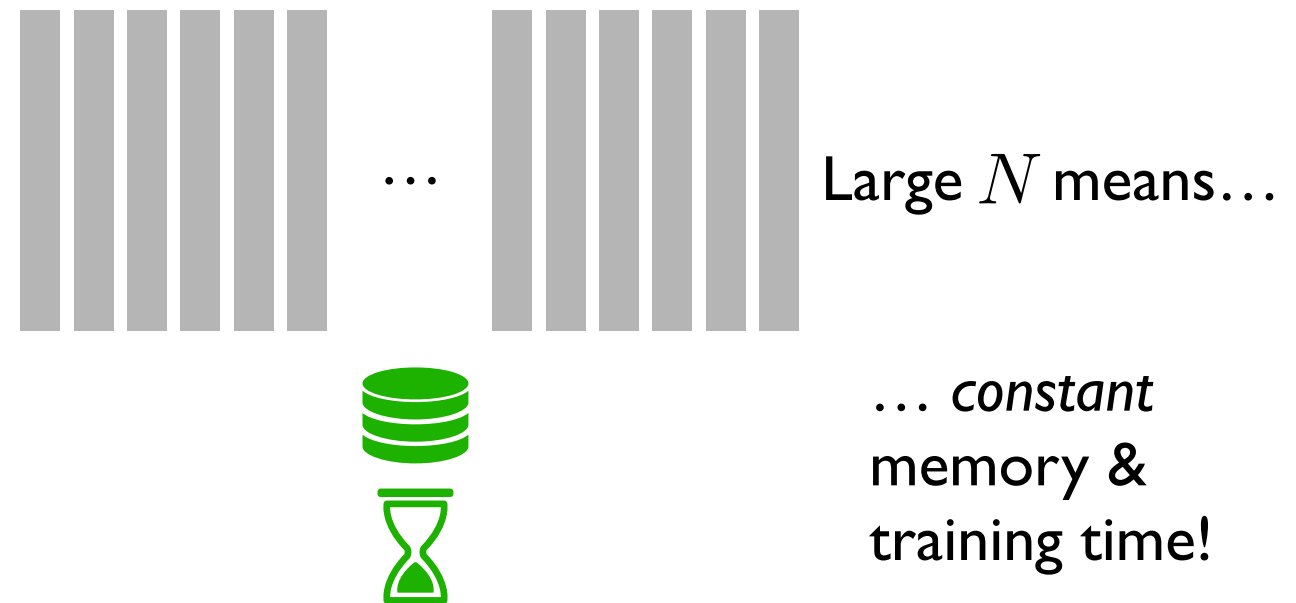
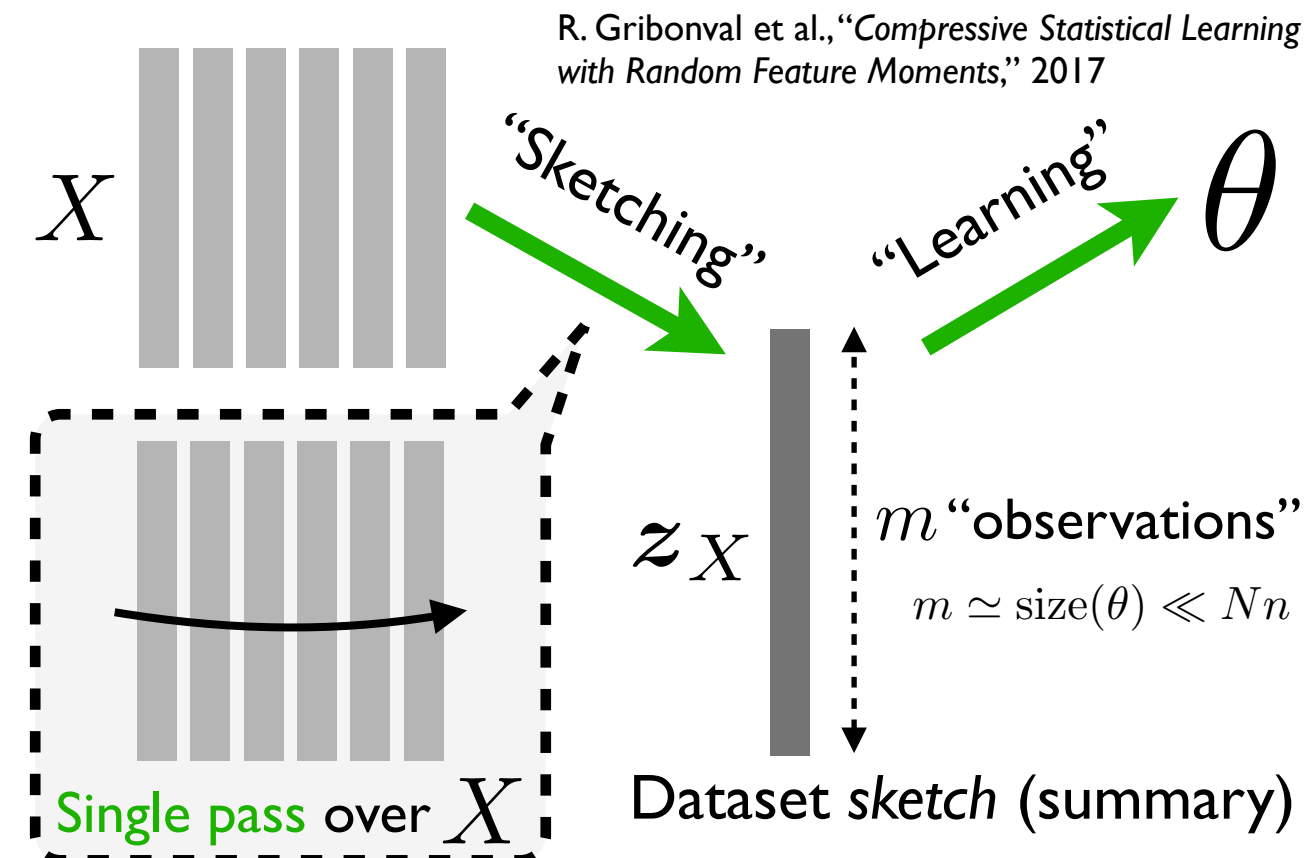
"Generalized moments"
(average of features of the examples)

Compressive Learning: What?

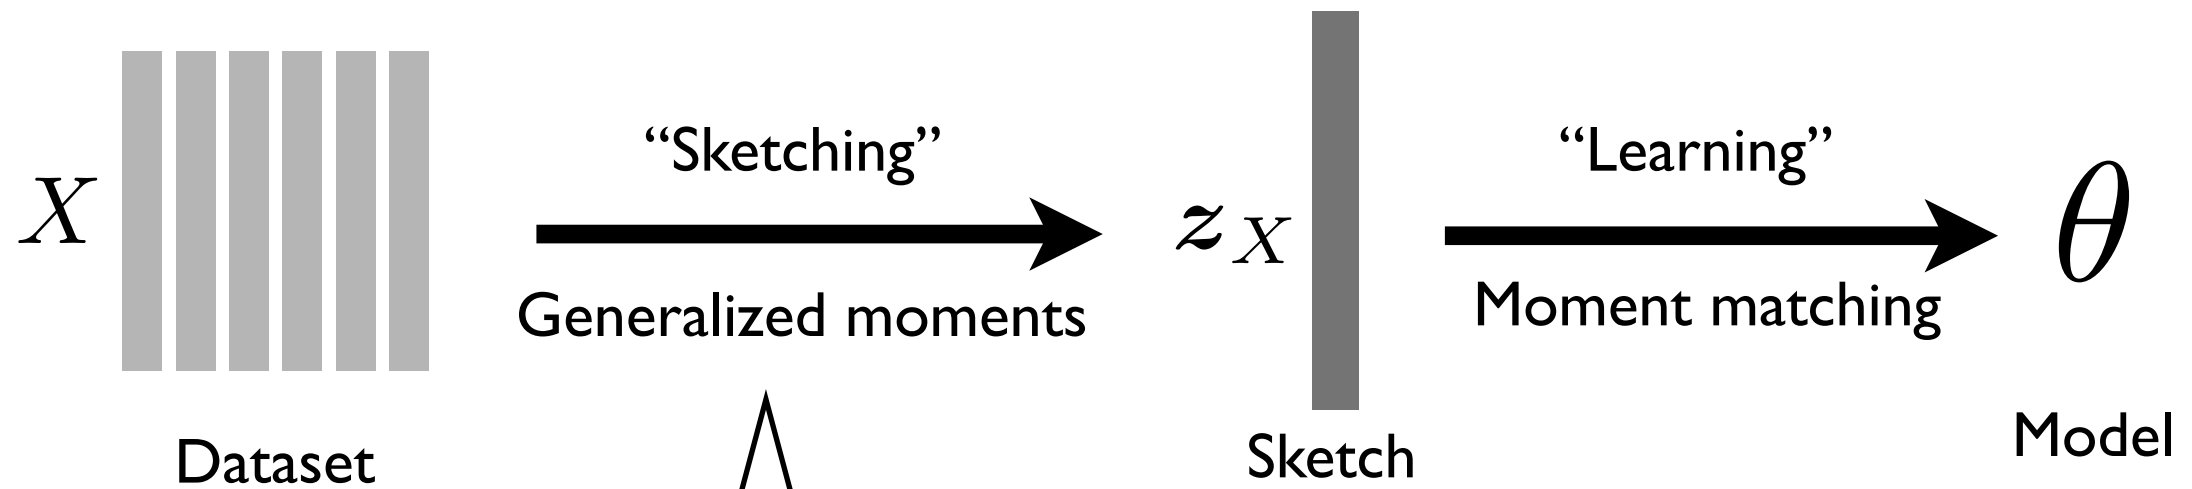
Usual machine learning



Compressive Learning



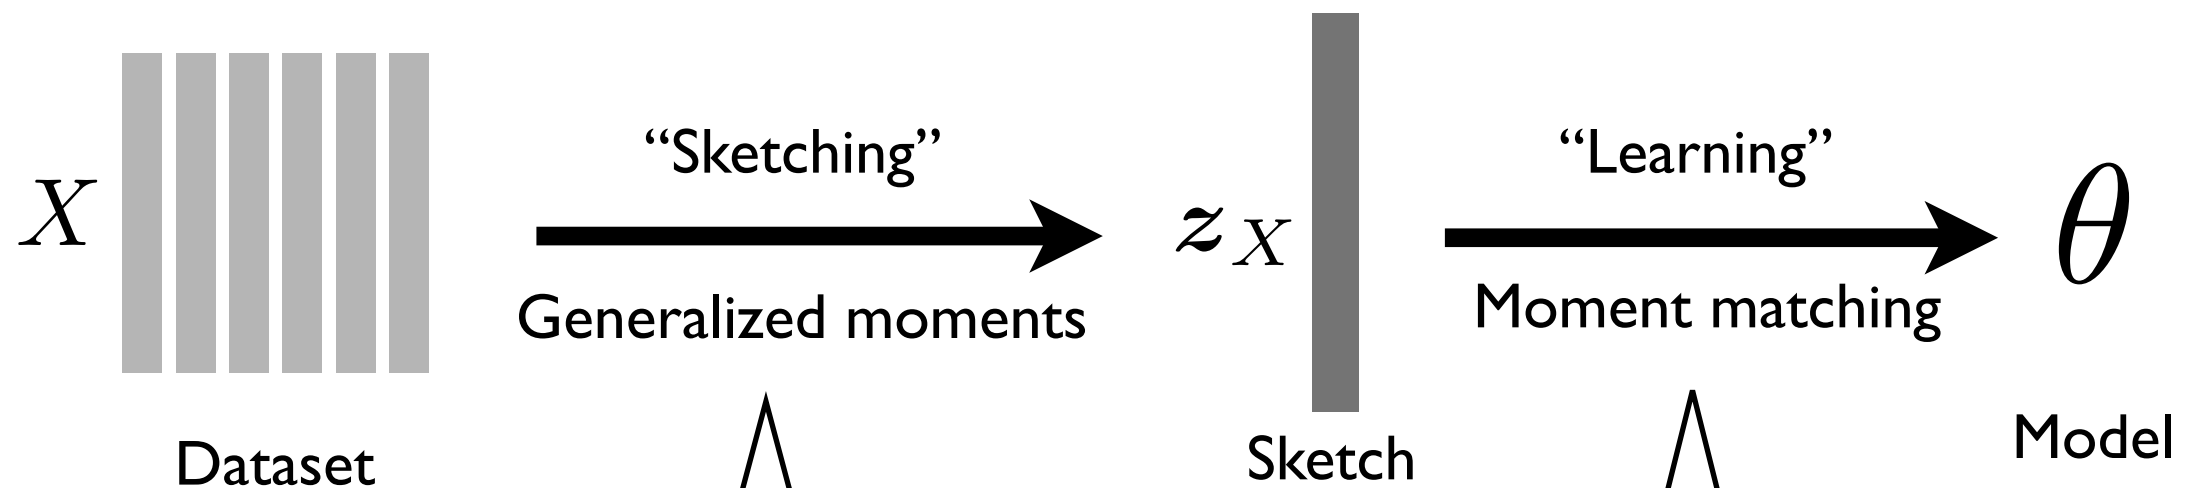
Compressive Learning in a nutshell



Advantages:

- **Harsh compression** (ideal for large-scale datasets, data streams)
- Only a **single pass on dataset** (parallelizable, distributable)
- Handy for **privacy preservation**

Compressive Learning in a nutshell



Advantages:

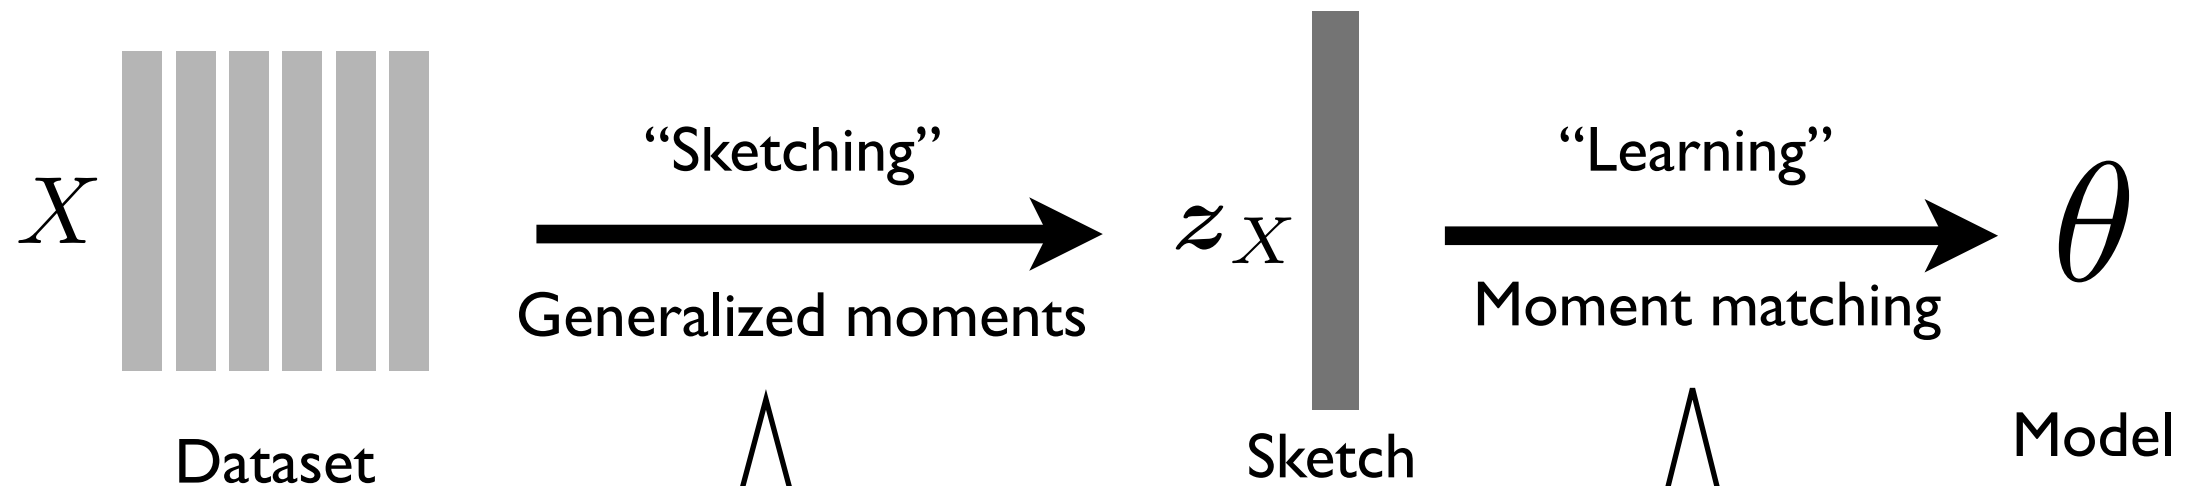
- **Harsh compression** (ideal for large-scale datasets, data streams)
- Only a **single pass on dataset** (parallelizable, distributable)
- Handy for **privacy preservation**

Existing algorithms for:

- k-means clustering
- mixture model estimation: GMM, alpha-stable distributions
- Principal Component Analysis
- Independent Component Analysis

Difficult to extend the framework to new tasks!

Compressive Learning in a nutshell



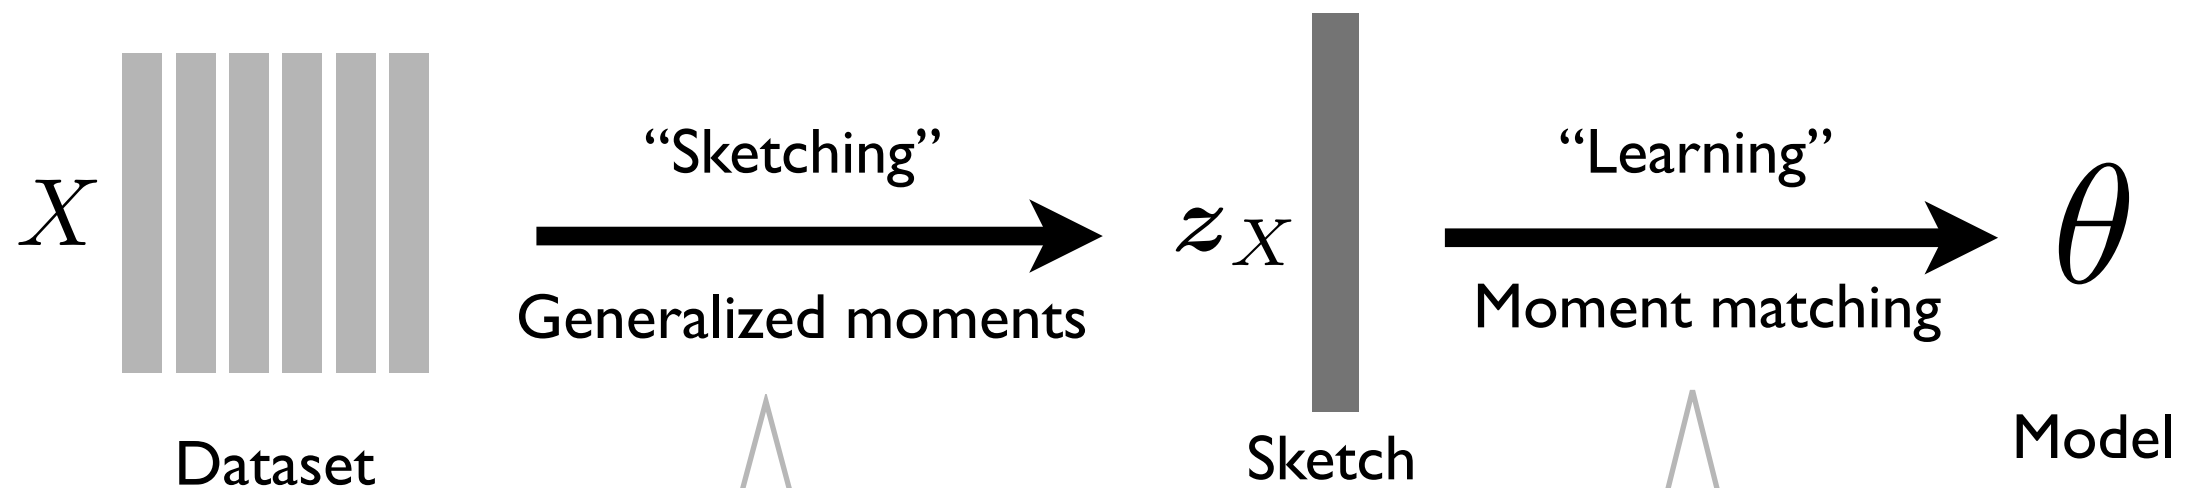
Advantages:

- **Harsh compression** (ideal for large-scale datasets, data streams)
- Only a **single pass on dataset** (parallelizable, distributable)
- Handy for **privacy preservation**

Existing algorithms for:

- k-means clustering
- mixture model estimation: GMM, alpha-stable distributions
- Principal Component Analysis
- Independent Component Analysis
- **This work: Generative NN**

Compressive Learning in a nutshell



Advantages:

- **Harsh compression** (ideal for large-scale datasets, data streams)
- Only a **single pass on dataset** (parallelizable, distributable)
- Handy for **privacy preservation**

Existing algorithms for:

- k-means clustering
- mixture model estimation: GMM, alpha-stable distributions
- Principal Component Analysis
- Independent Component Analysis
- **This work: Generative NN**

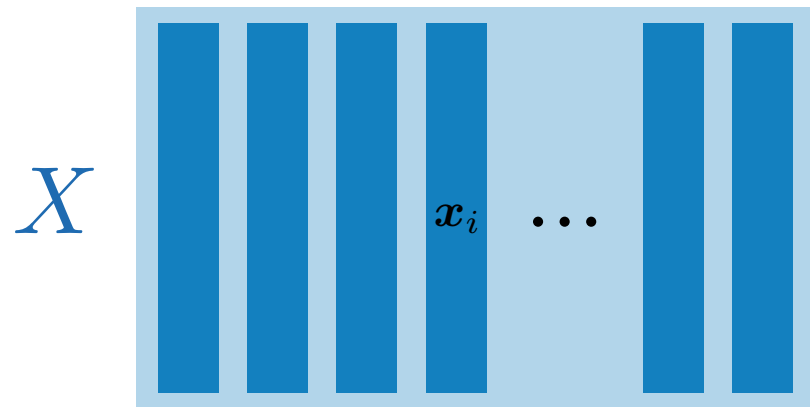
Contributions:

- Connecting literature (Compressive Learning & Generative Networks)
- Practical learning algorithm
- Qualitative proof-of-concept (toy example)
- Quantitative experimental validation (real data)

Compressive Learning of
Generative Networks

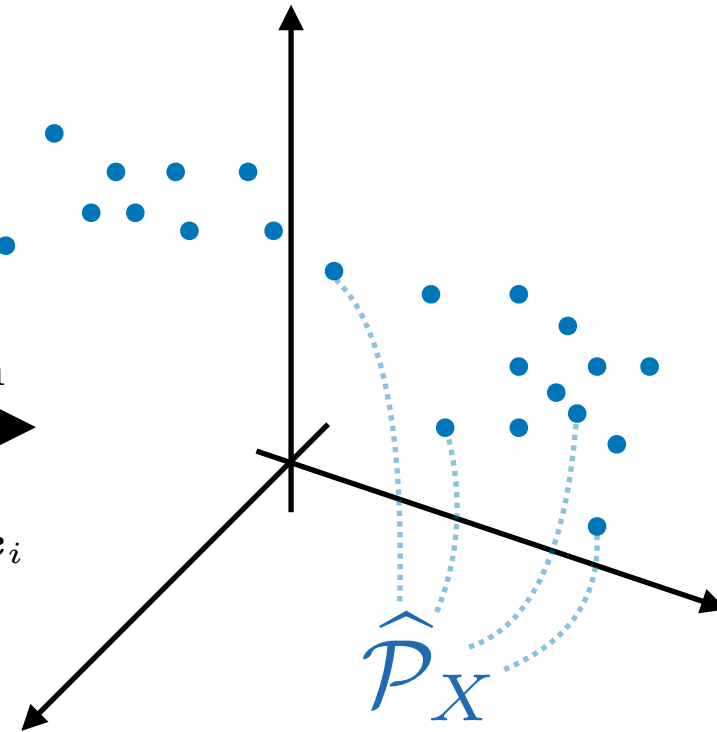
Generative Networks: What?

Learn a generative network \mathcal{G}_θ



Empirical distribution

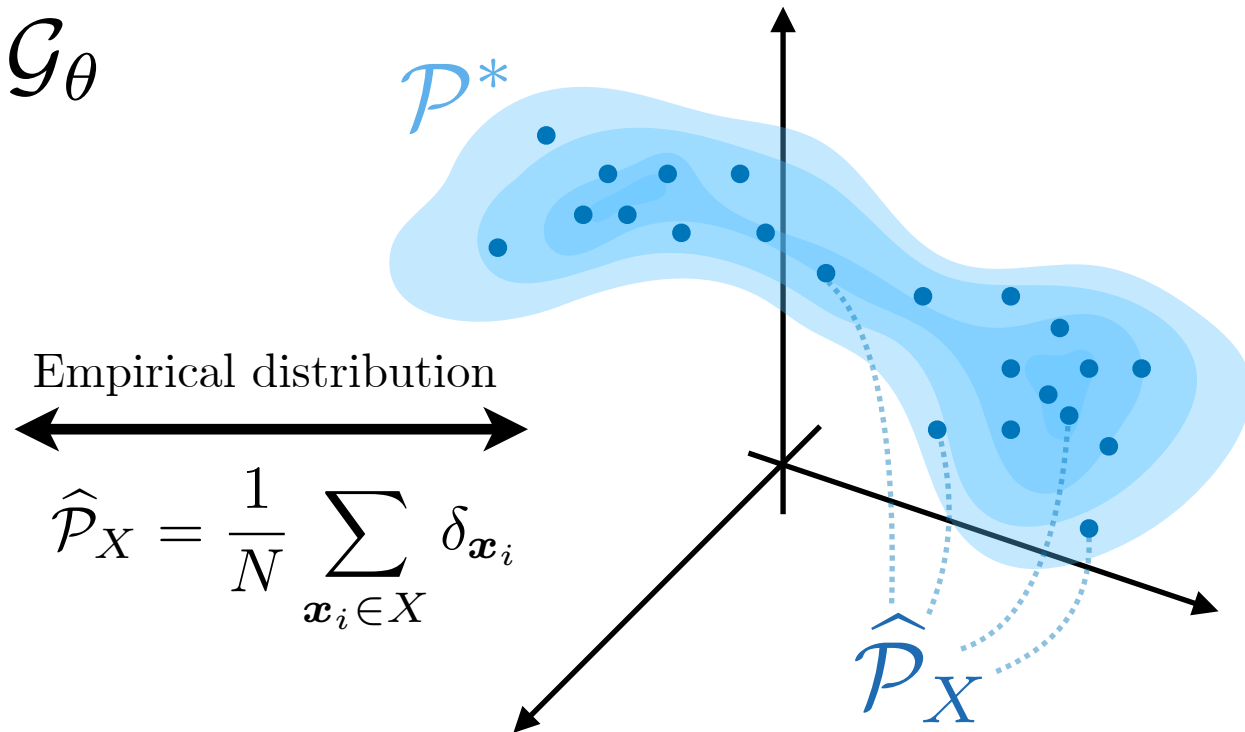
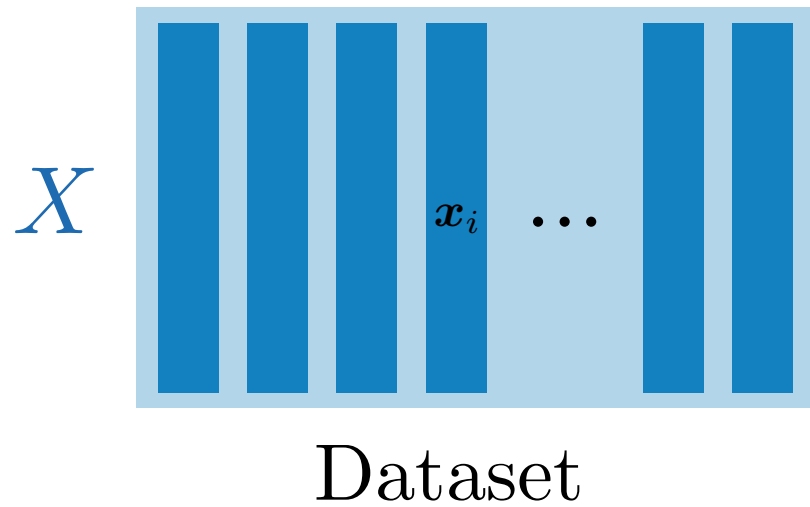
$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{\mathbf{x}_i \in X} \delta_{\mathbf{x}_i}$$



We have **samples** (signals)

Generative Networks: What?

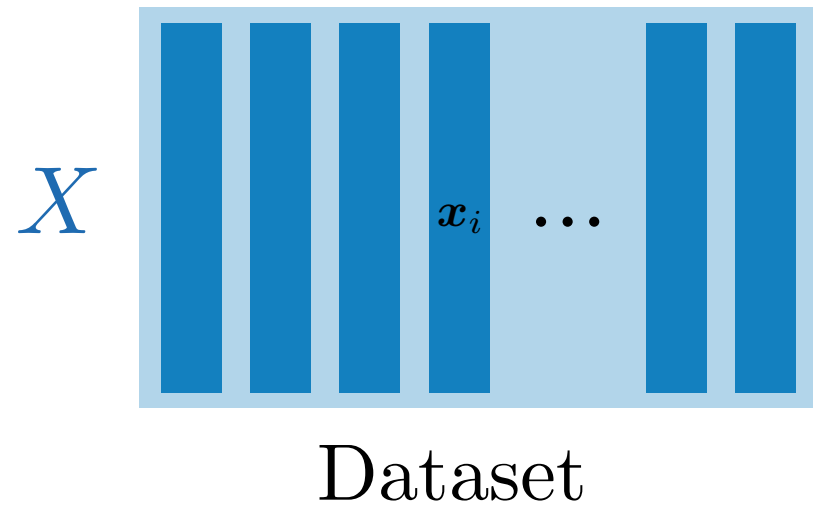
Learn a generative network \mathcal{G}_θ



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

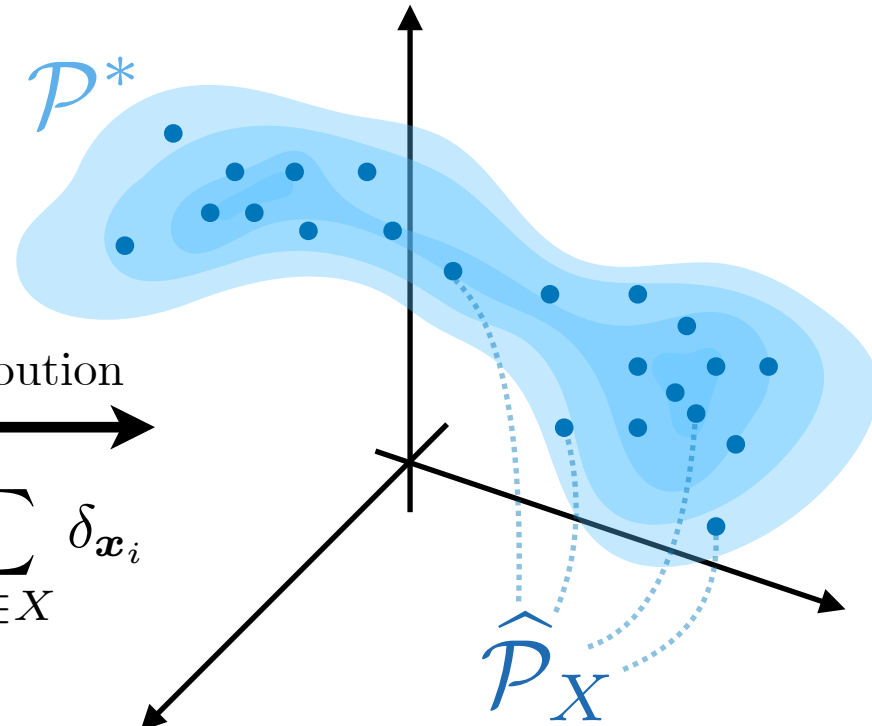
Generative Networks: What?

Learn a generative network \mathcal{G}_θ



Empirical distribution

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

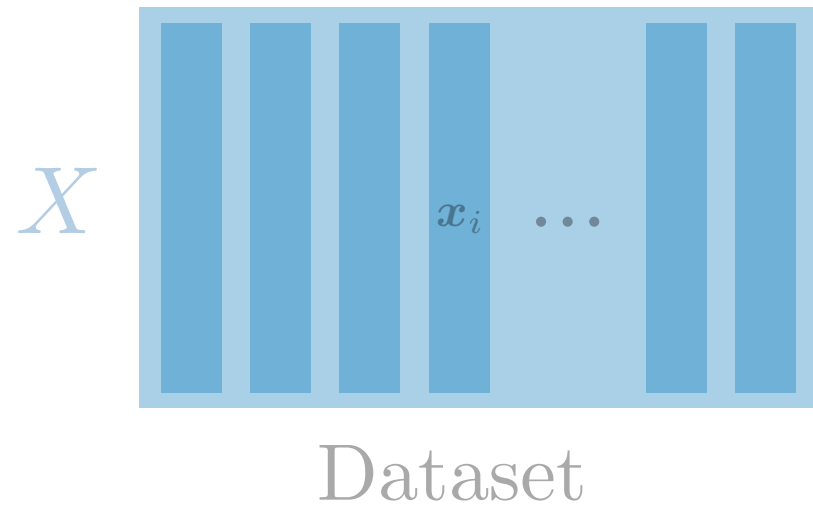


We have **samples** (signals) from a high-dimensional **distribution**...

Idea: approximate true prior distribution from learning examples

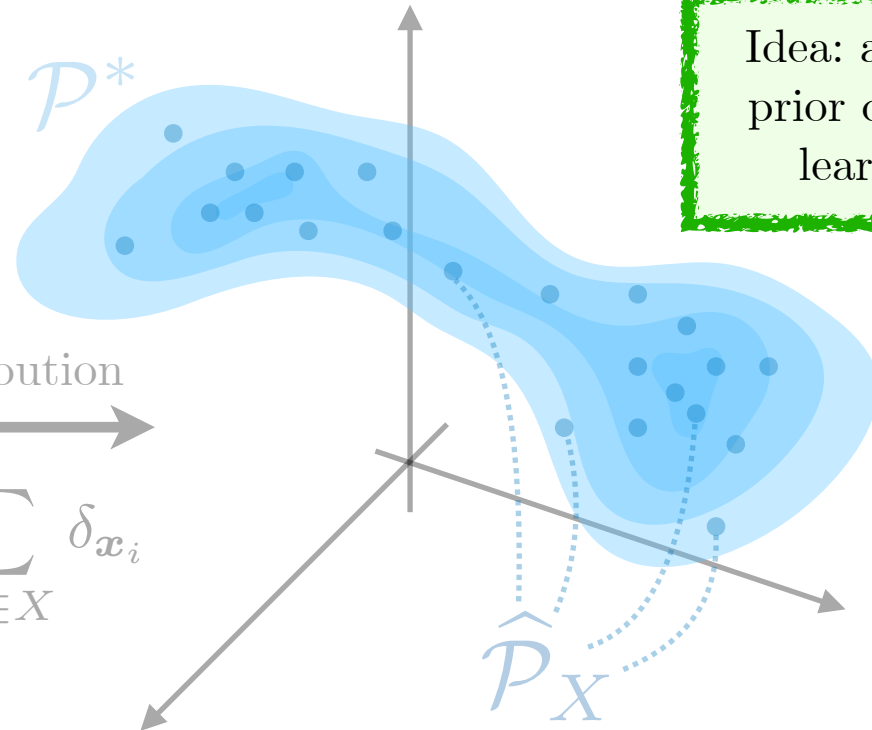
Generative Networks: What?

Learn a generative network \mathcal{G}_θ



Empirical distribution

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

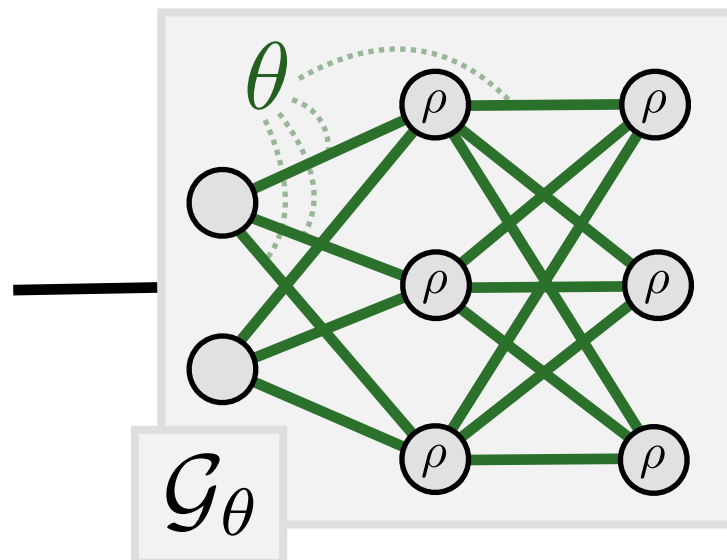
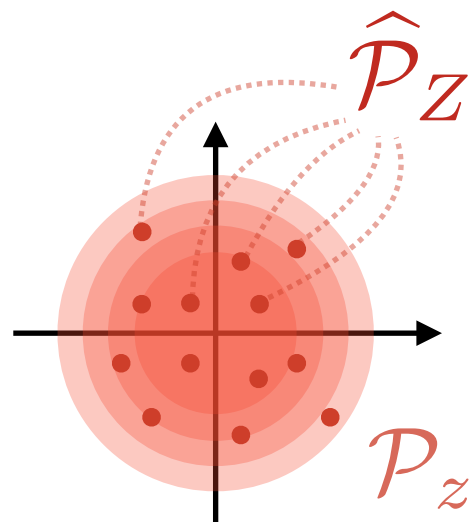


Idea: approximate true prior distribution from learning examples

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

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

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

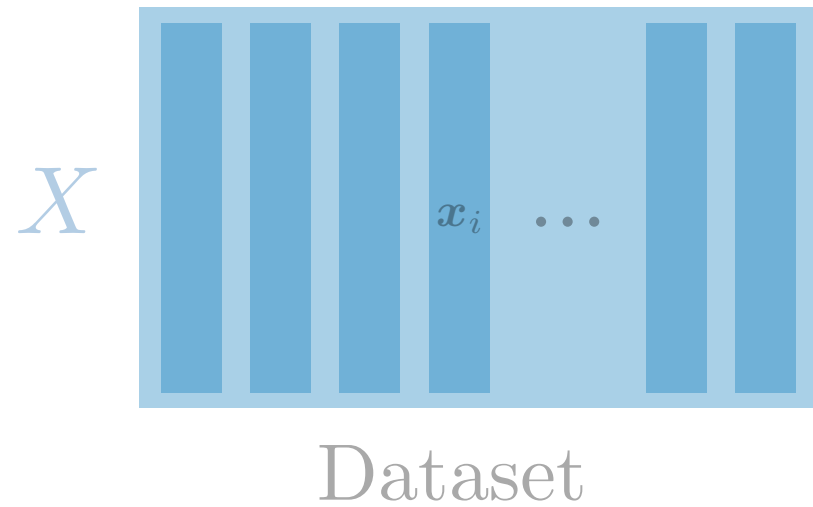


Latent space
Random “noise”

Generative network

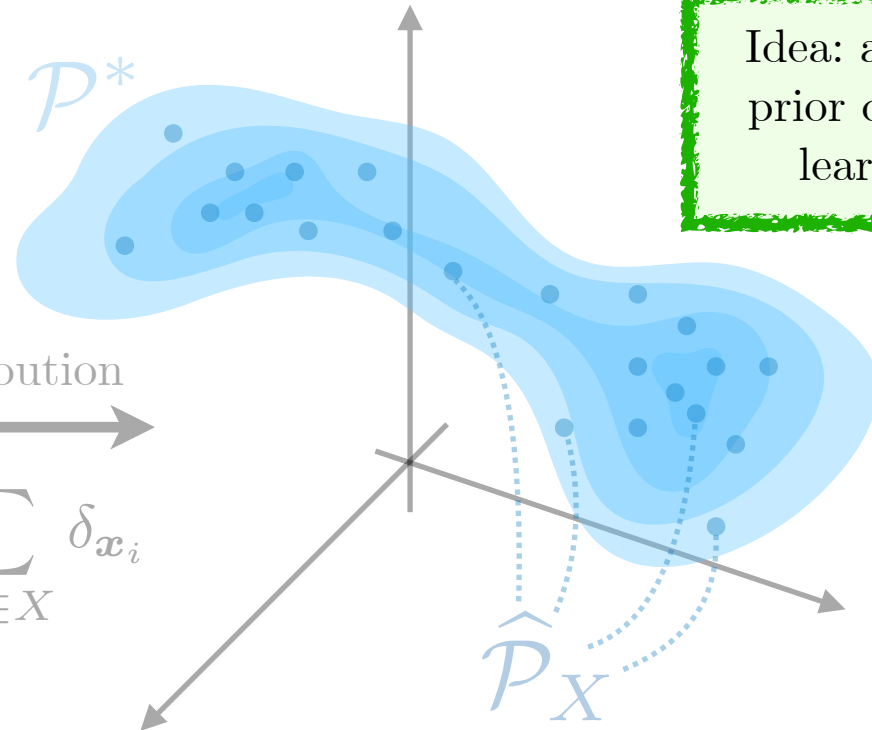
Generative Networks: What?

Learn a generative network \mathcal{G}_θ



Empirical distribution

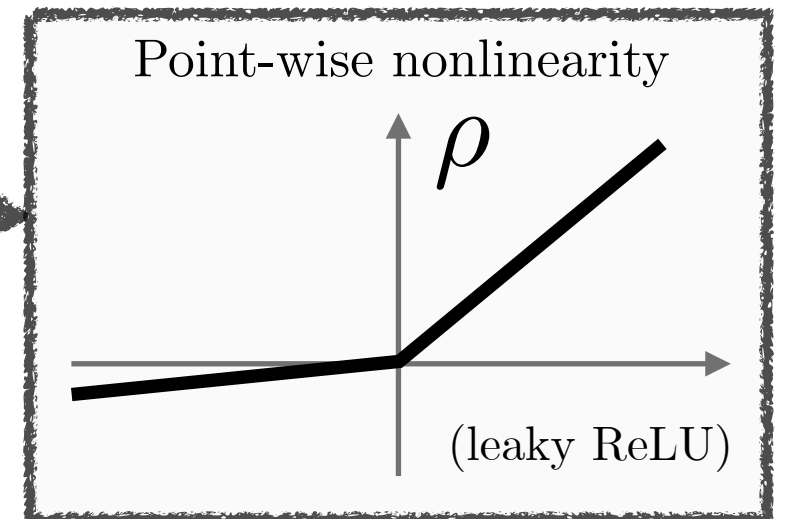
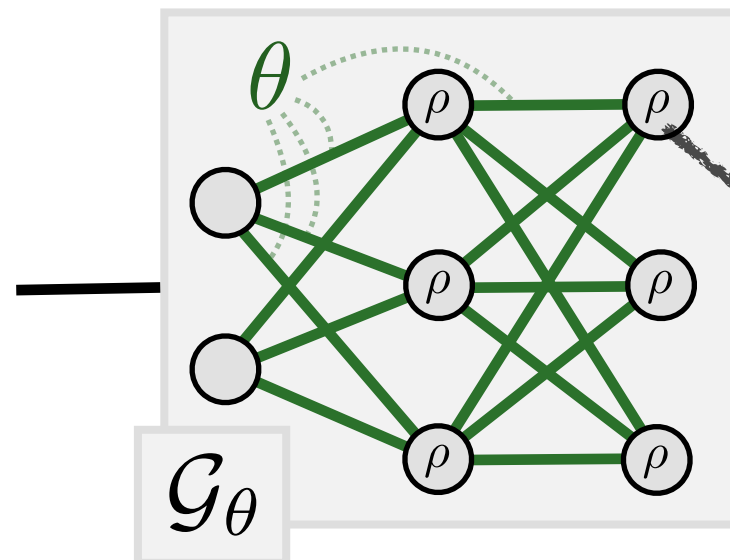
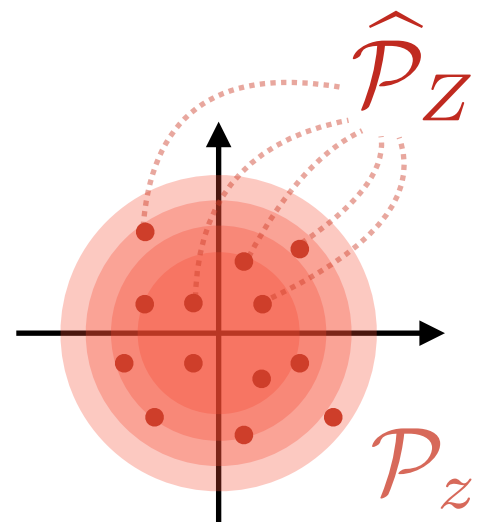
$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{x_i \in X} \delta_{x_i}$$



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

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

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

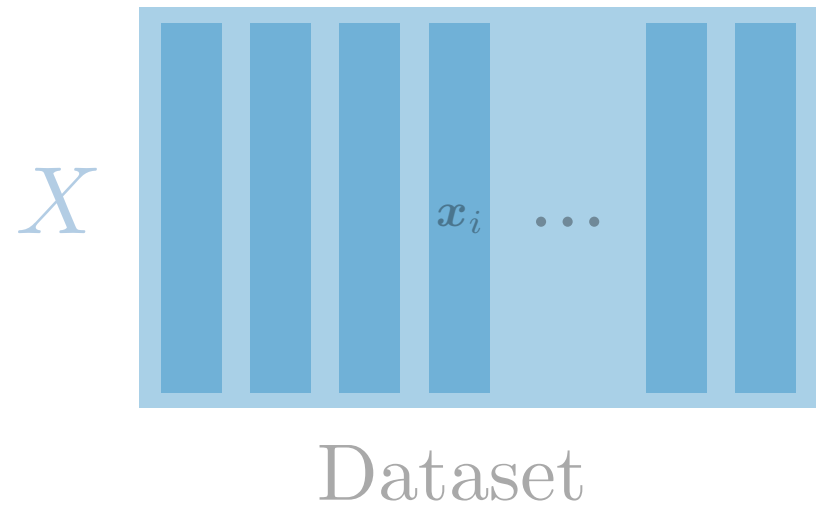


Latent space
Random “noise”

Generative network

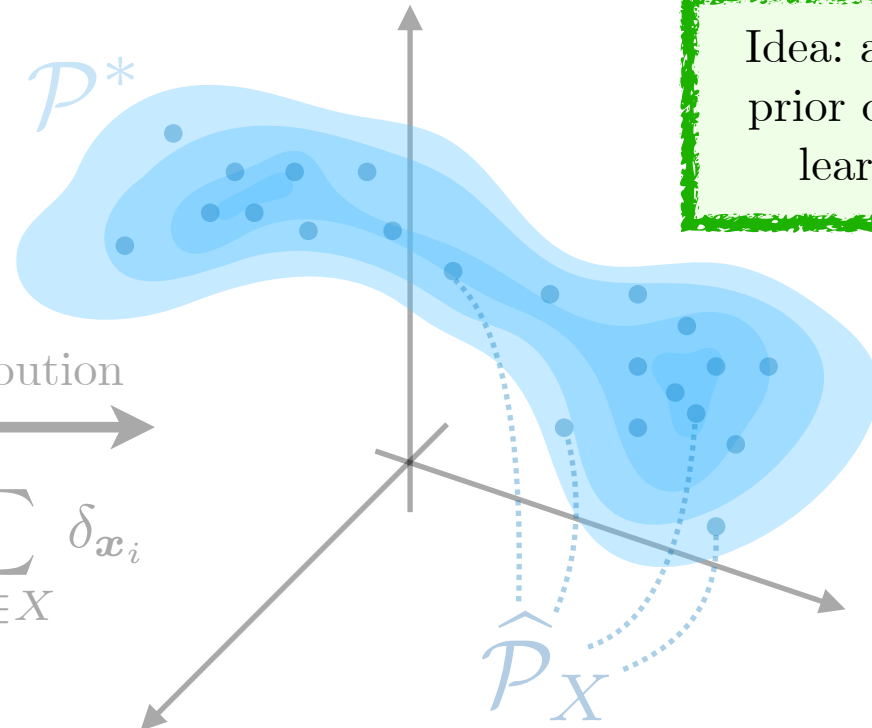
Generative Networks: What?

Learn a generative network \mathcal{G}_θ



Empirical distribution

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

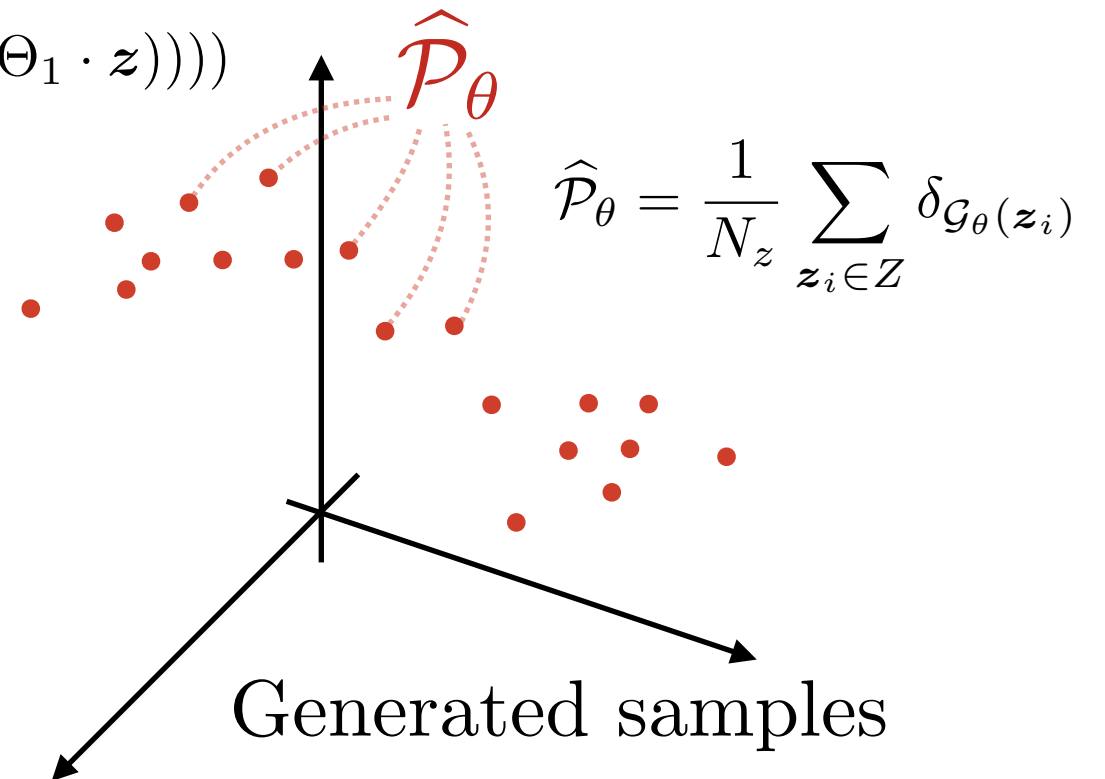
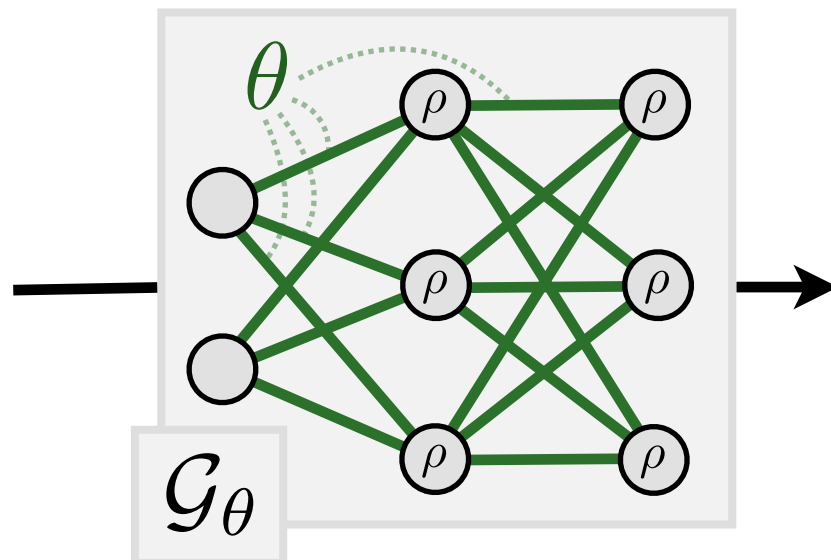
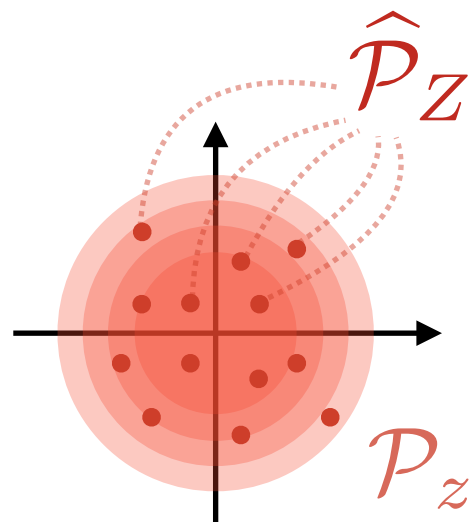


Idea: approximate true prior distribution from learning examples

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

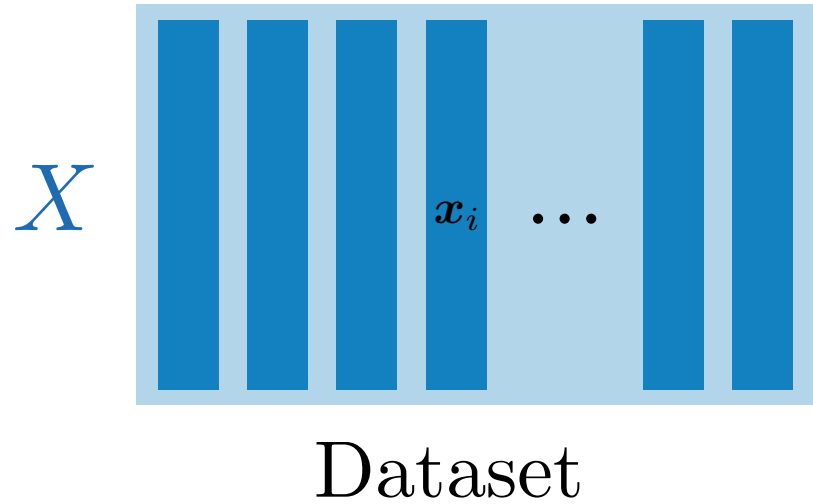
... and a way to generate “artificial” samples...

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



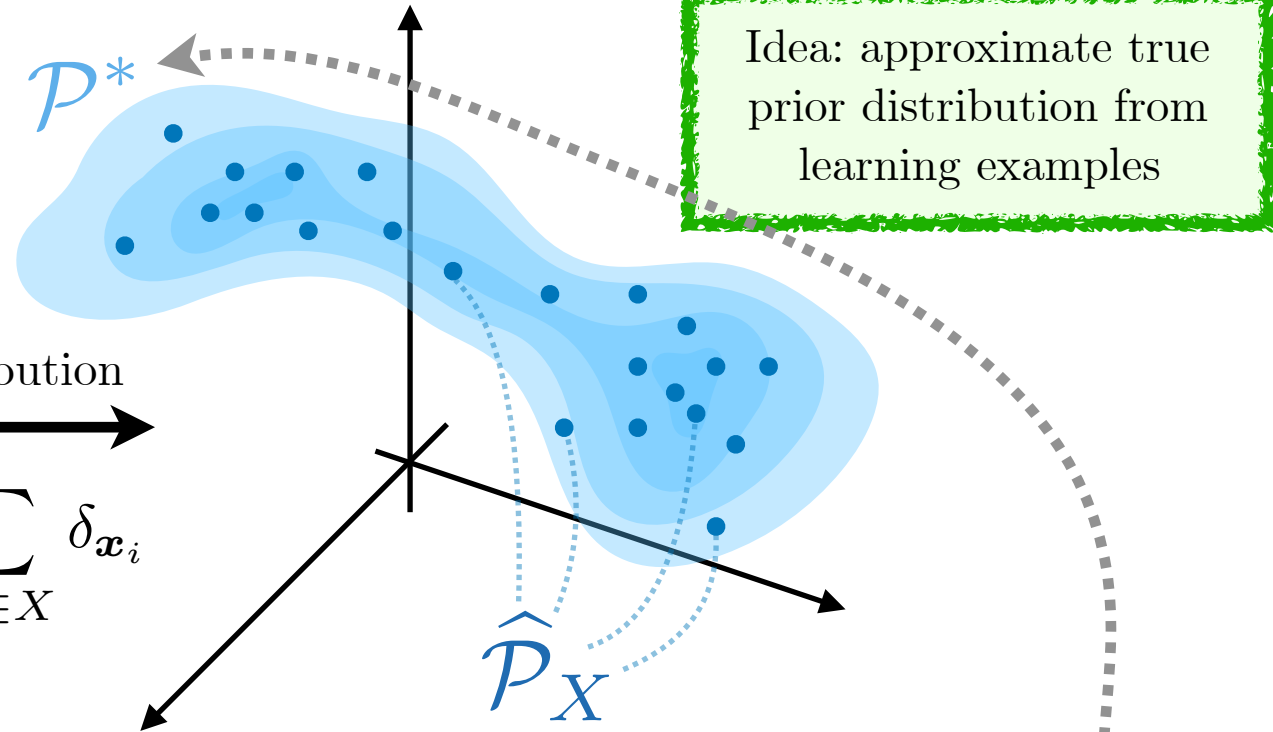
Generative Networks: What?

Learn a generative network \mathcal{G}_θ

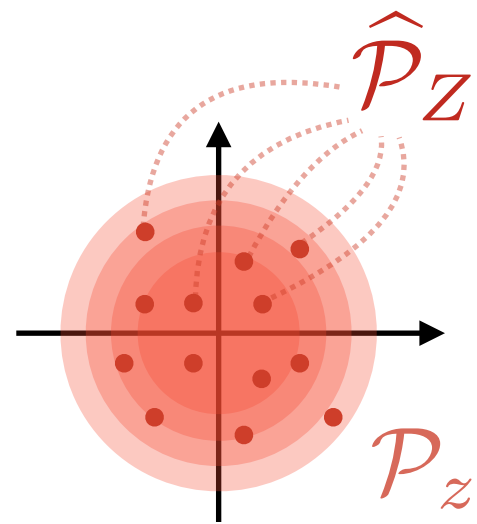


Empirical distribution

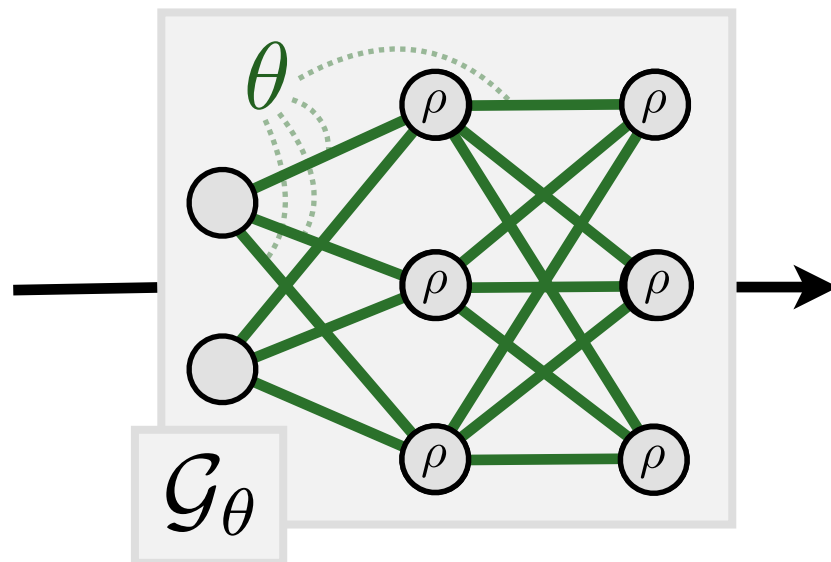
$$\hat{\mathcal{P}}_X = \frac{1}{N} \sum_{\mathbf{x}_i \in X} \delta_{\mathbf{x}_i}$$



Goal: mimic sampling from the data-generating distribution (implicit manifold learning)

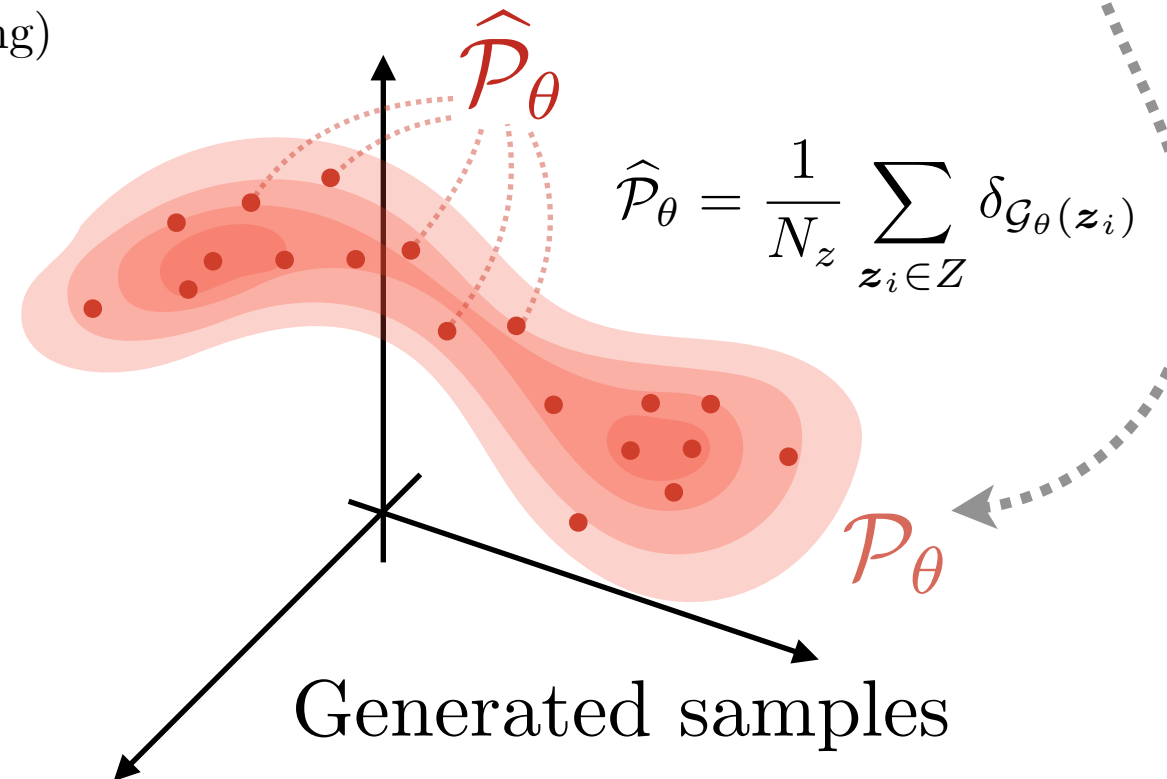


Latent space
Random “noise”



Generative network

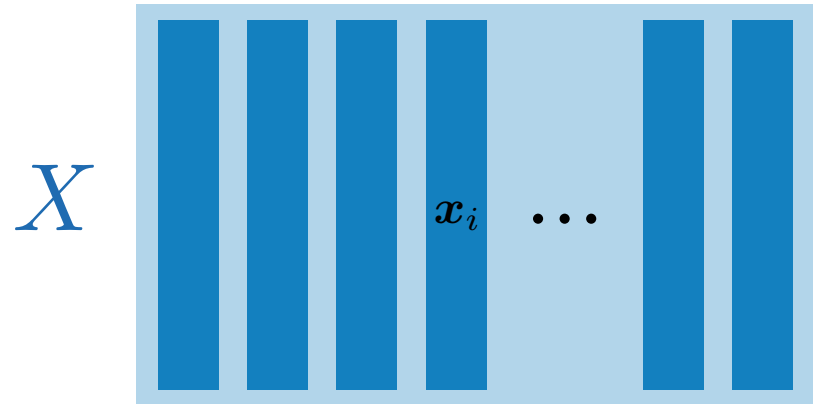
20



Generated samples

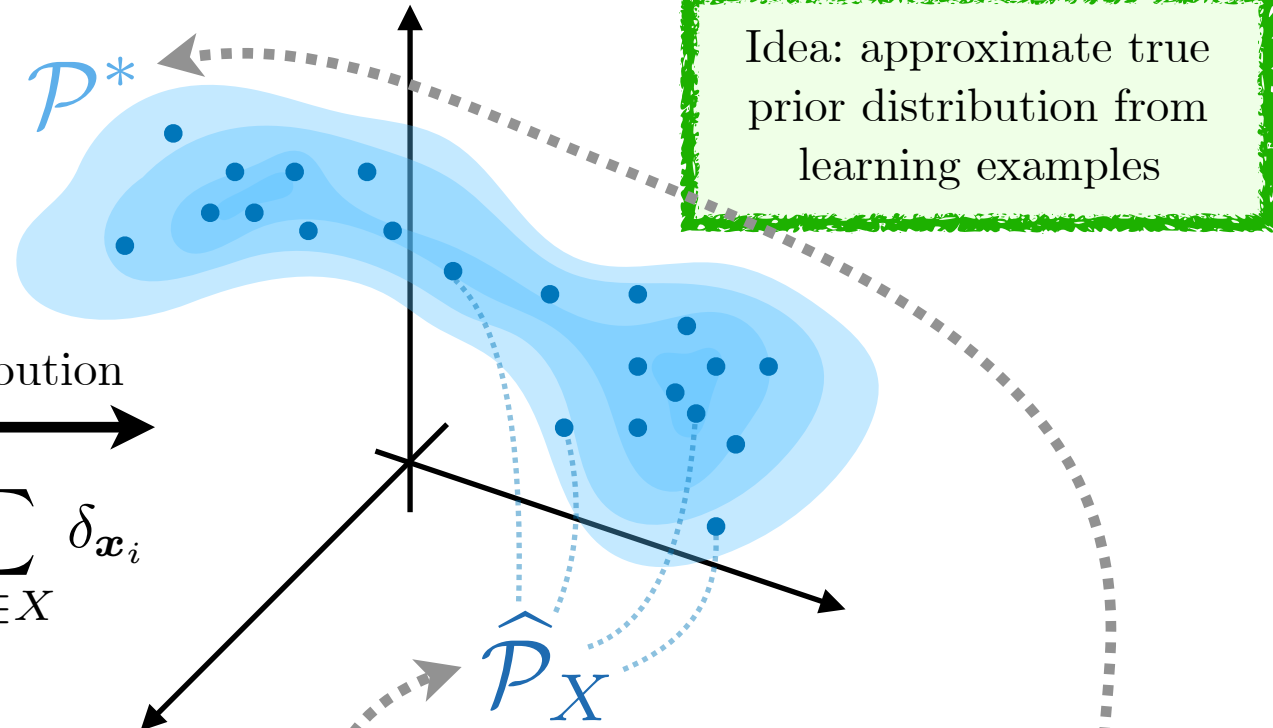
Generative Networks: What?

Learn a generative network \mathcal{G}_θ

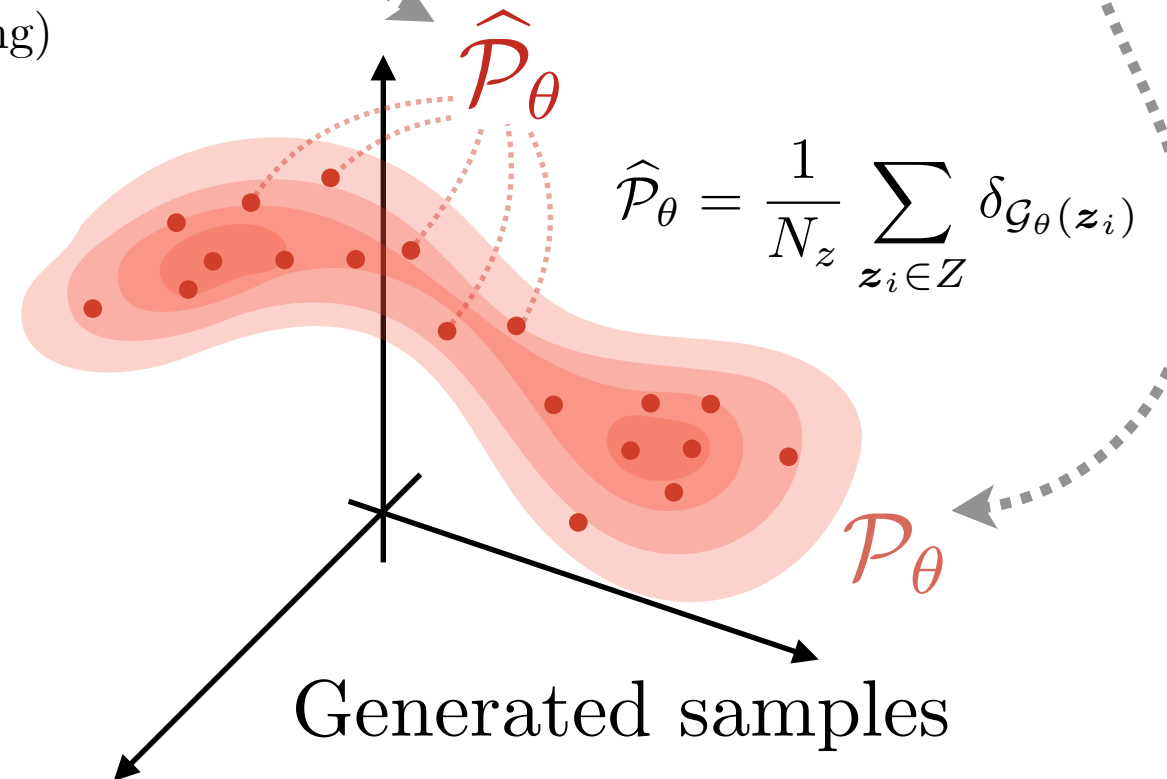
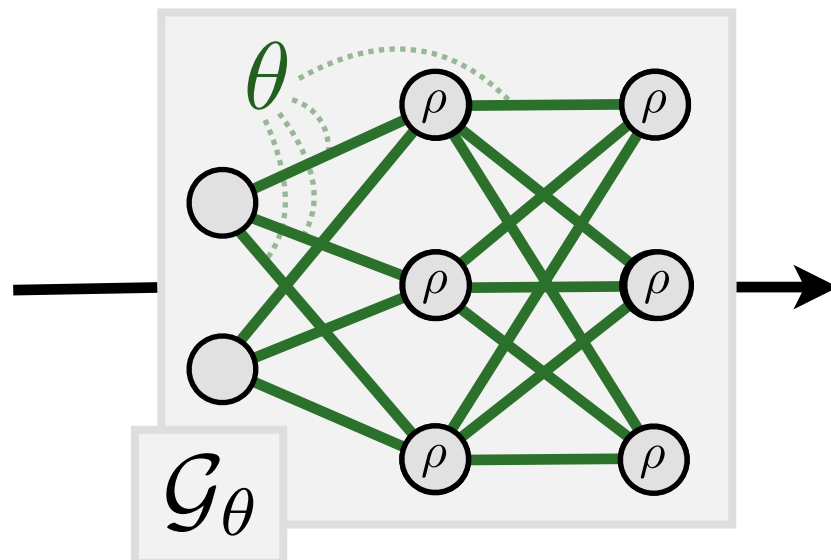
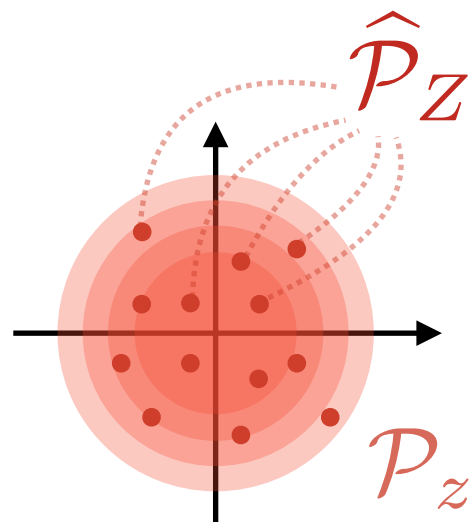


Empirical distribution

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



Goal: mimic sampling from the data-generating distribution (implicit manifold learning)



Latent space
Random “noise”

Generative Networks: Why?

Motivation (amongst others): priors in inverse problems

SUNLayer: Stable denoising with generative networks

Dustin G. Mixon* Soledad Villar†

Abstract

It has been experimentally established that deep neural networks are good priors 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 focus on the problem of image denoising. 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] (referred to as empirical risk minimization) in the compressed sensing task by using a generative network with random weights and ReLU activation functions.

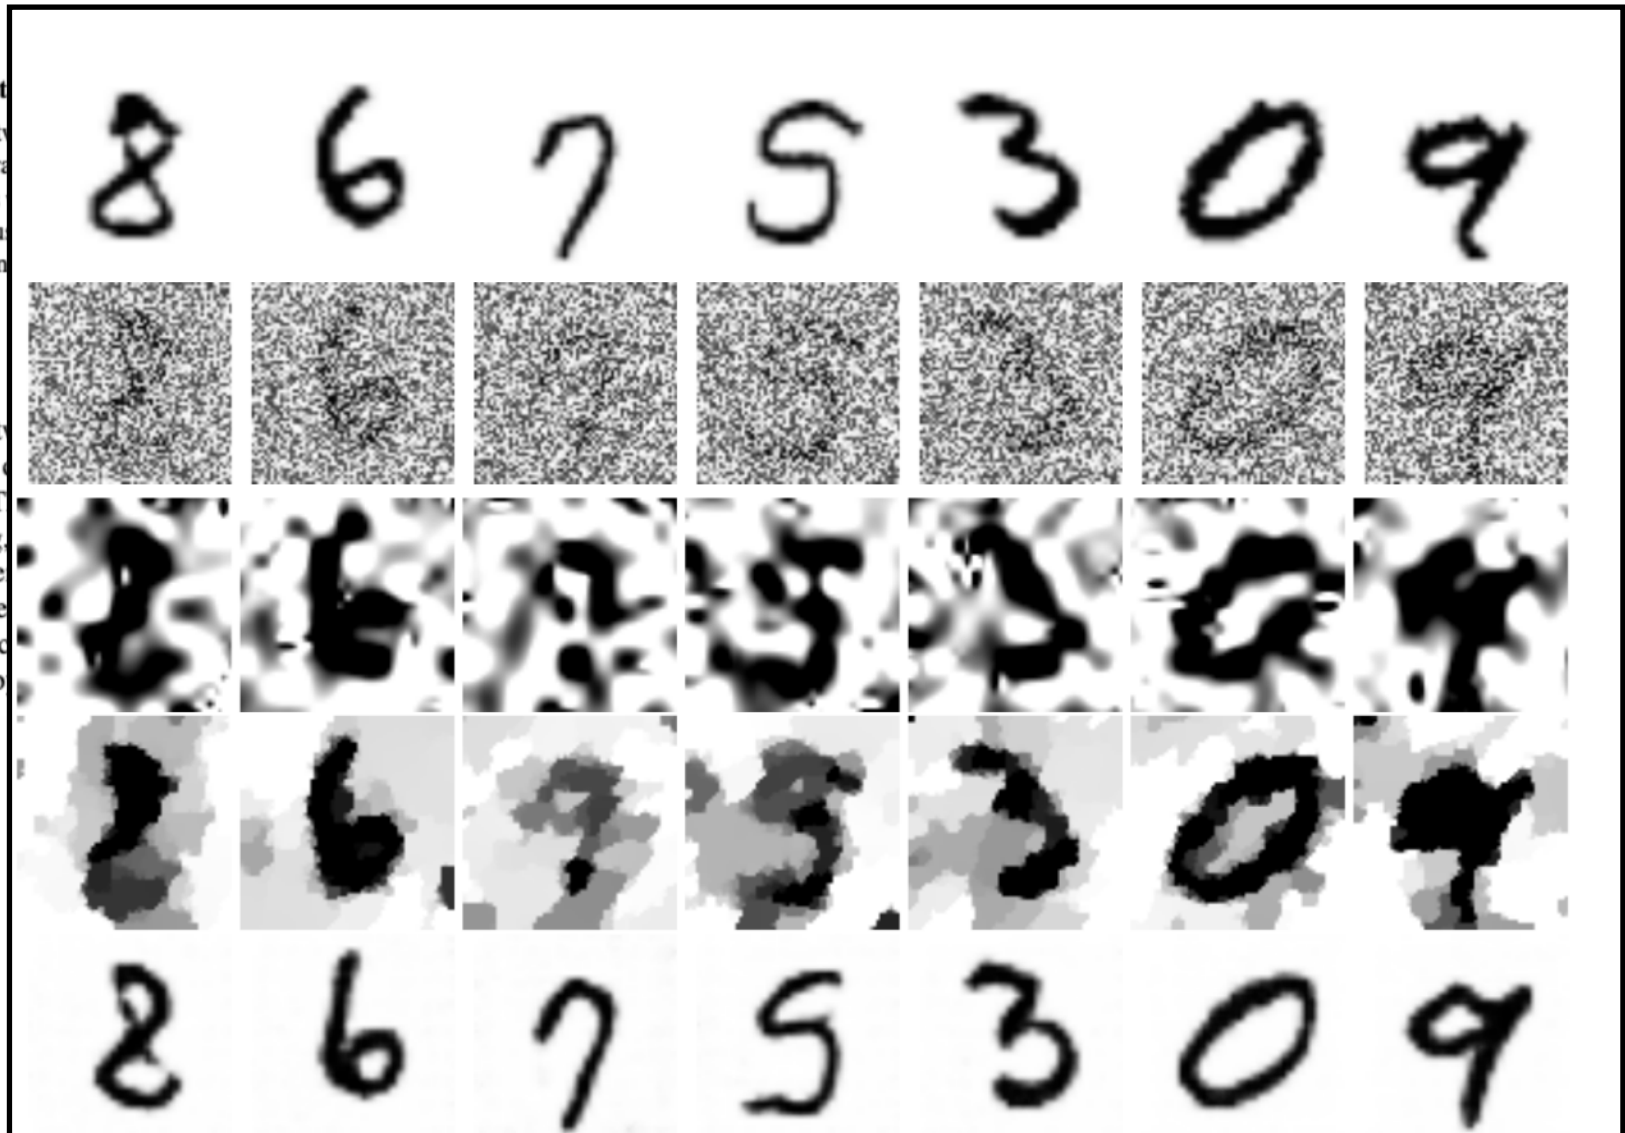


Figure 1: Denoising with generative priors

(**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 Networks: Why?

Motivation (amongst others): priors in inverse problems

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.

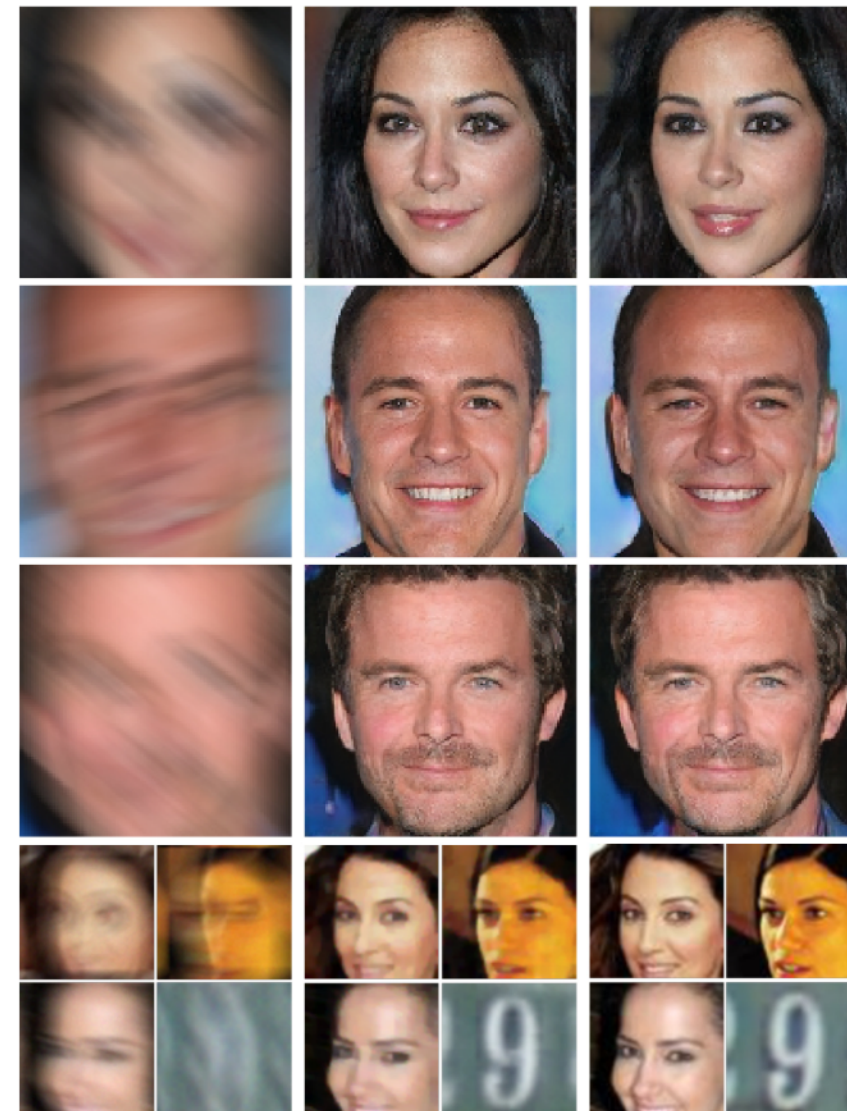
Index Terms—Blind image deblurring, generative adversarial networks, variational autoencoders, deep image prior.

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 . For a thorough discussion on

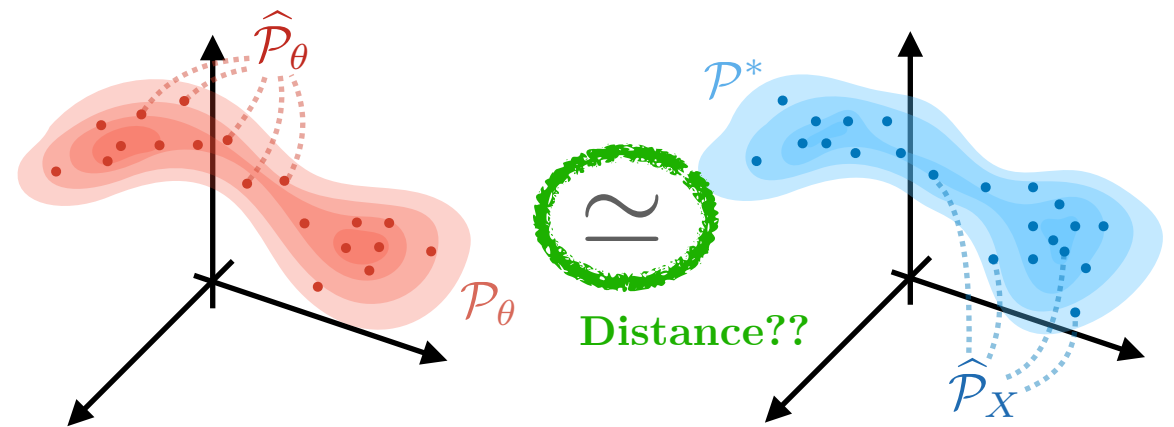


(a) Blurry (b) Ours (c) Original

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

Generative Networks: How?

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

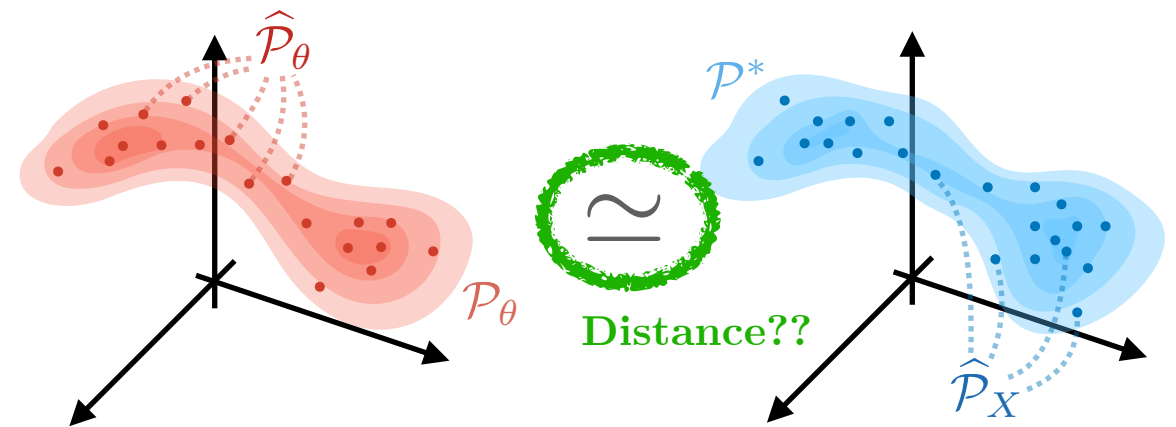
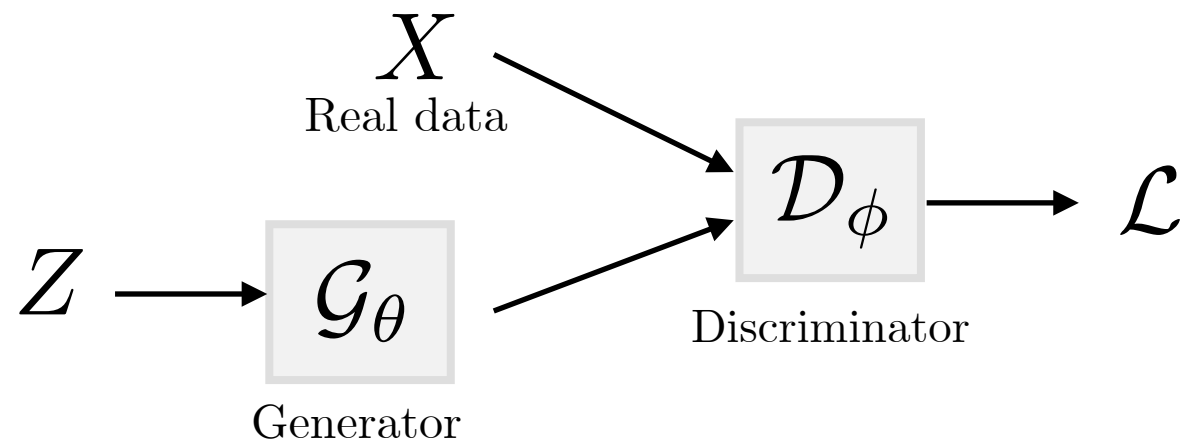


Generative Networks: How?

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

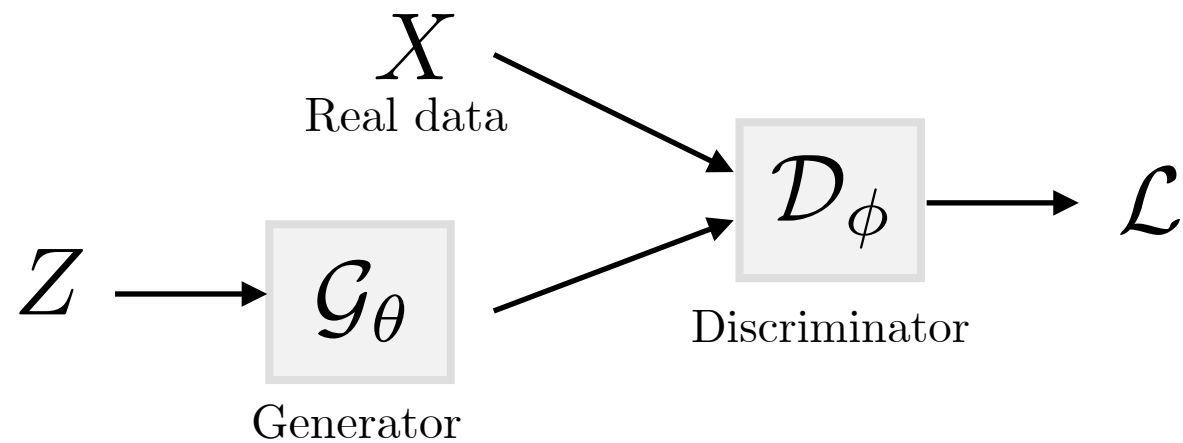


Generative Networks: How?

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

1) Golden standard: Generative Adversarial Networks

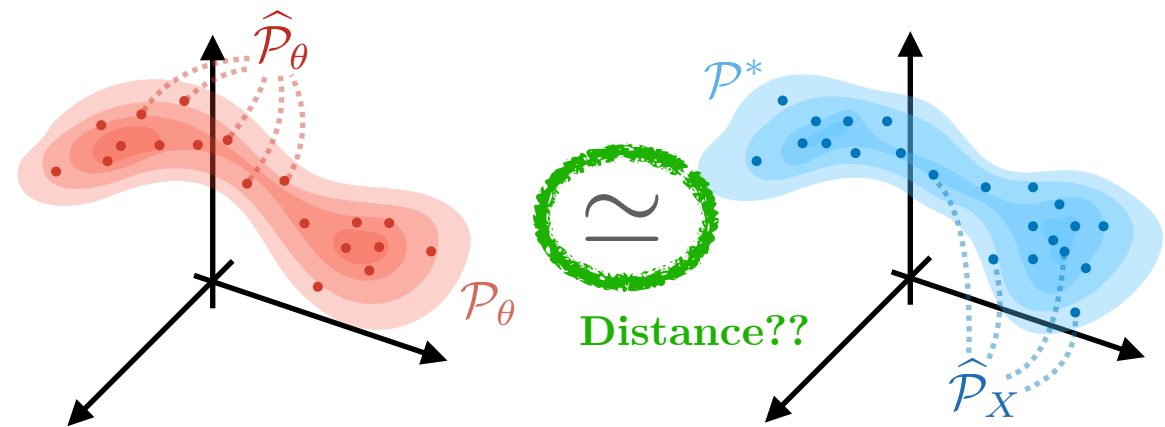
Learn a second “discriminator” network that classifies real/fake at the same time as the generator



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

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

Lots of research interest in other (easier) methods

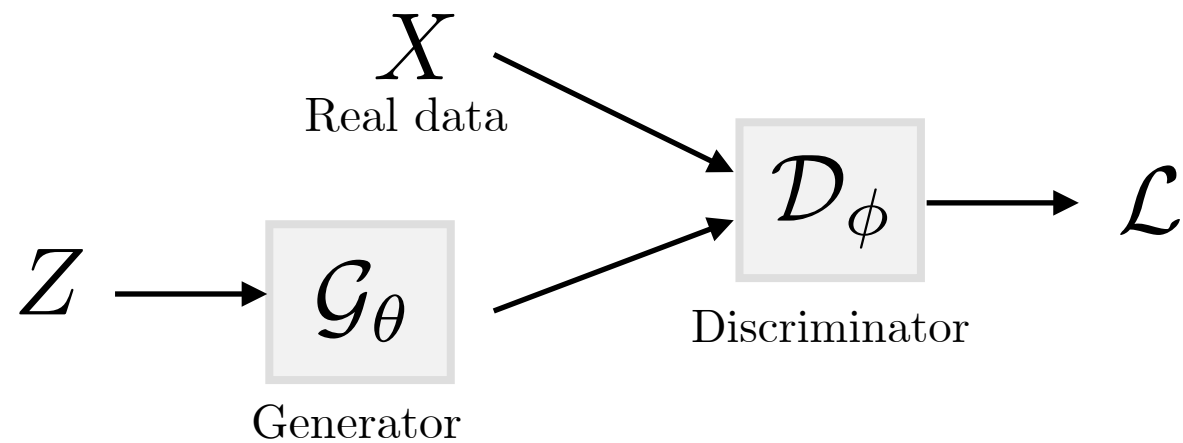


Generative Networks: How?

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



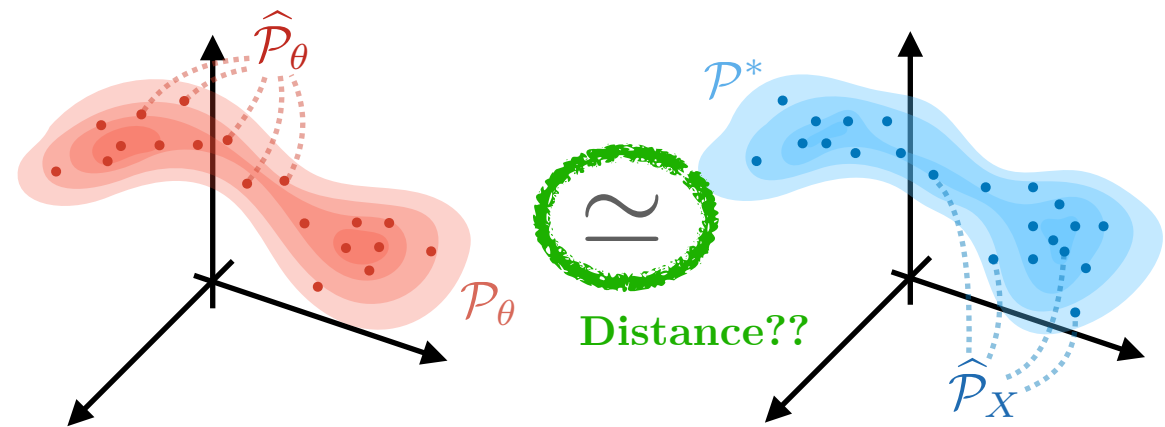
Very difficult to train (due to balancing of training discriminator/generator)

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

2) Maximum Mean Discrepancy

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

Easier to train (no balancing)...

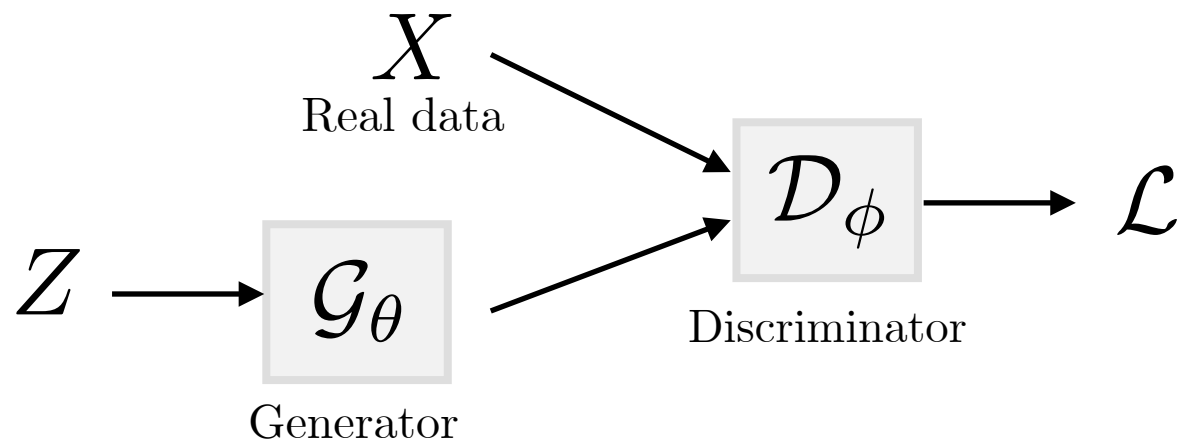
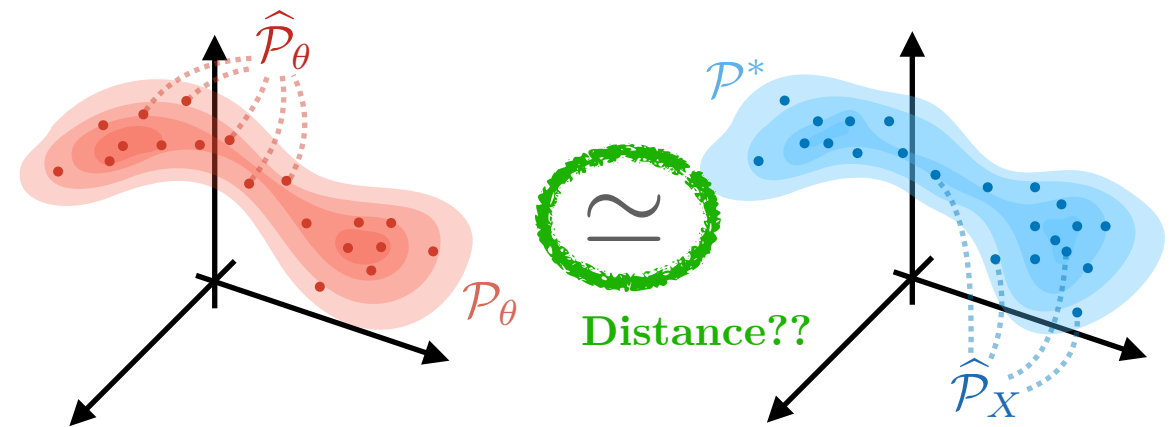


Generative Networks: How?

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

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

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

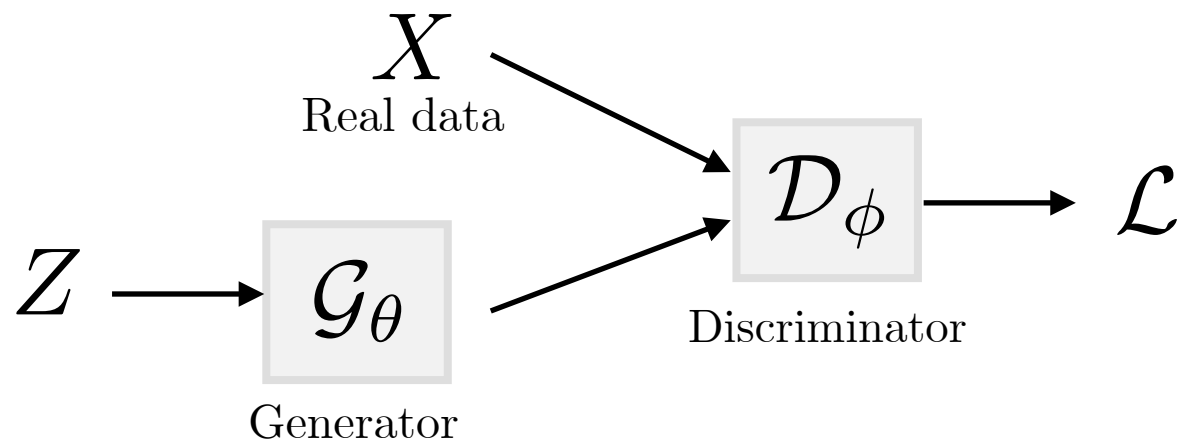
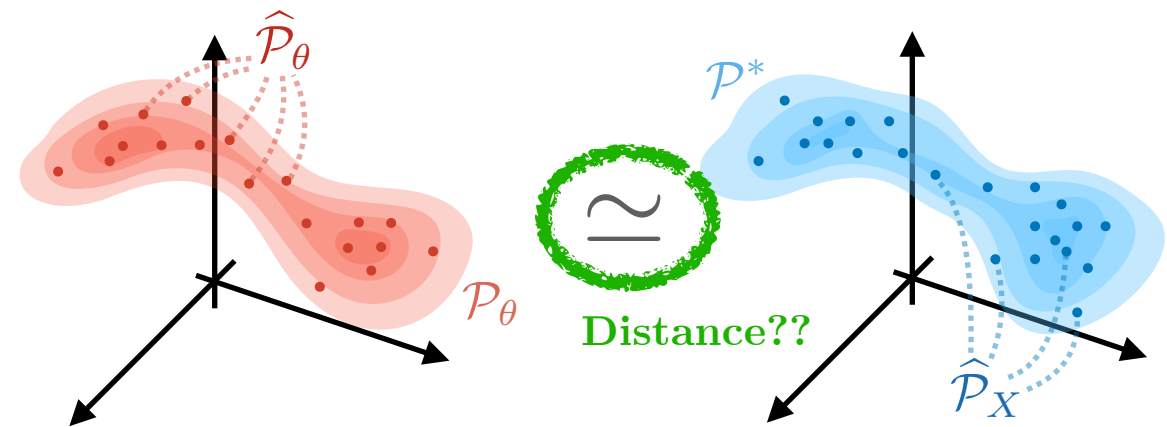
Easier to train (no balancing)...

Generative Networks: How?

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

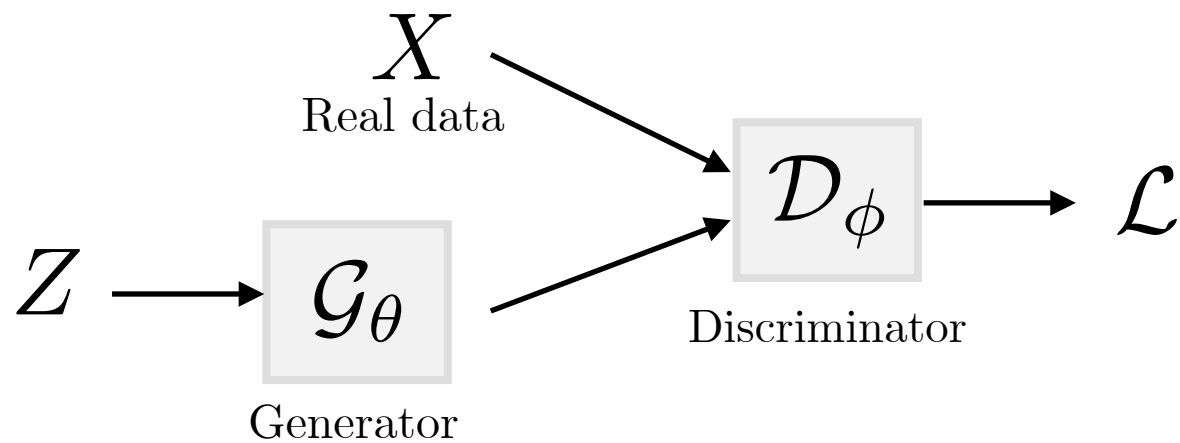
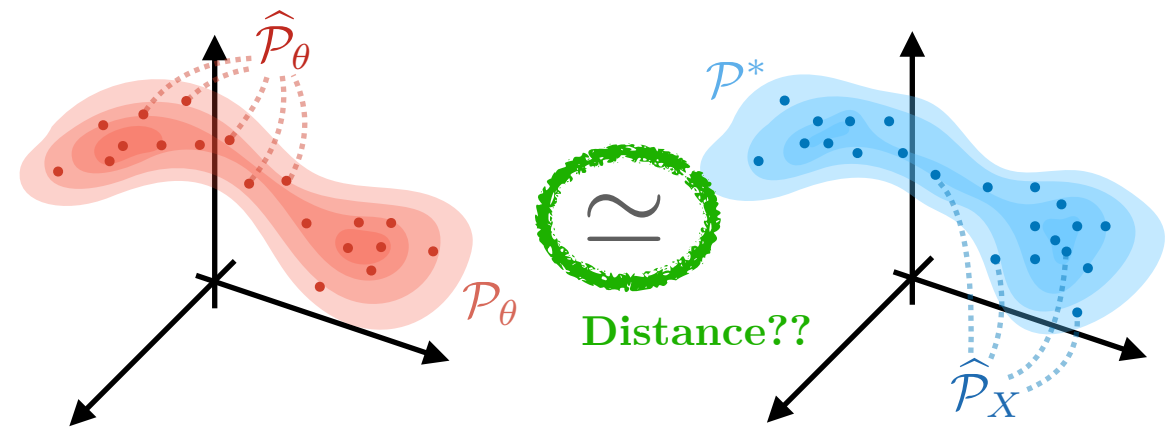
Easier to train (no balancing)...

Generative Networks: How?

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)...

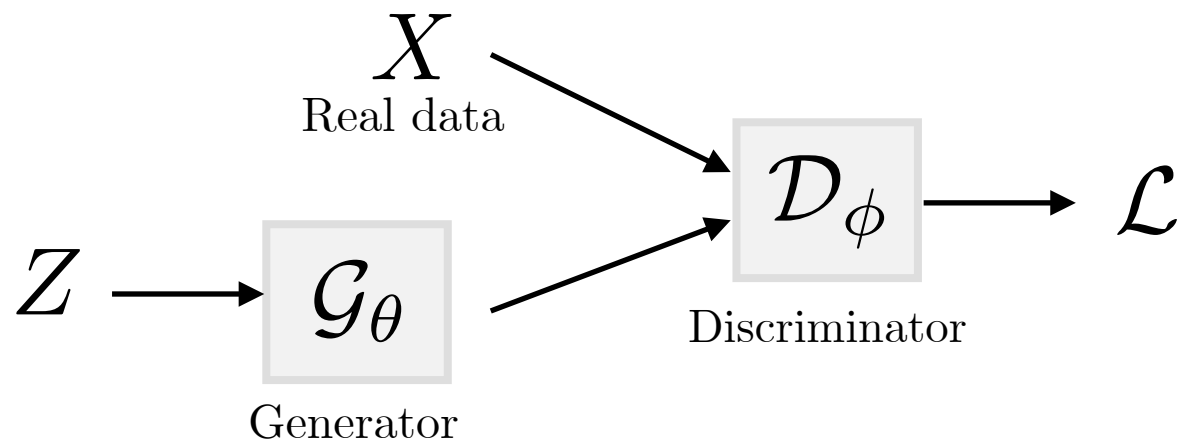
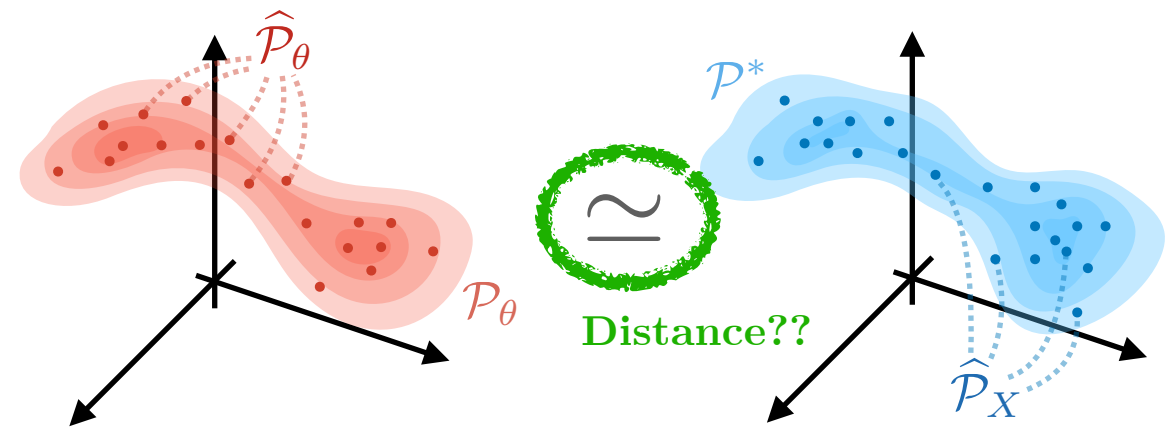
With $\Lambda = \mathcal{F}\kappa$

Generative Networks: How?

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

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

Similarity between samples

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

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) but quadratic complexity...

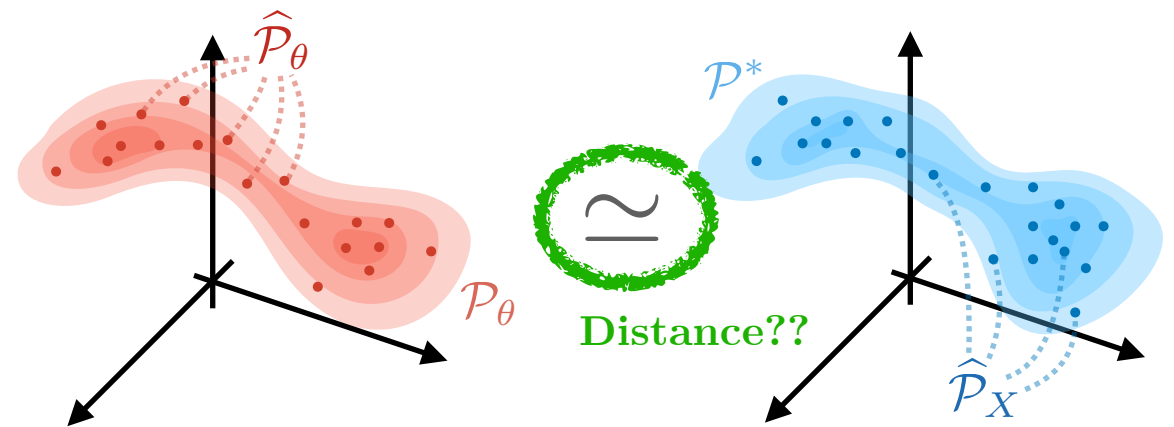
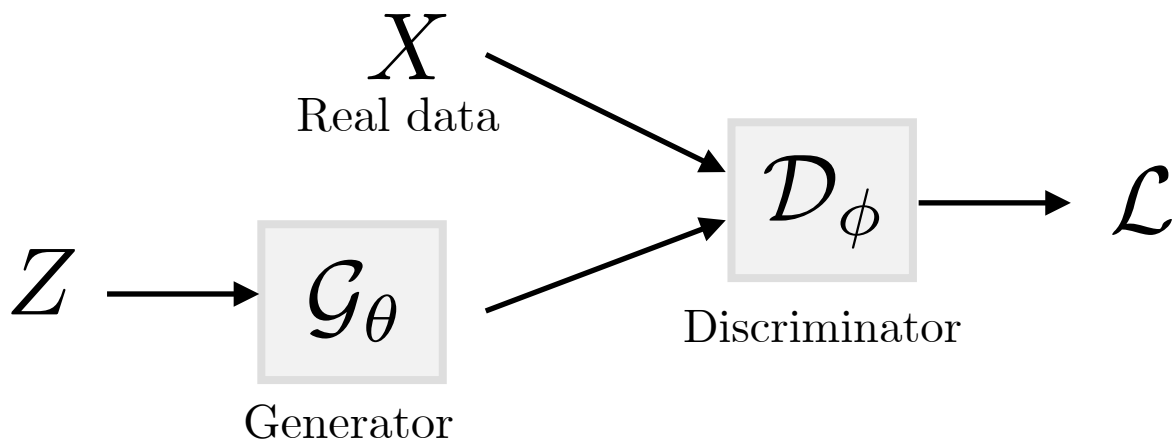
With $\Lambda = \mathcal{F}\kappa$

Generative Networks: How?

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

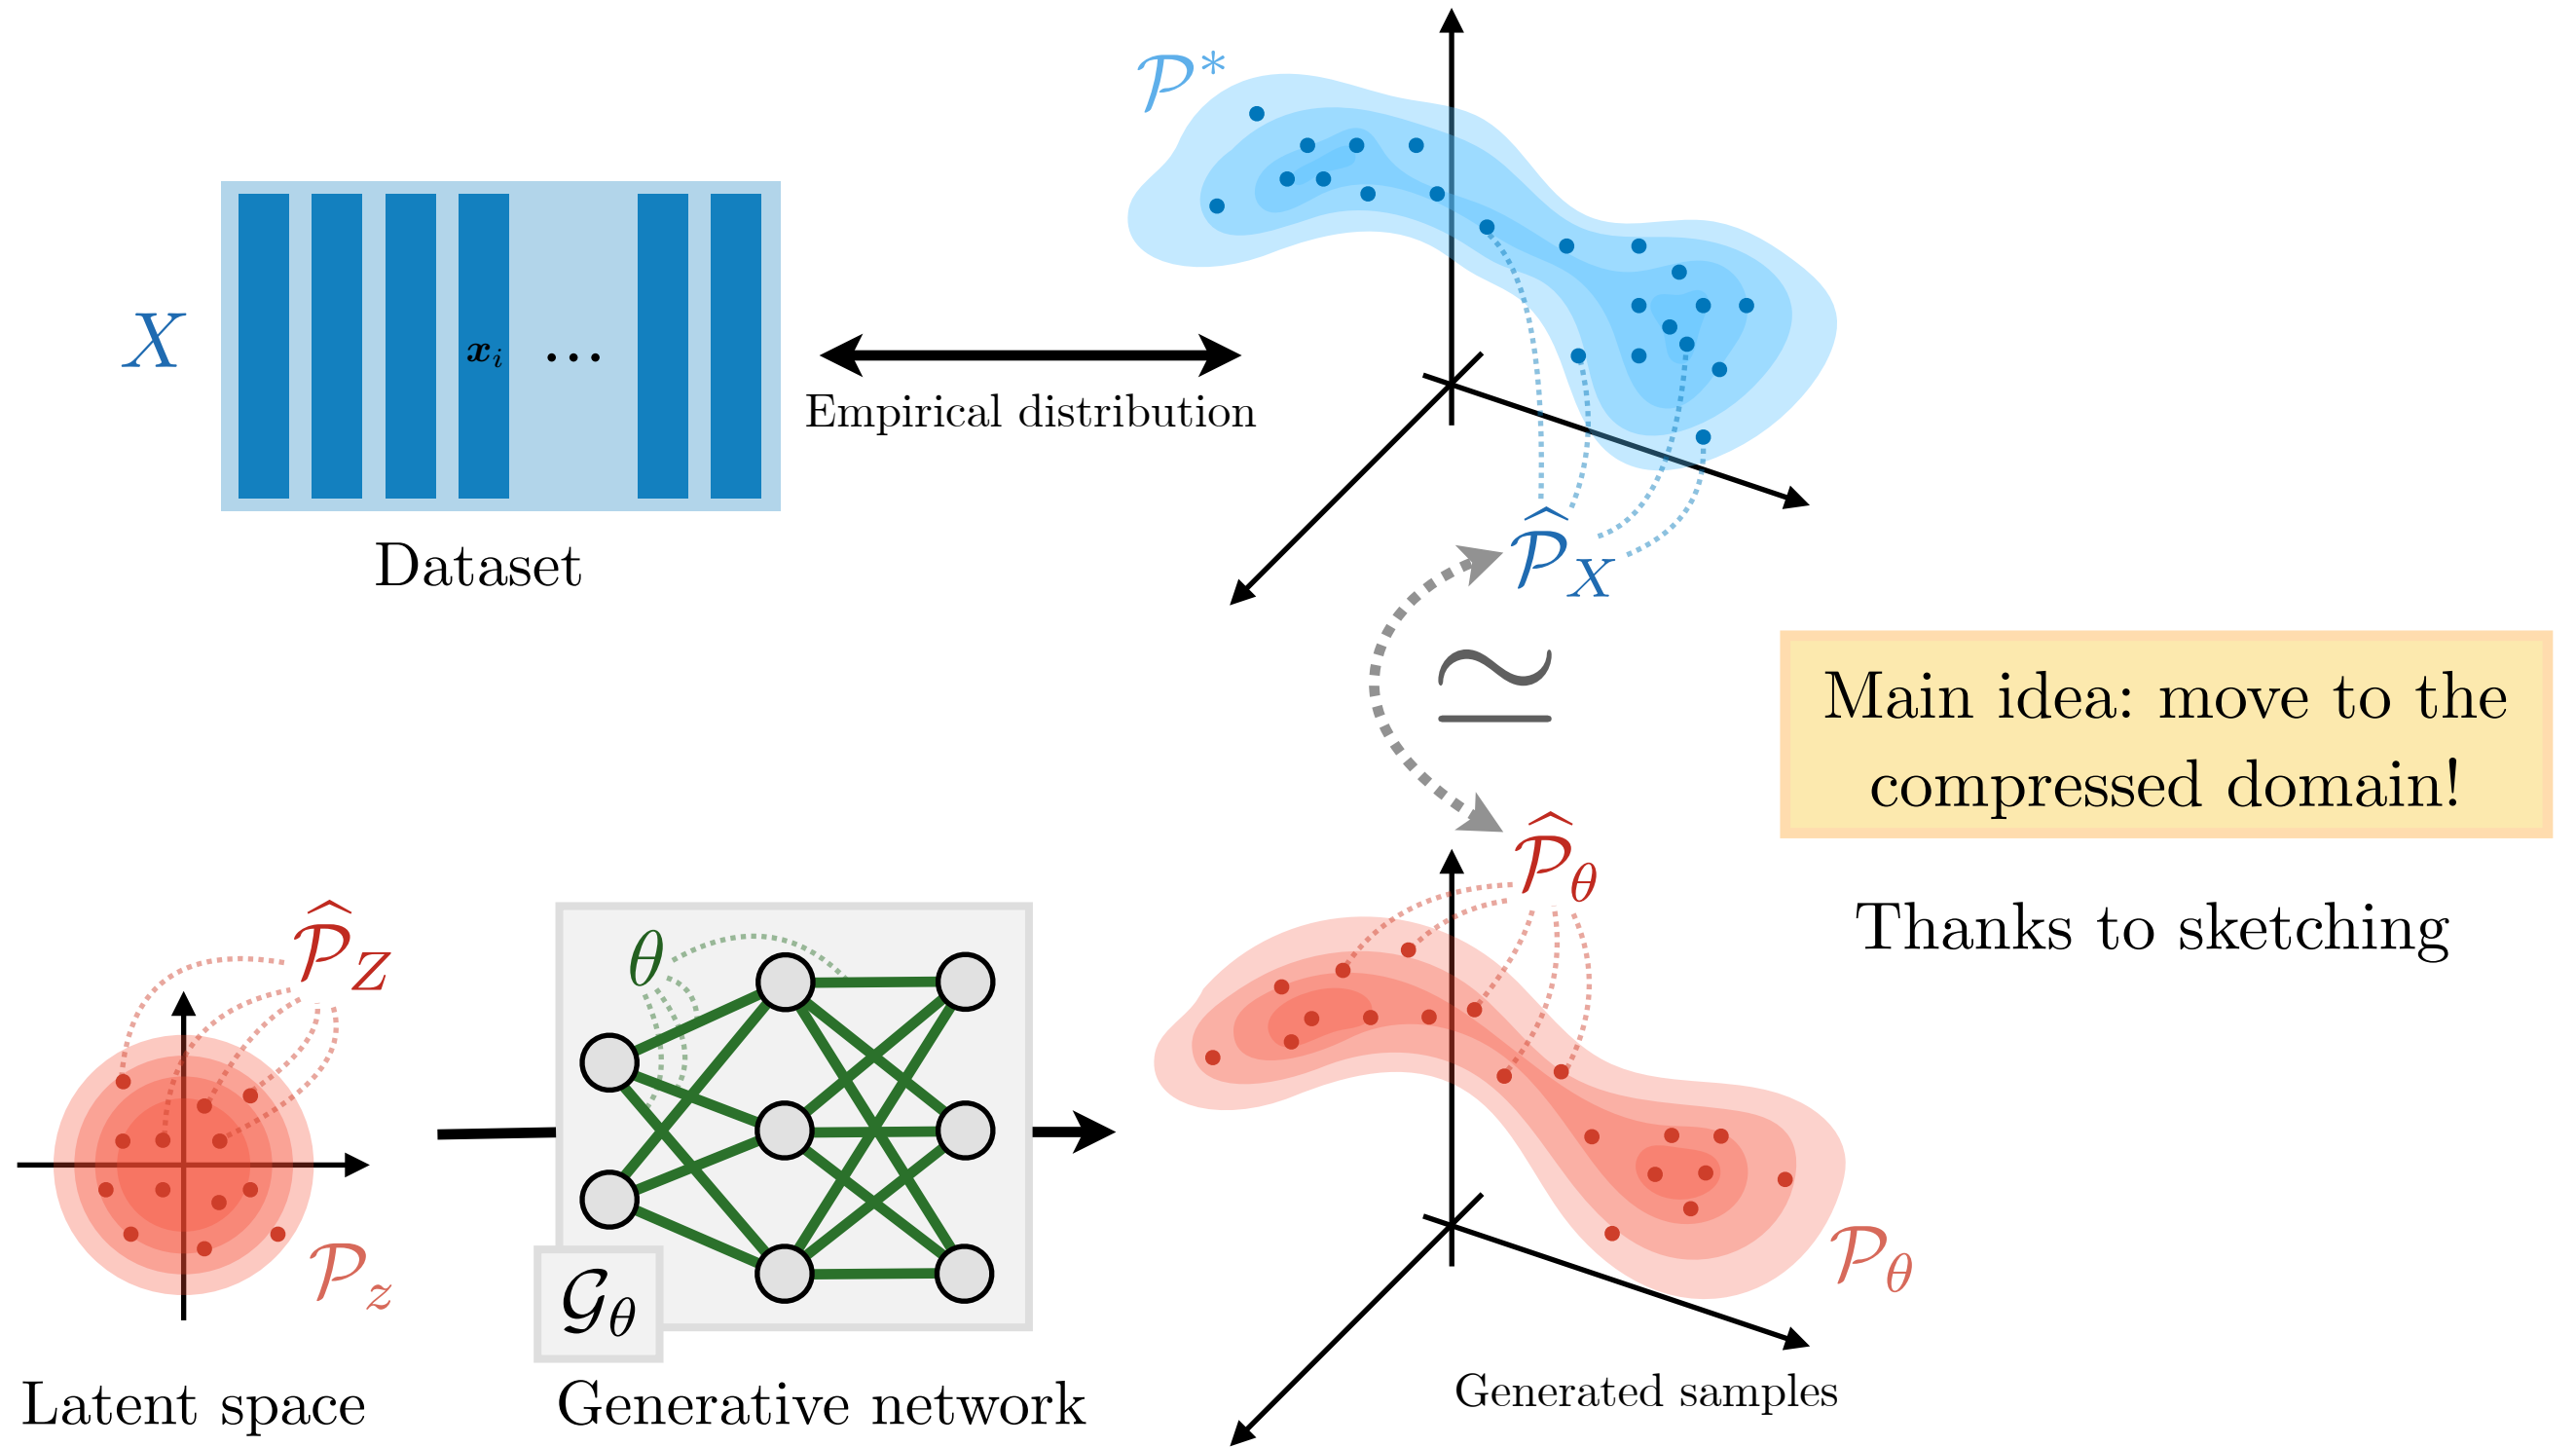
Training generative networks typically requires massive amounts of data!

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

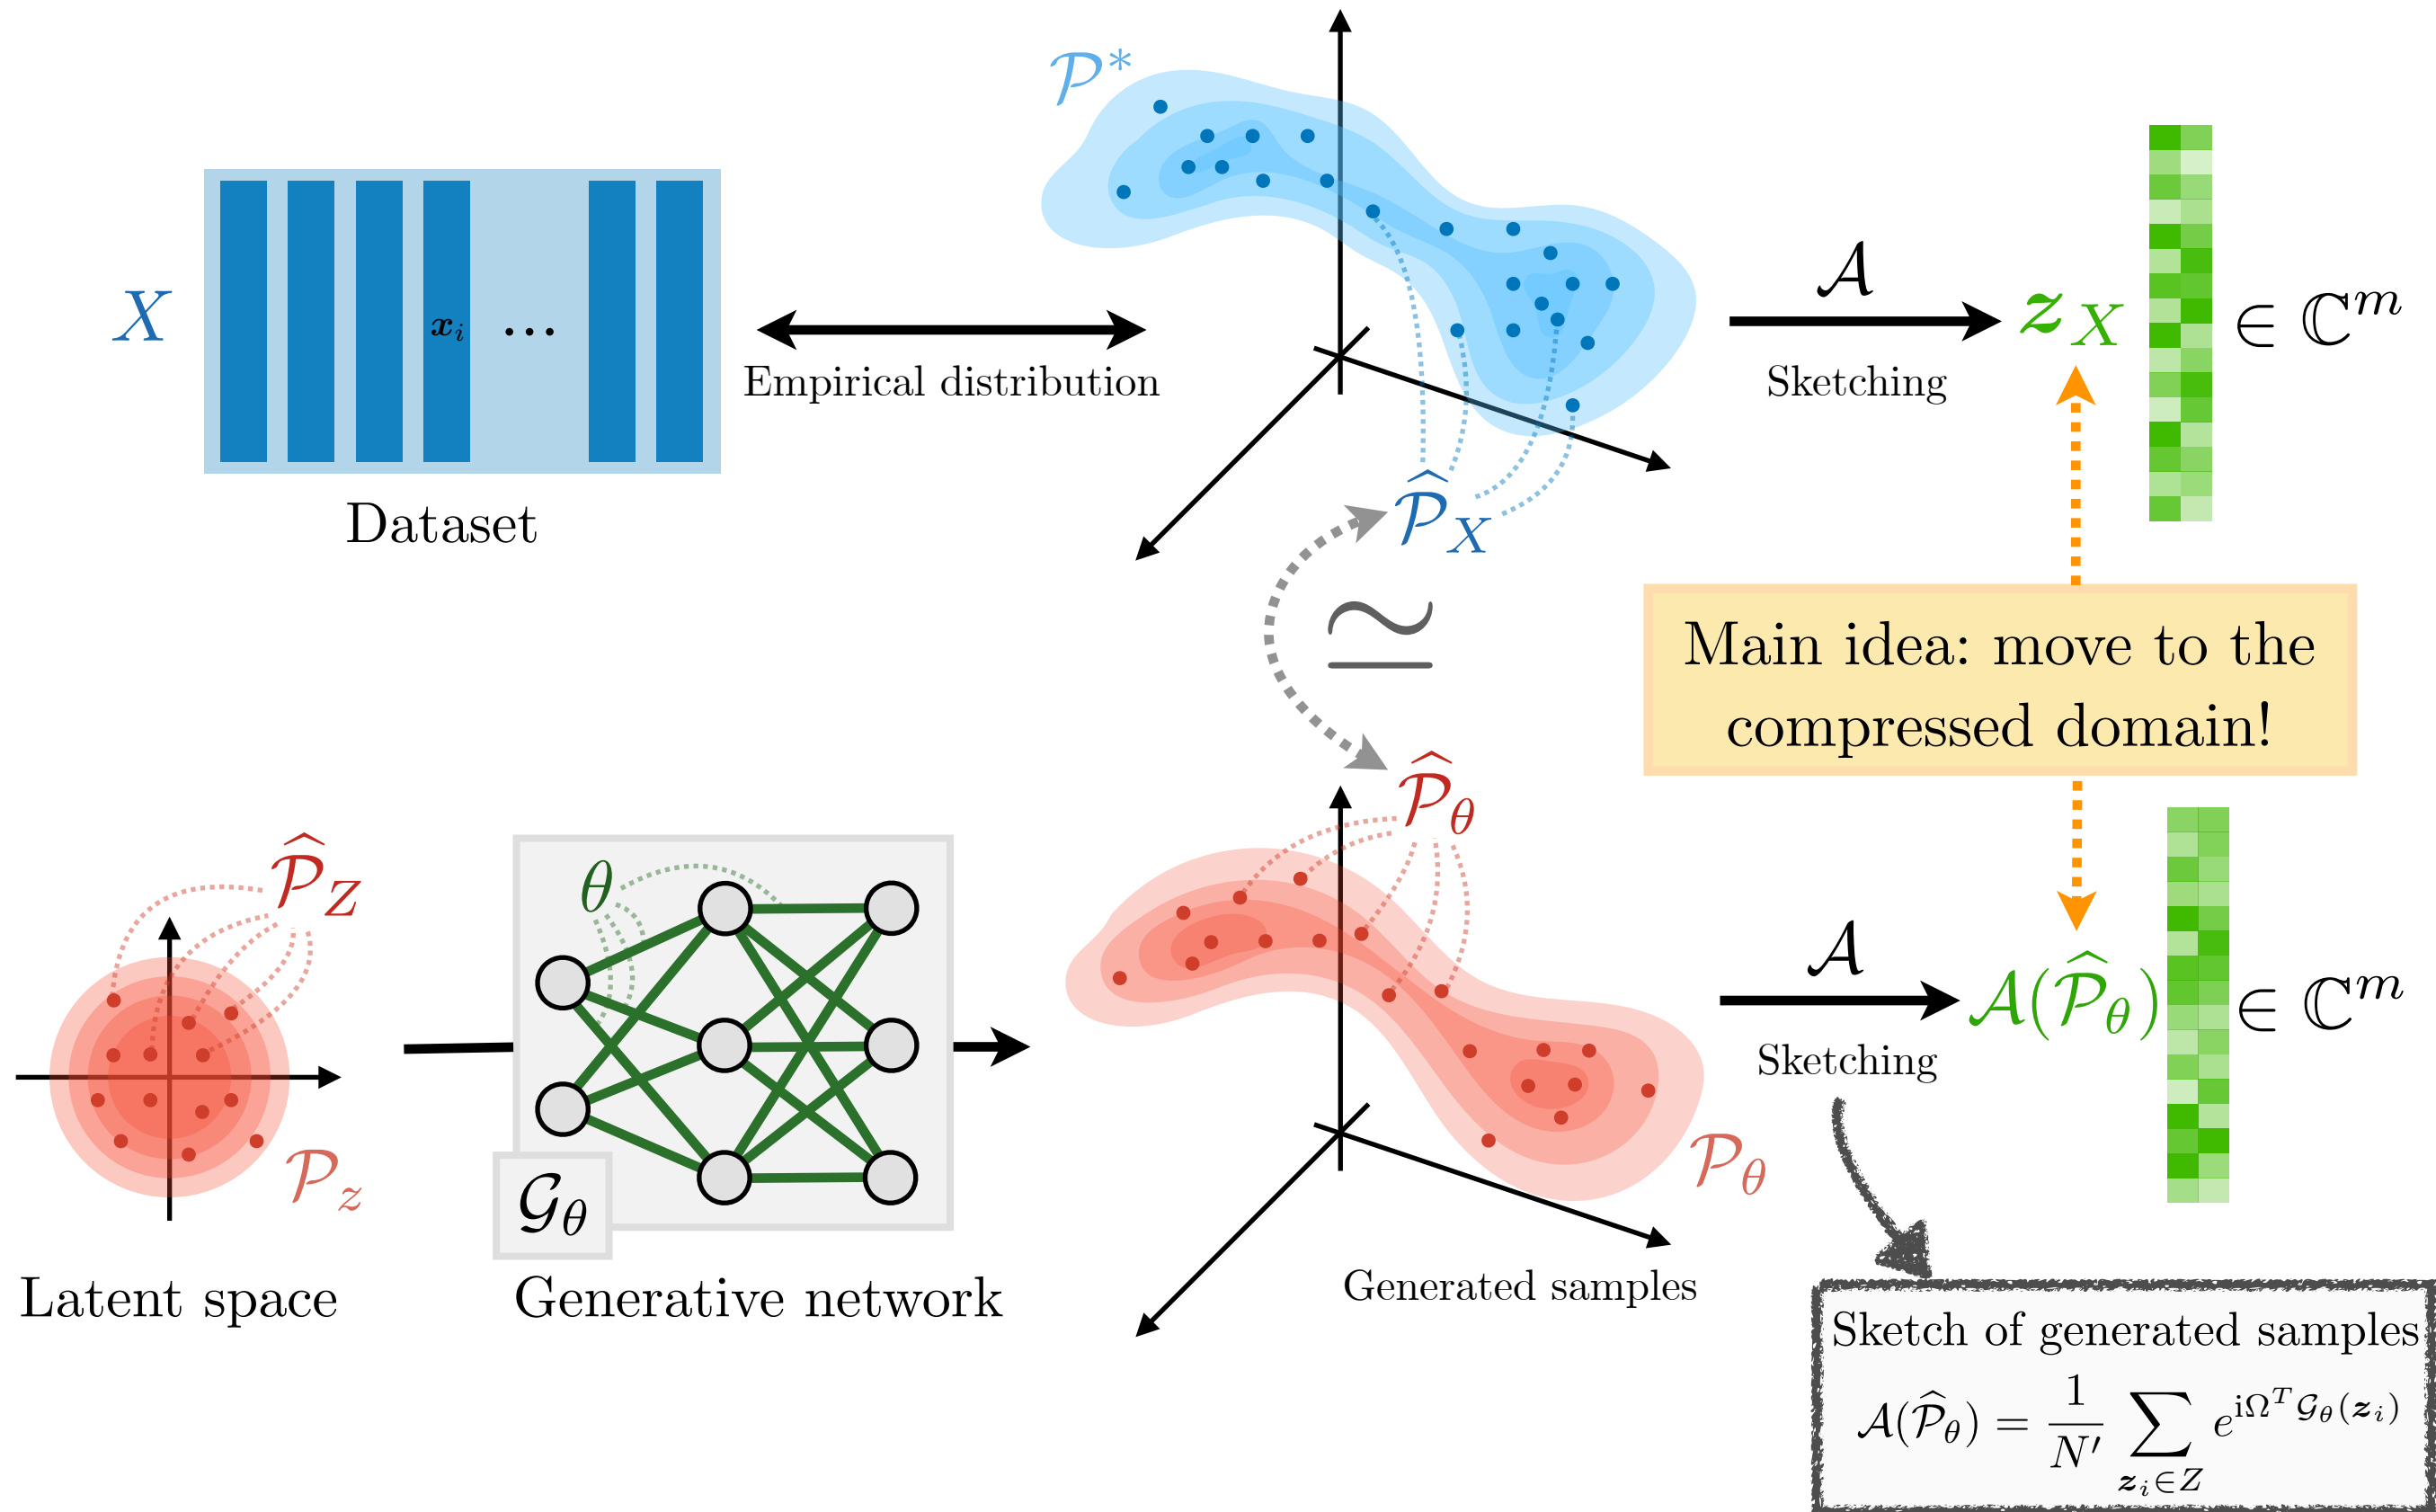
$$\mathcal{O}(N'(N + N'))$$

Compressive Learning of Generative Networks

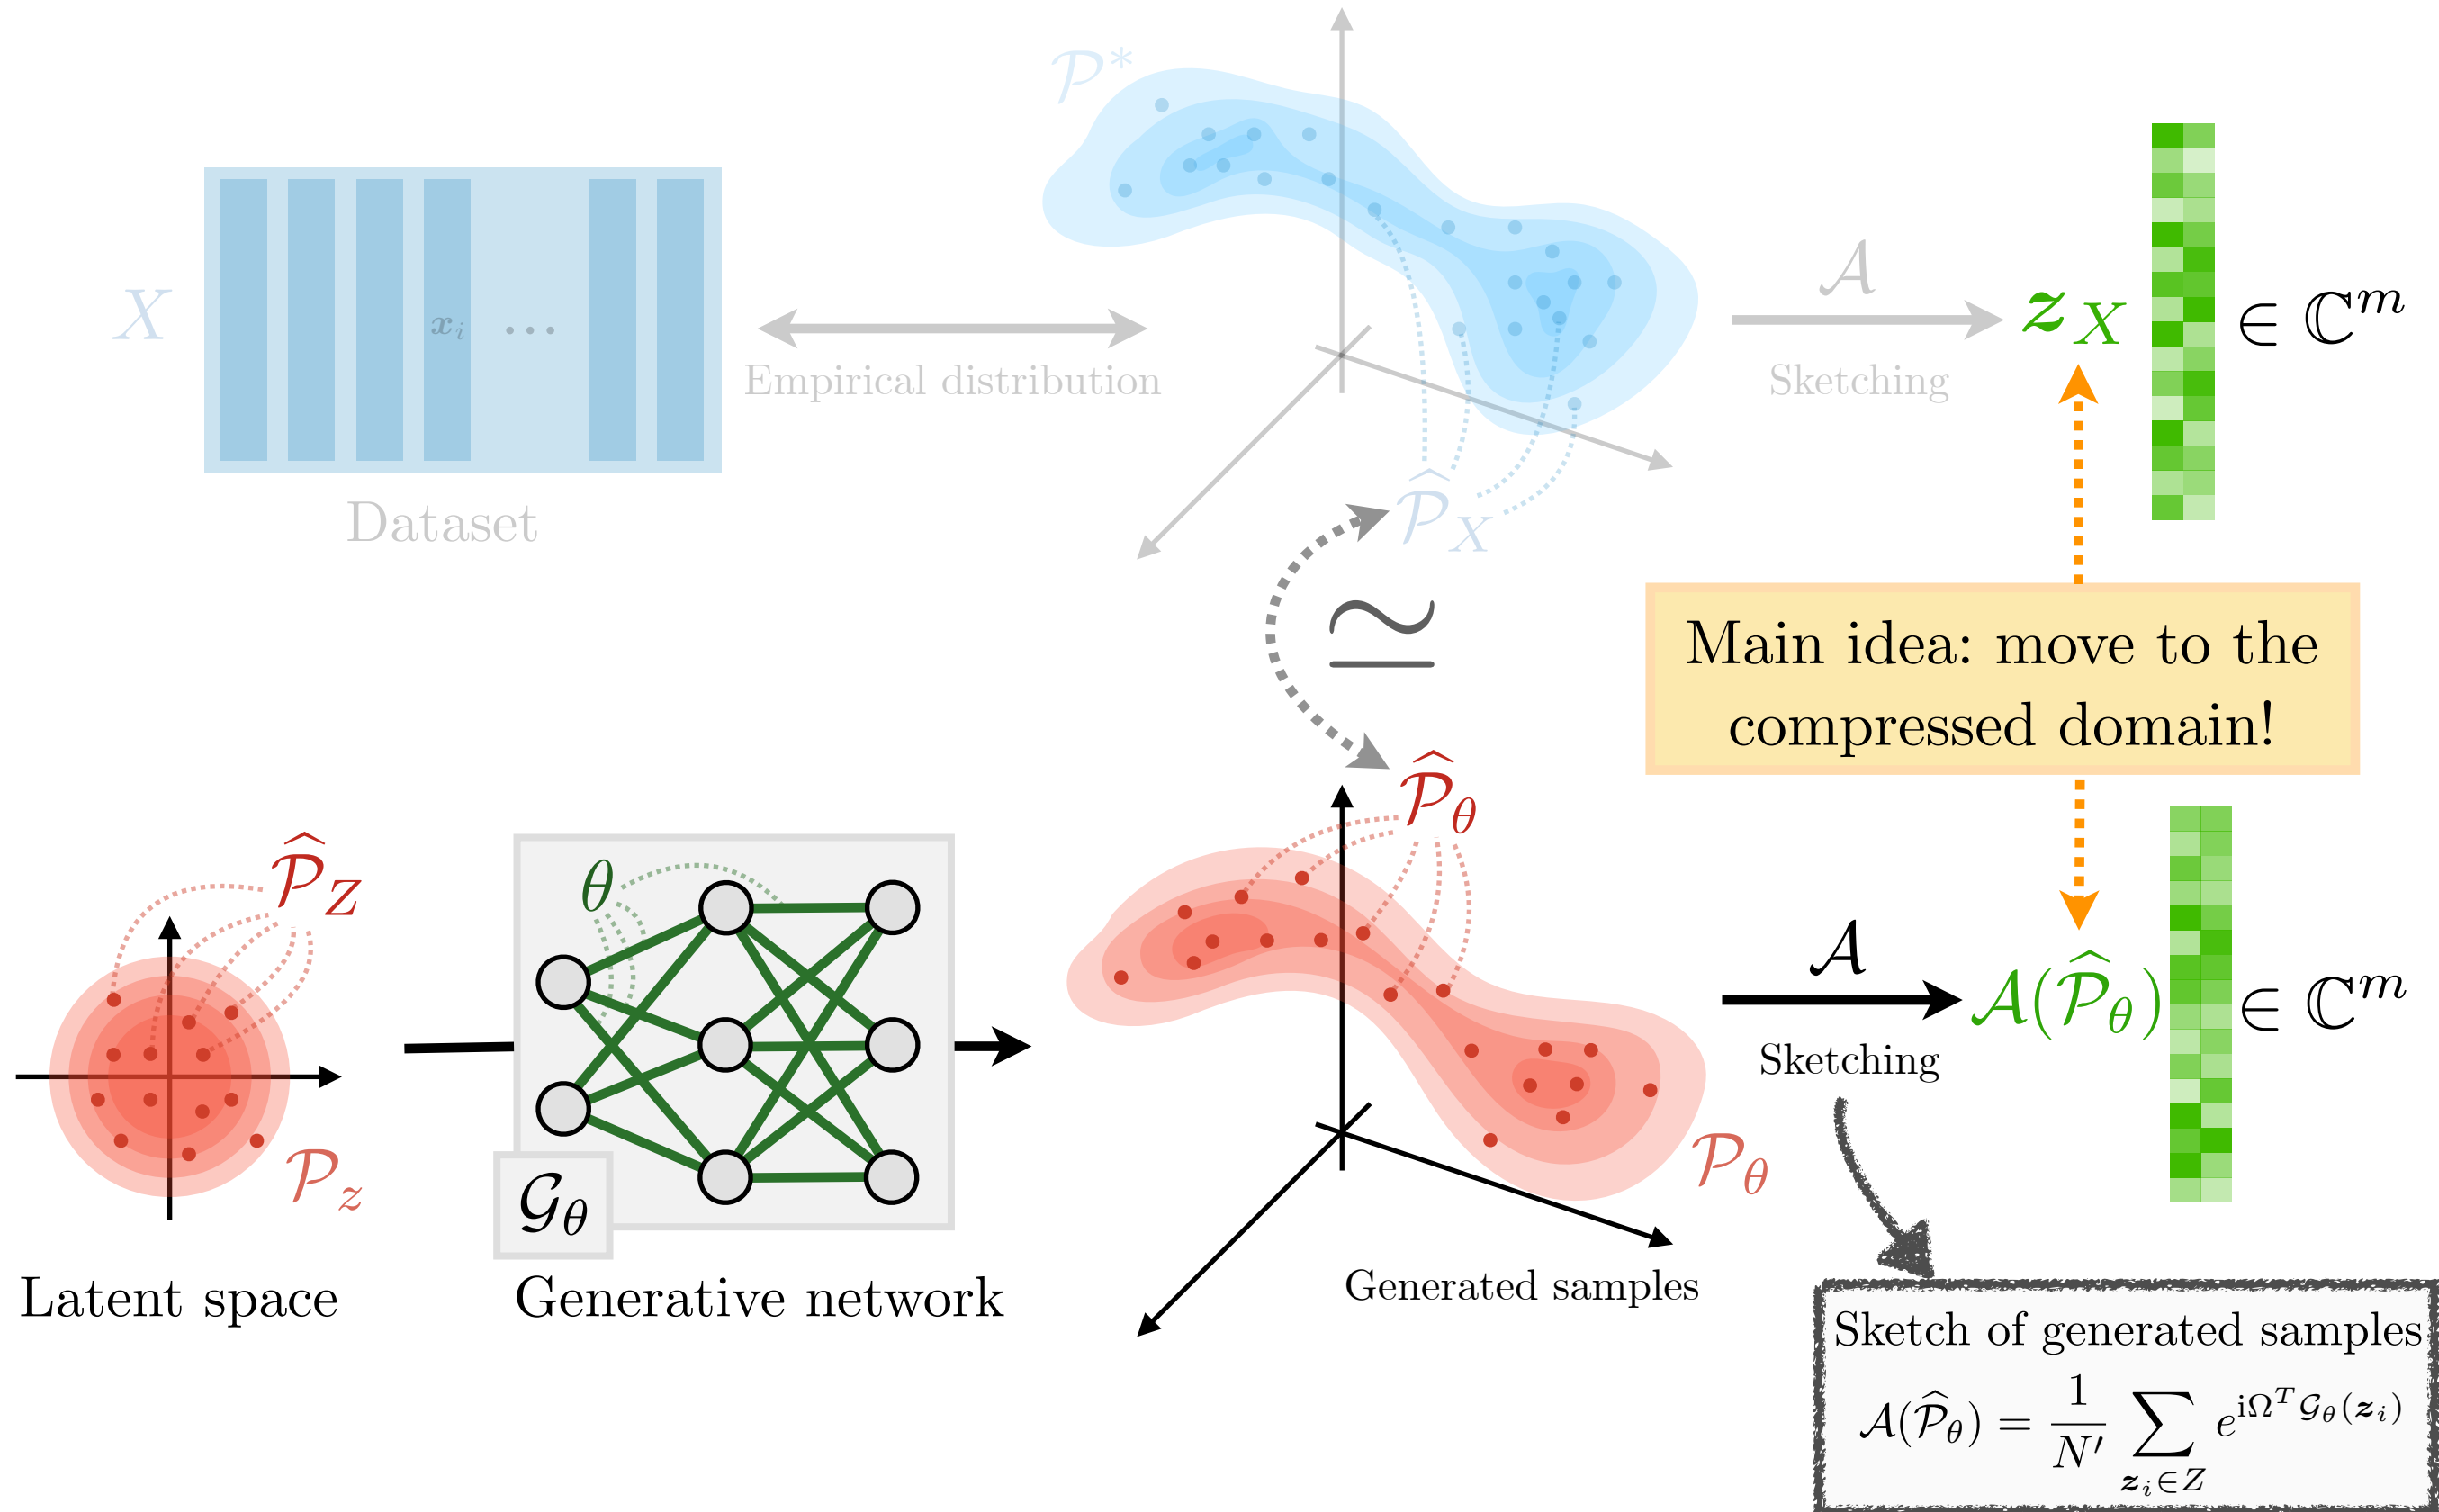
Compressively learning generative networks



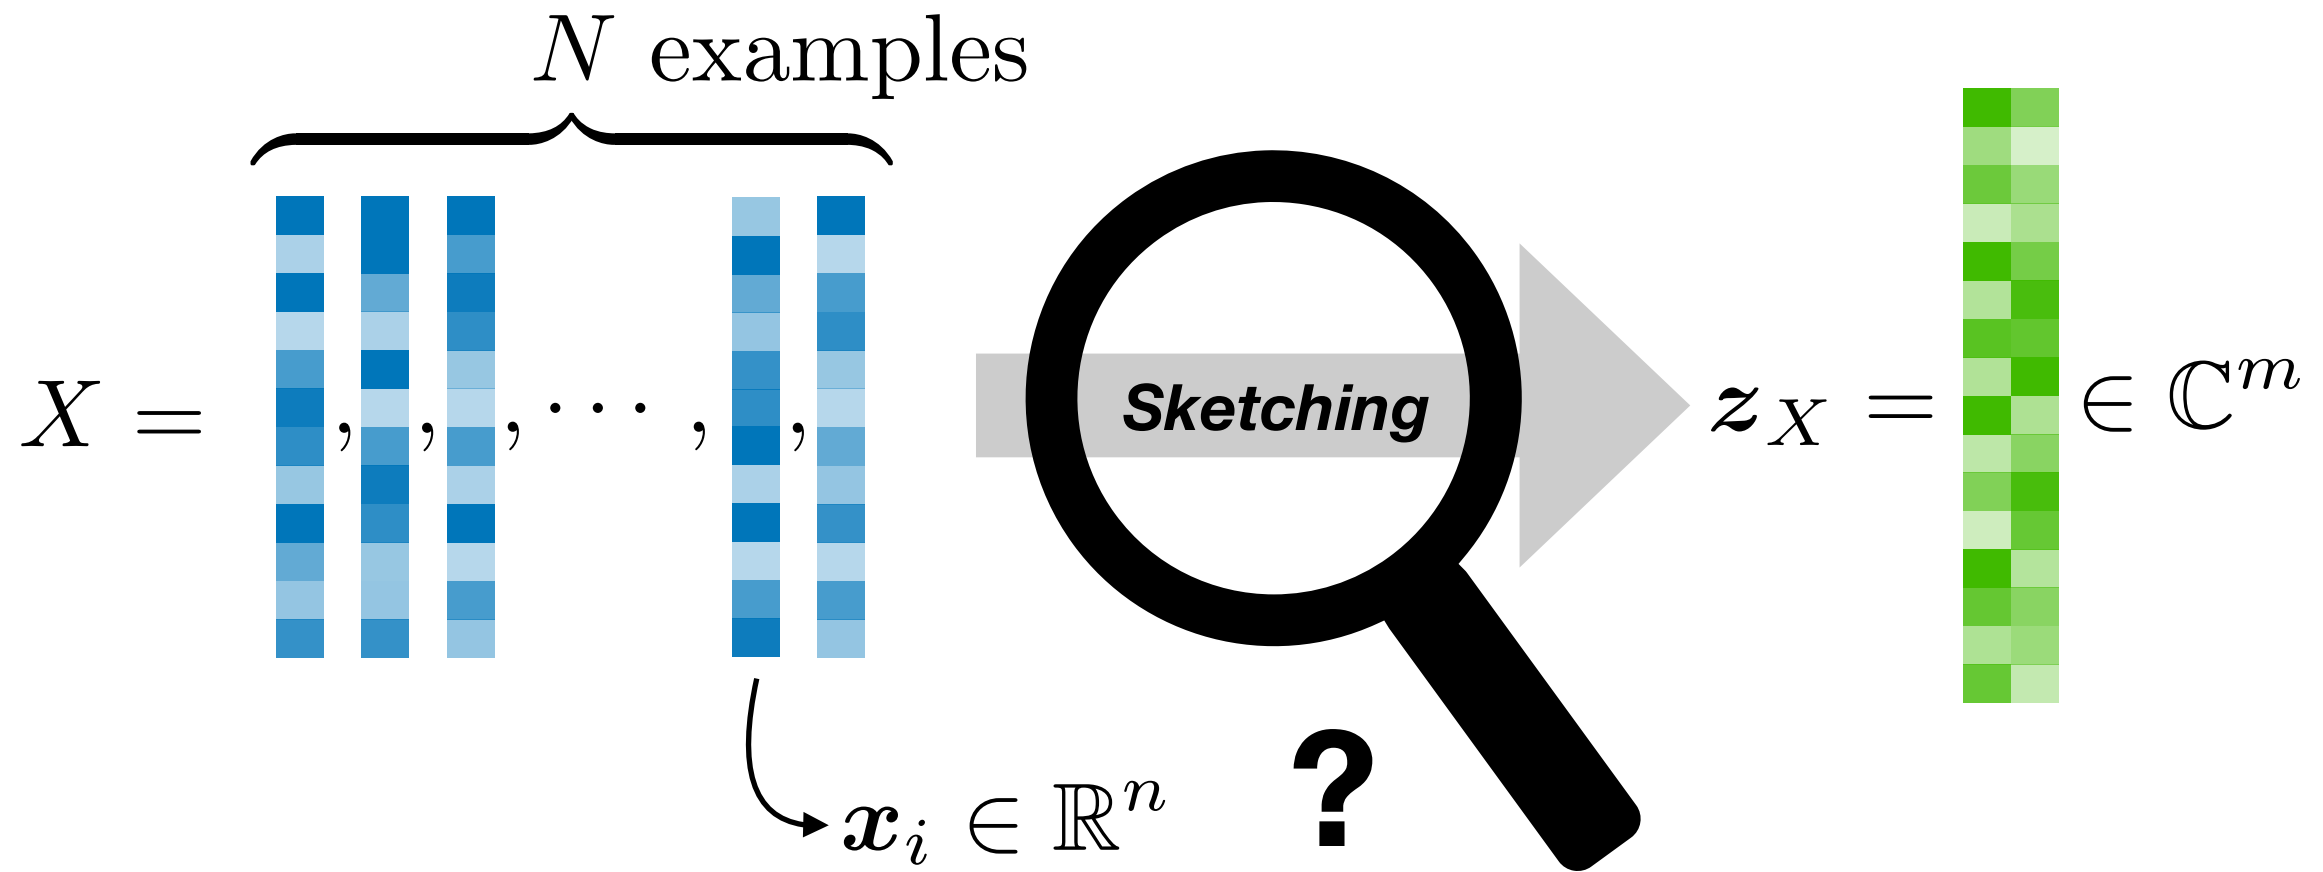
Compressively learning generative networks



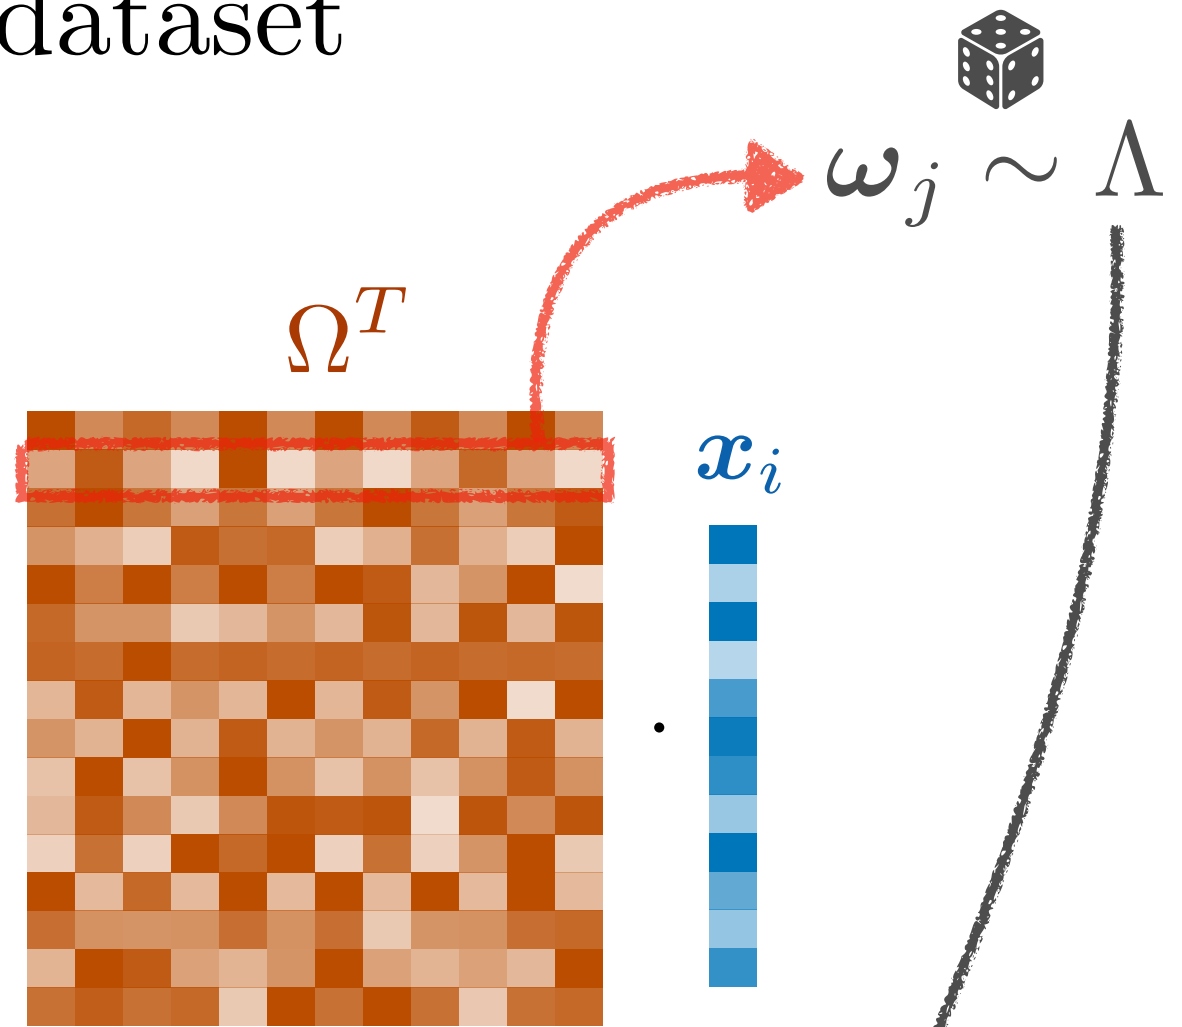
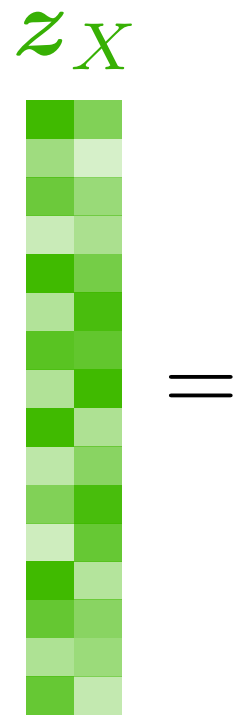
Compressively learning generative networks



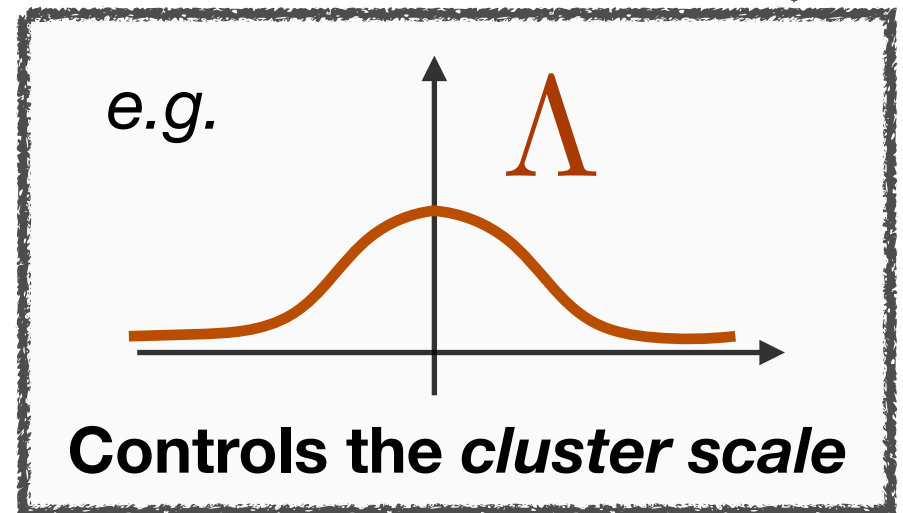
Sketching a dataset



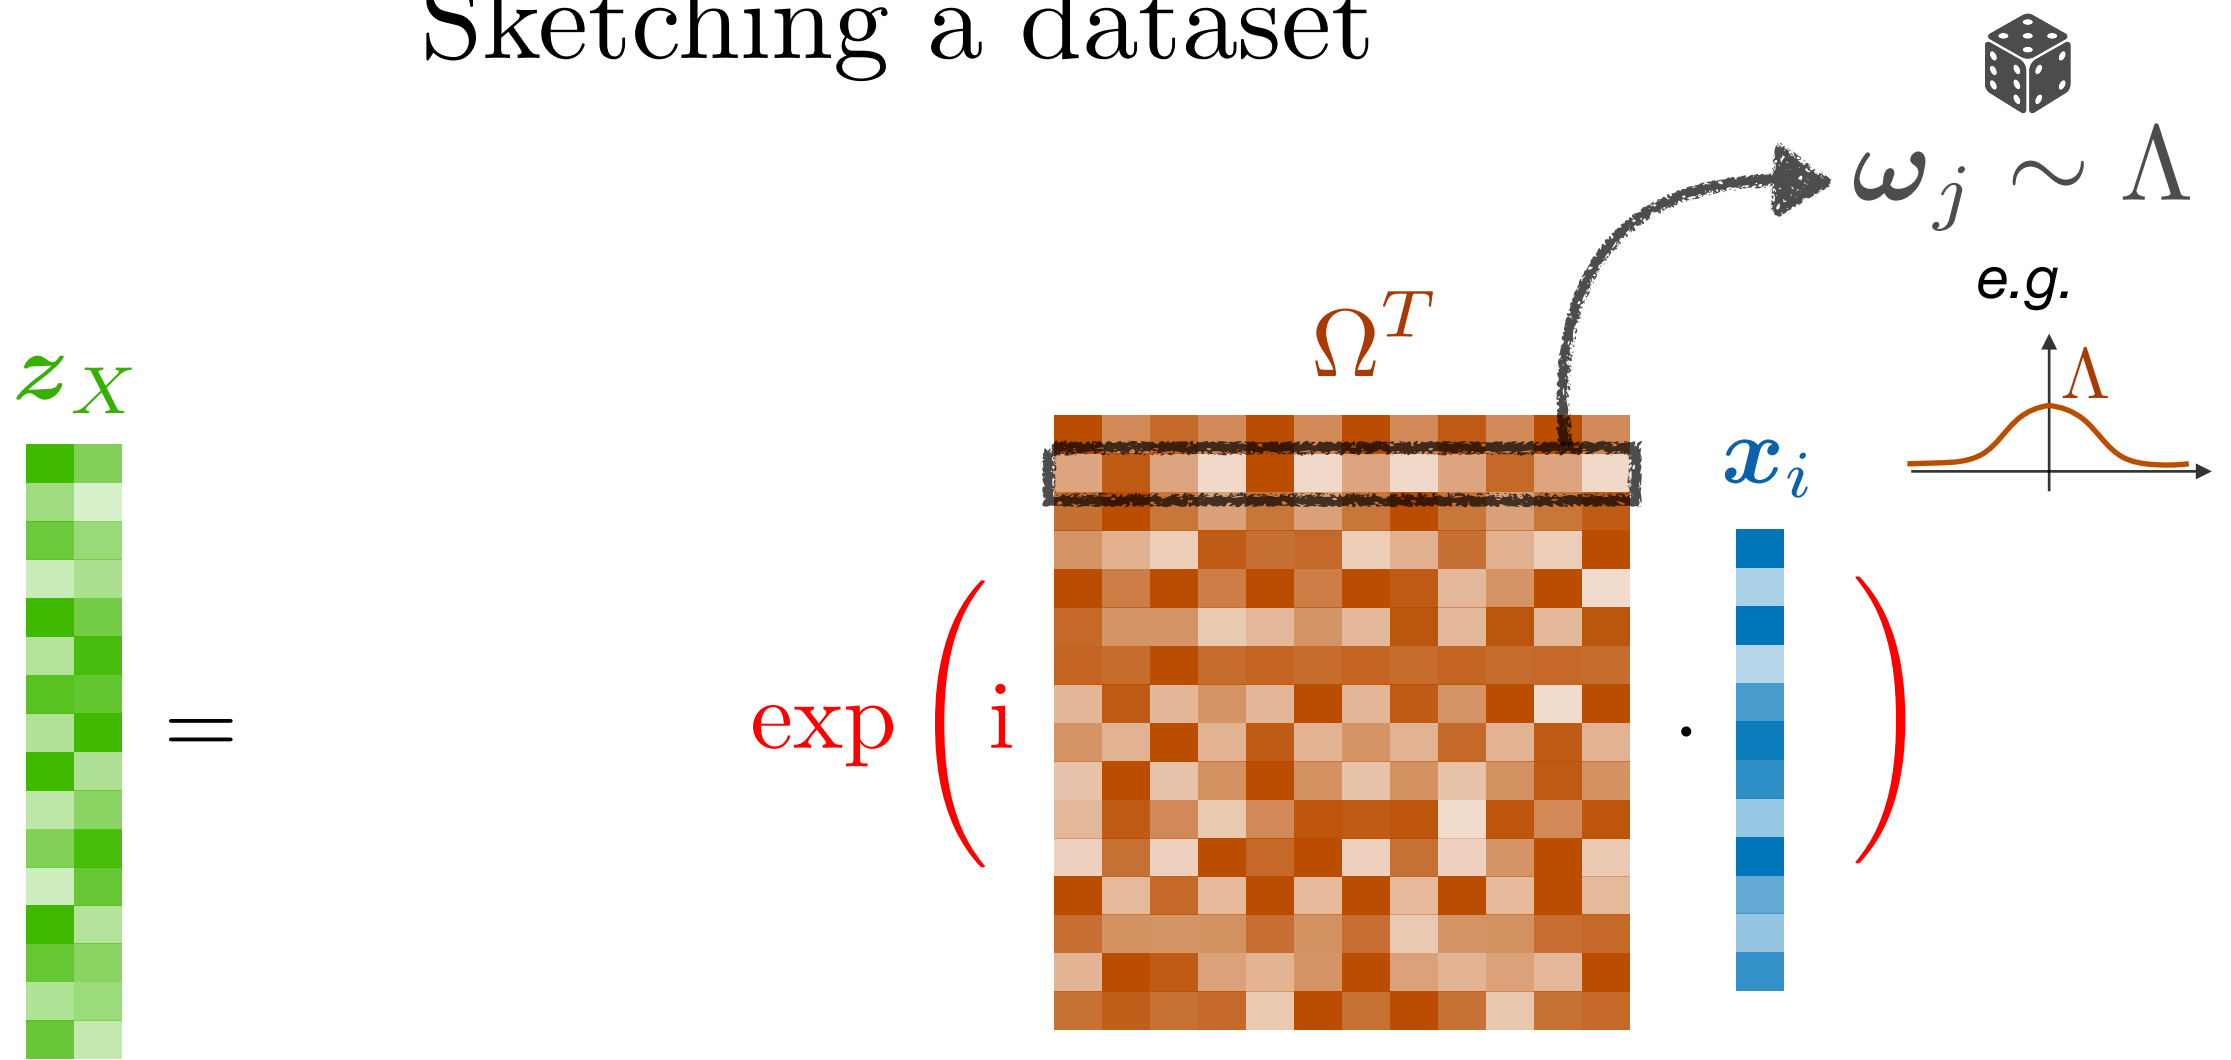
Sketching a dataset



1. Project on m (random) vectors

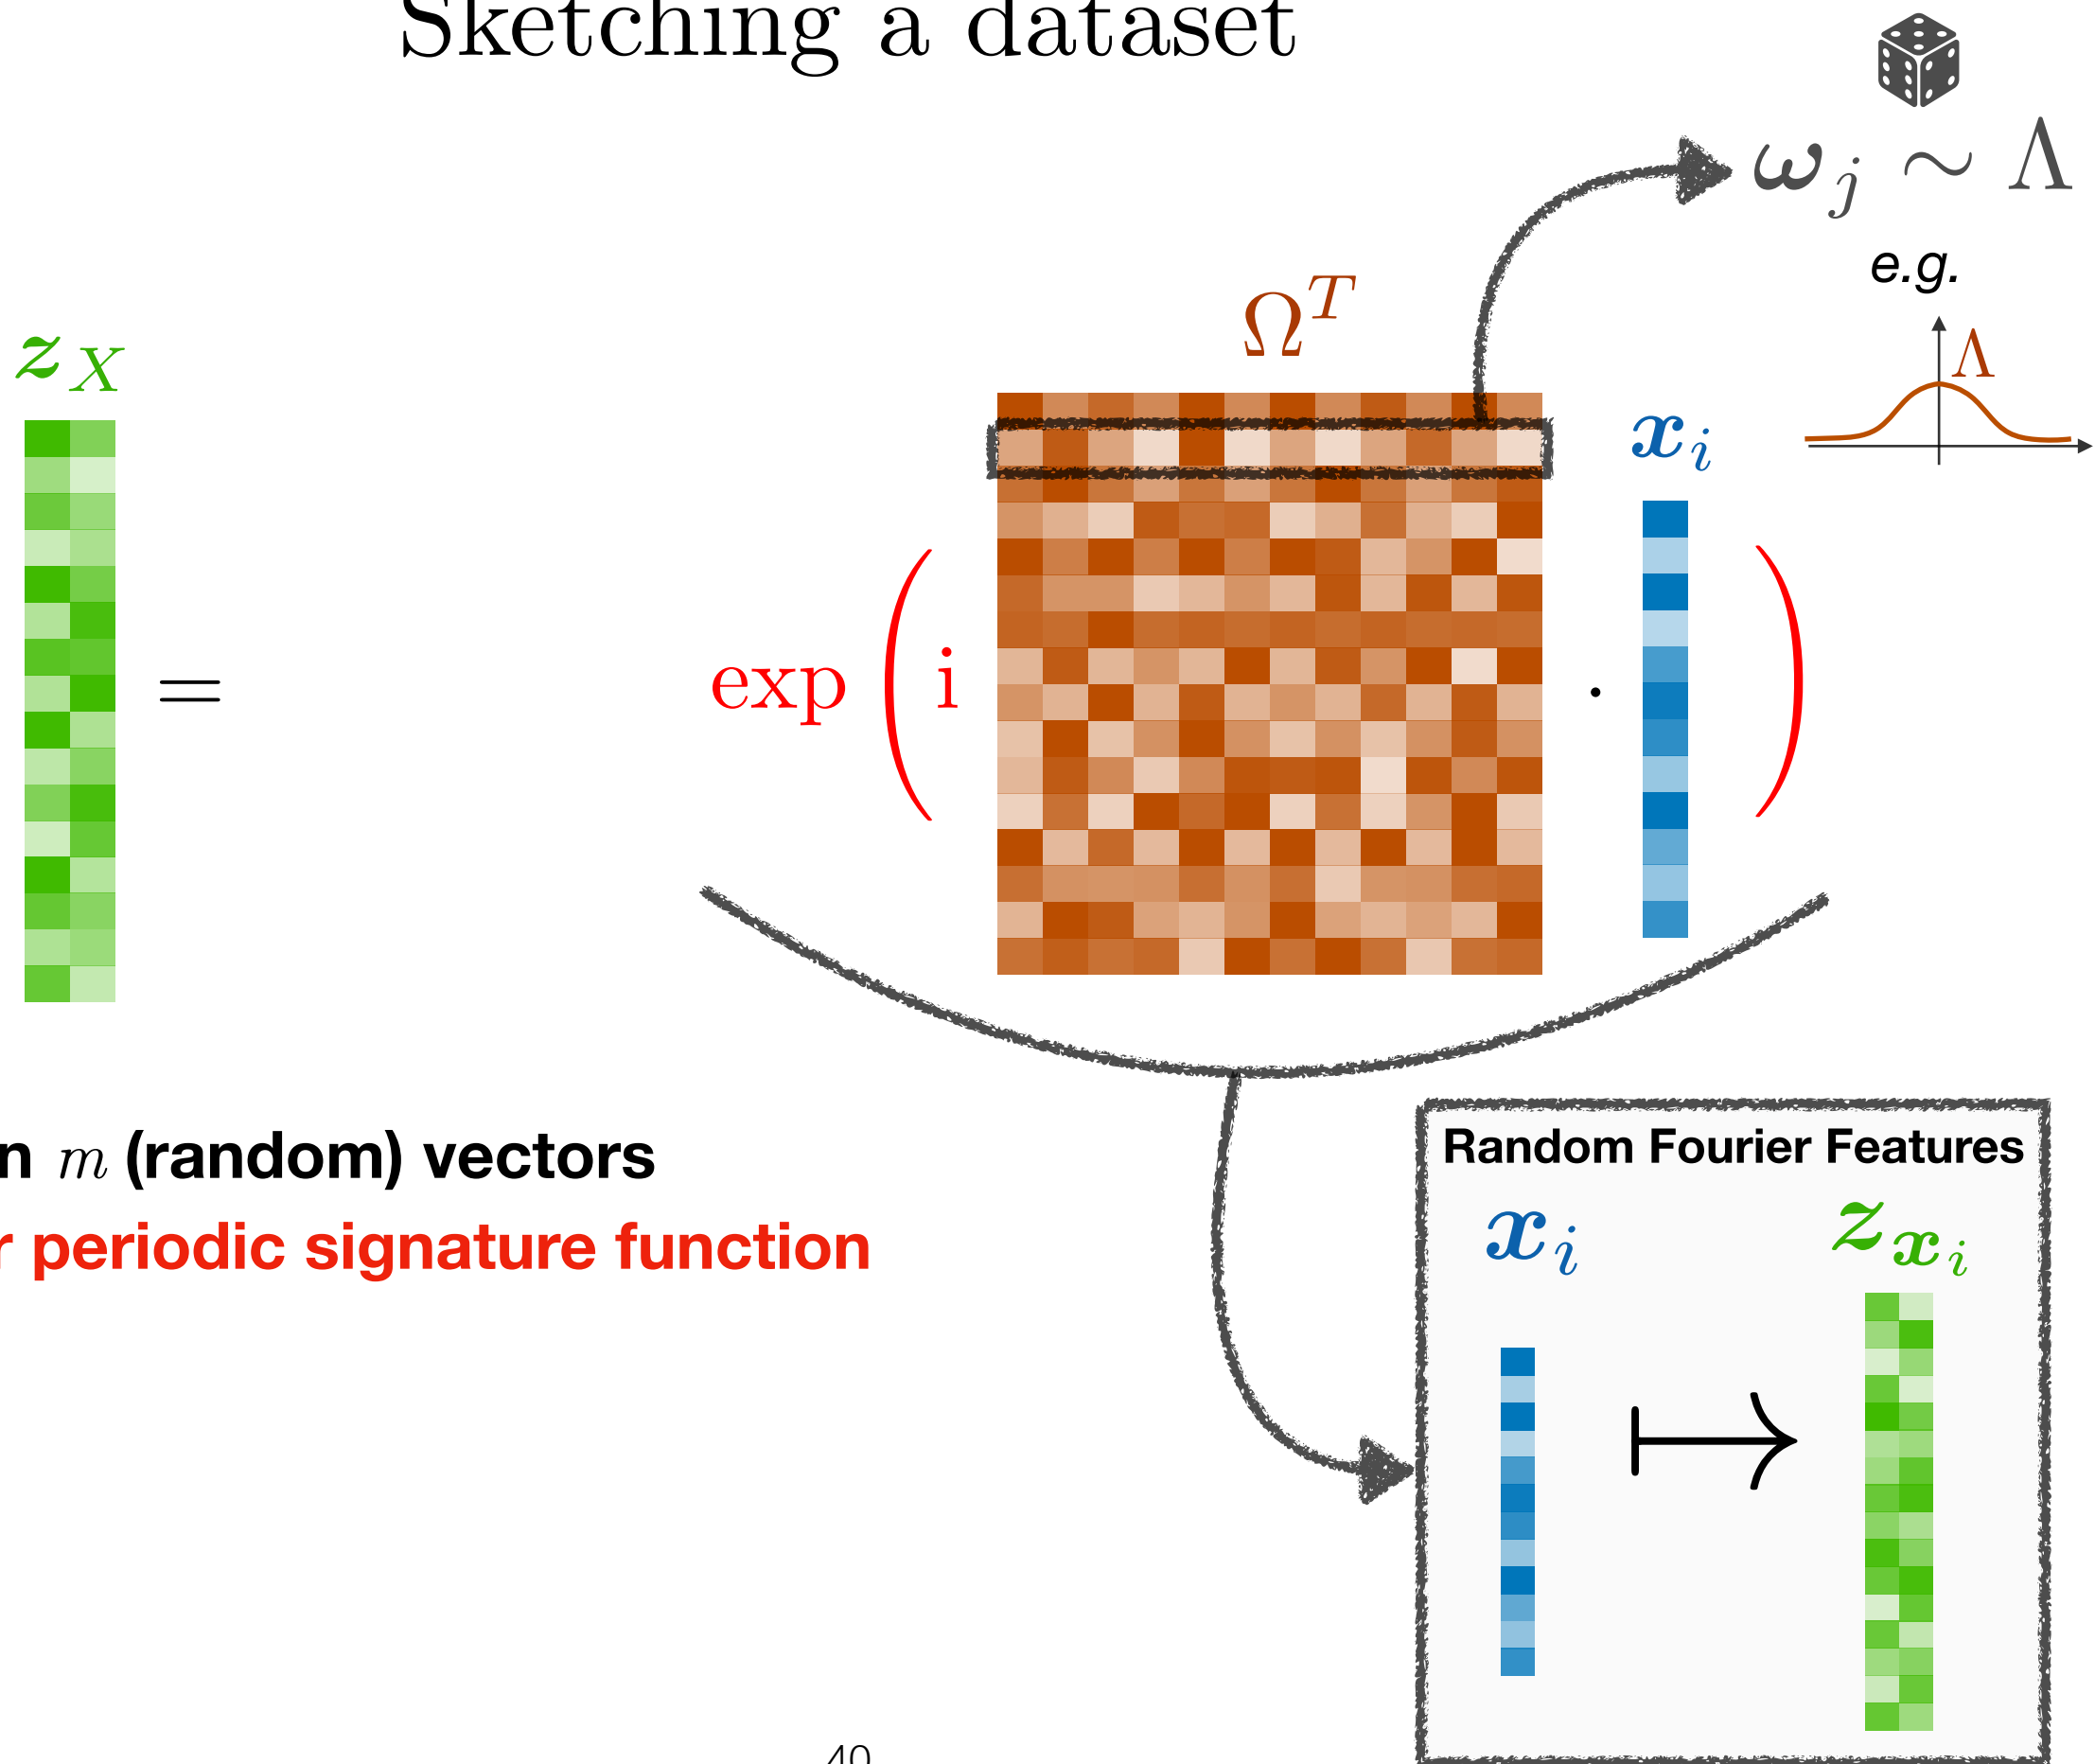


Sketching a dataset



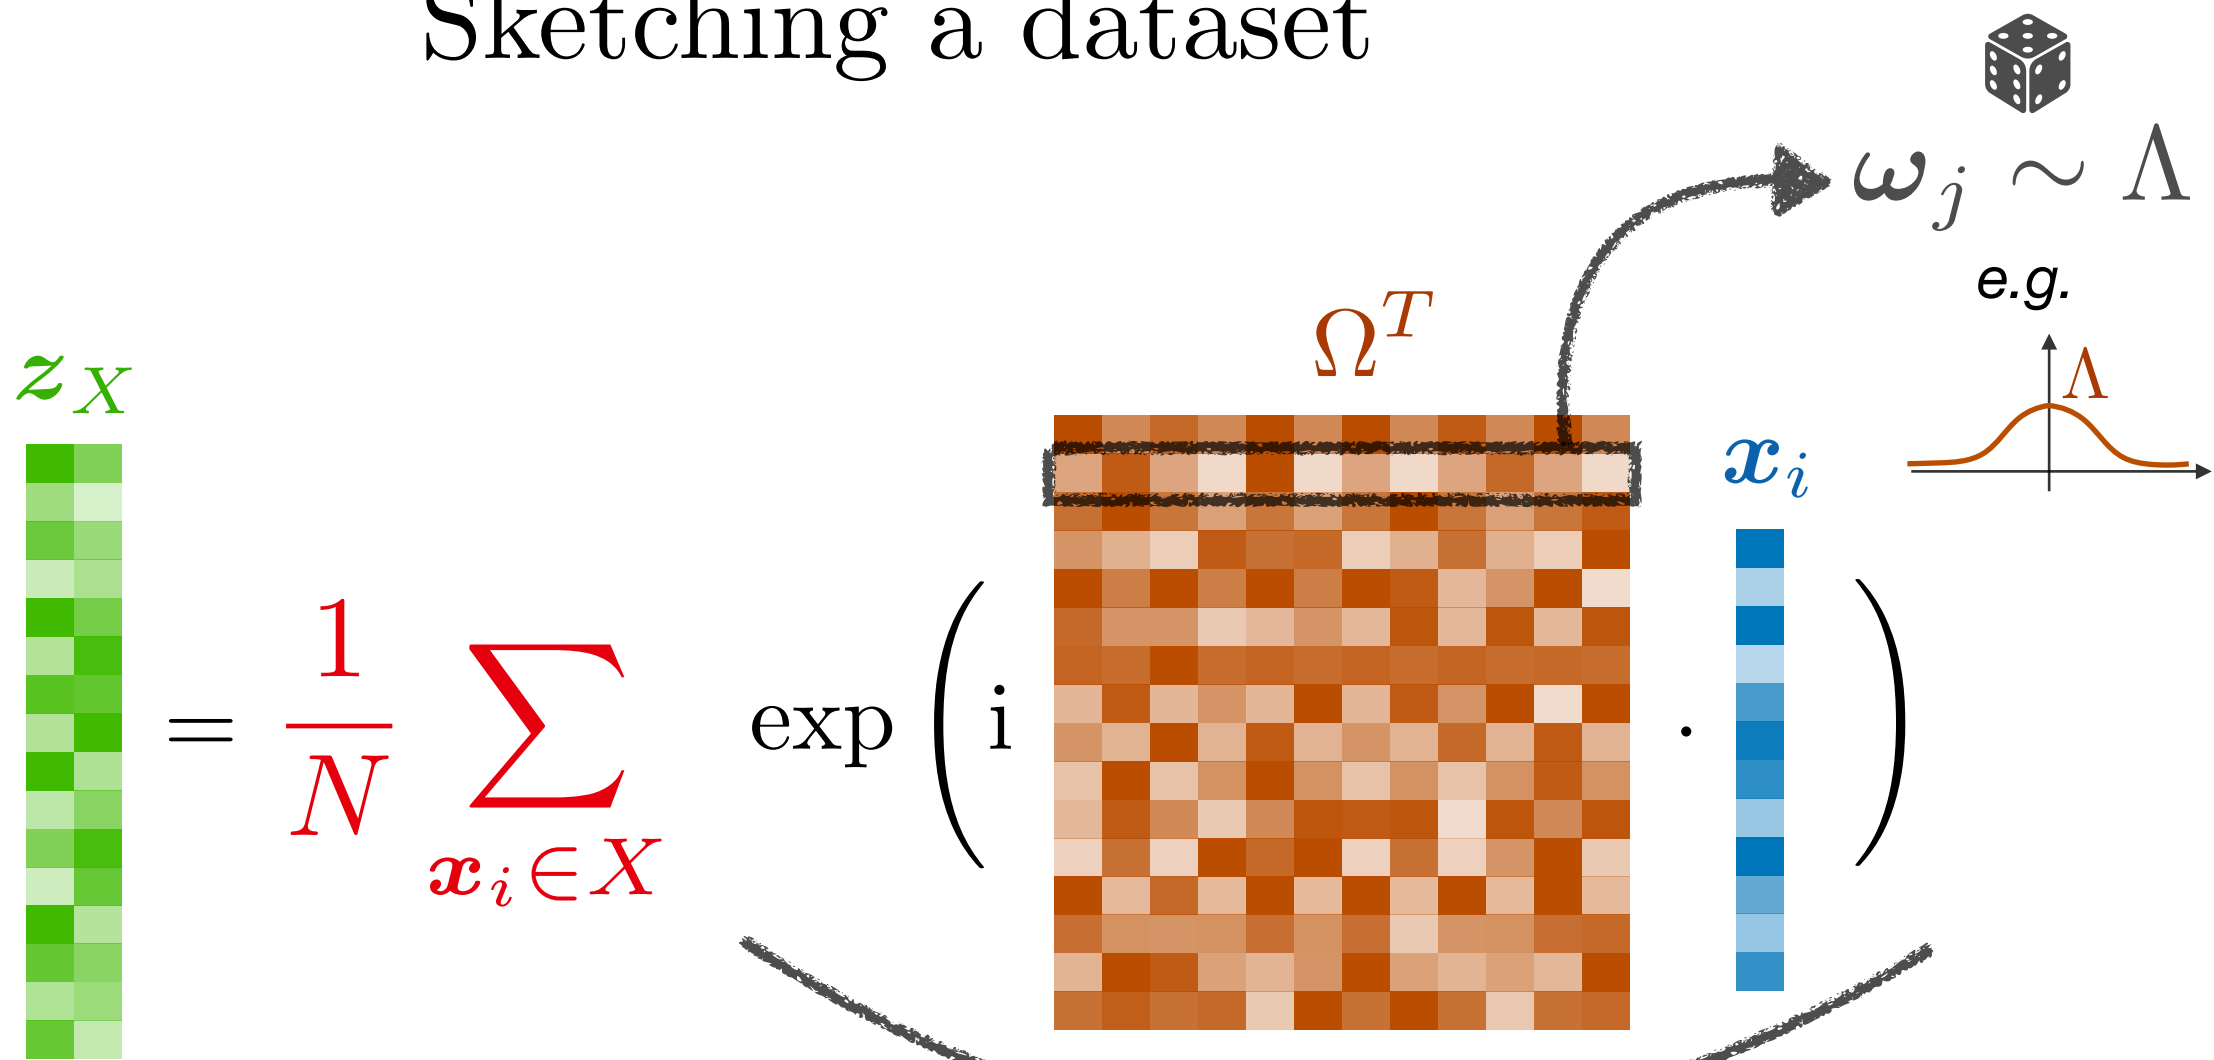
1. Project on m (random) vectors
2. Nonlinear periodic signature function

Sketching a dataset

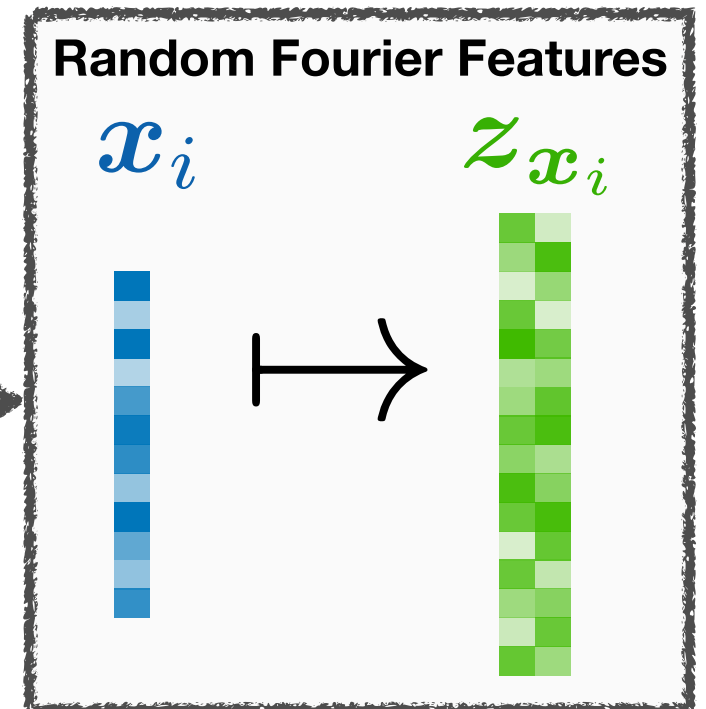


1. Project on m (random) vectors
2. Nonlinear periodic signature function

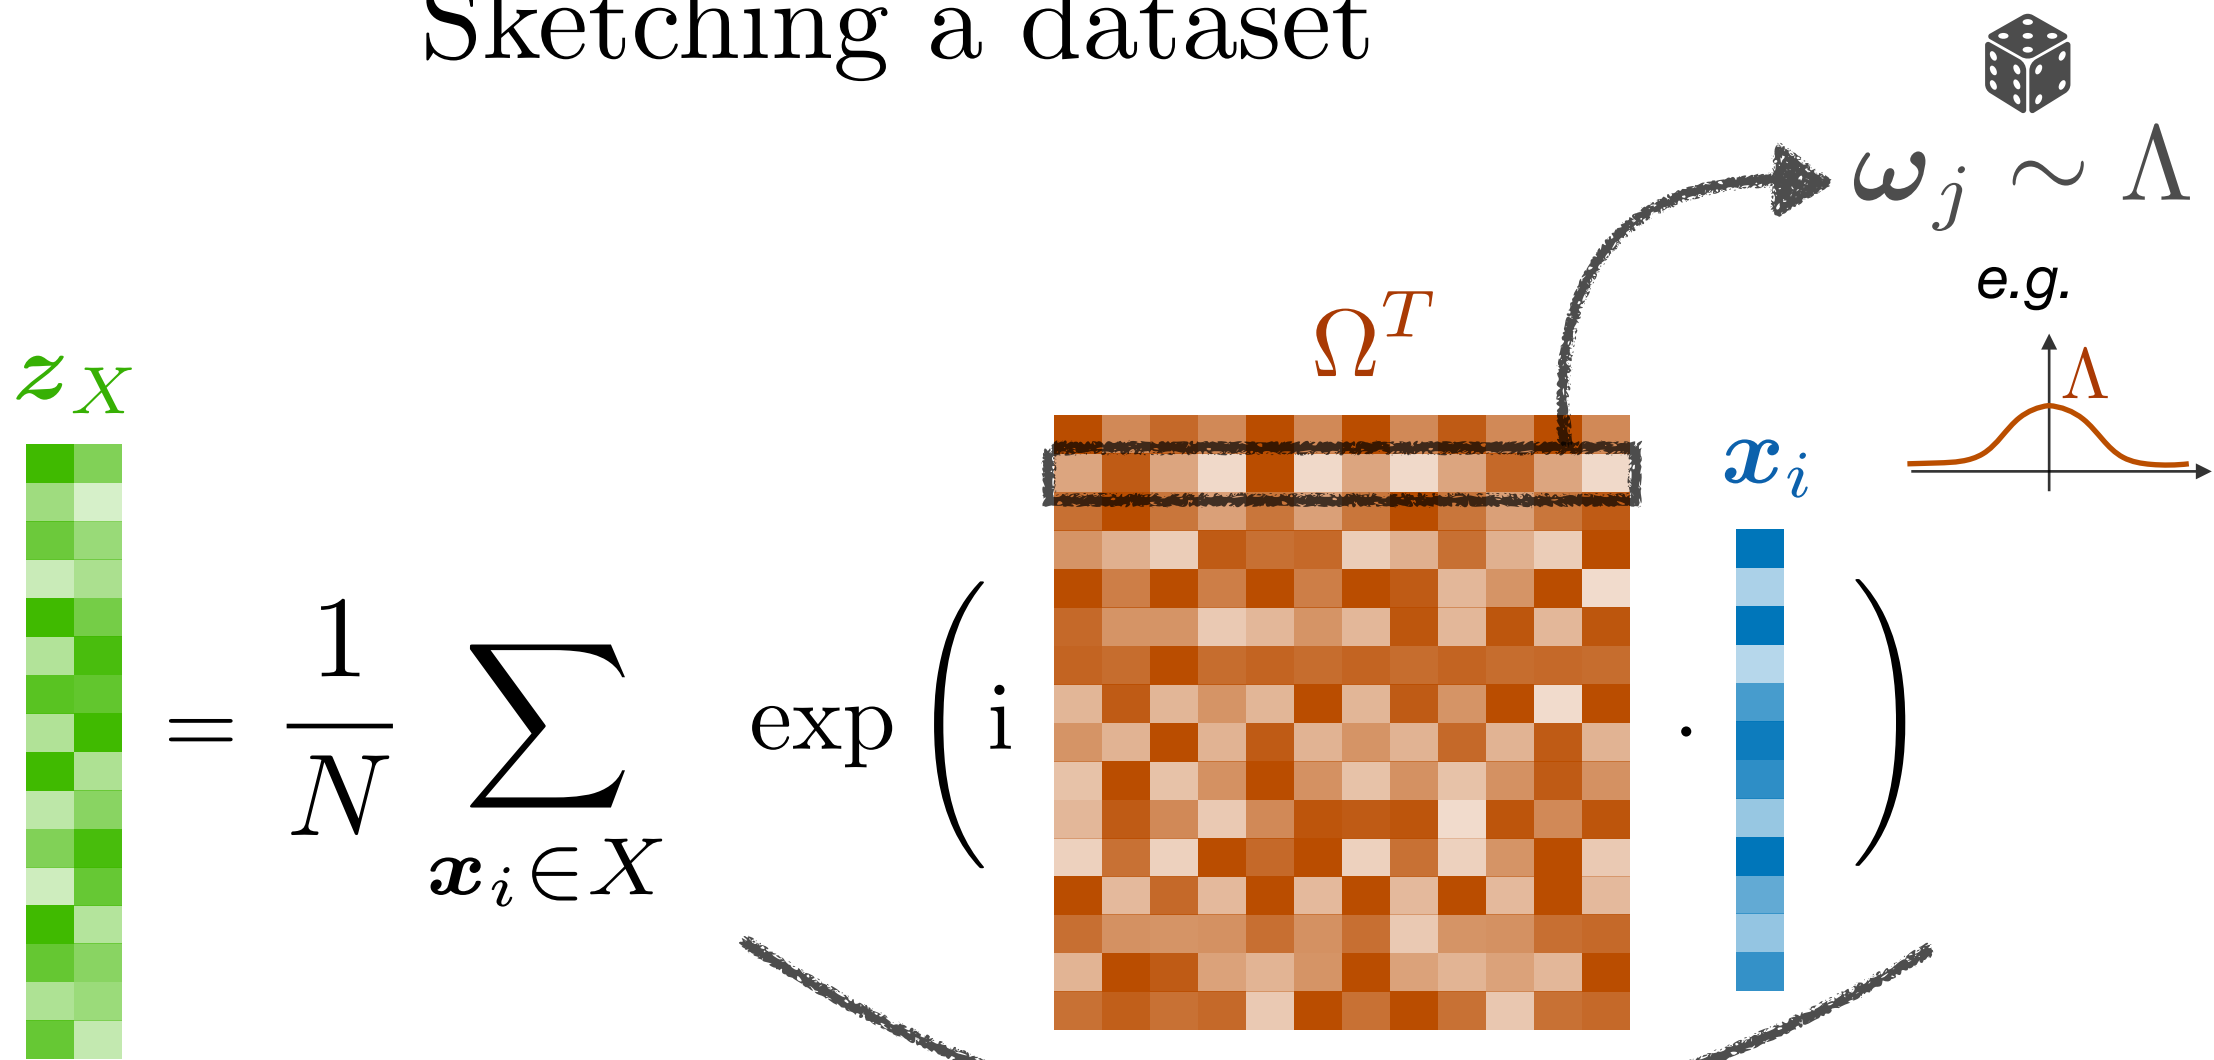
Sketching a dataset



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

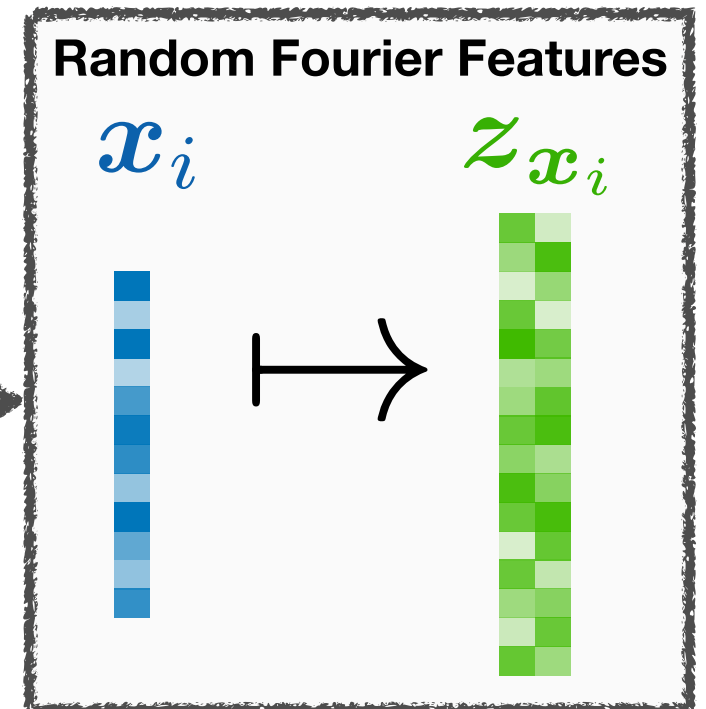


Sketching a dataset



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

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



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) 

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

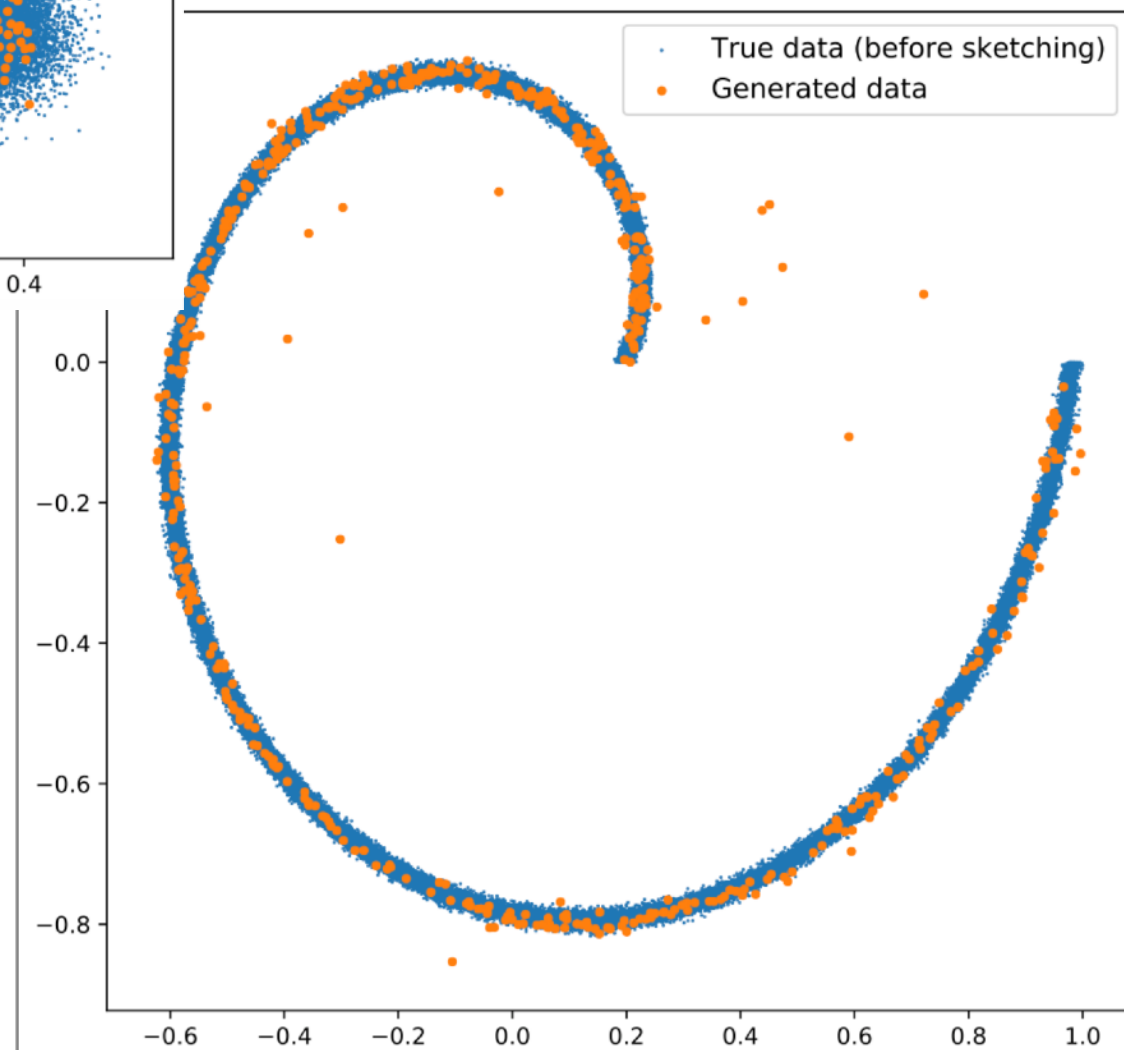
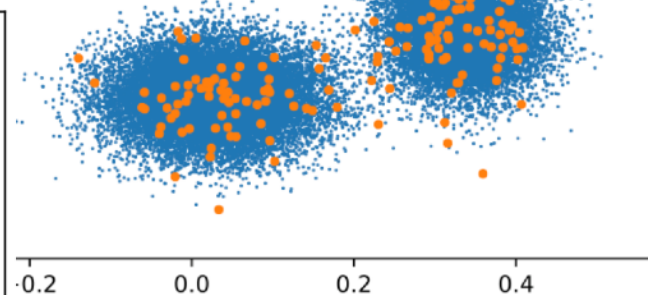
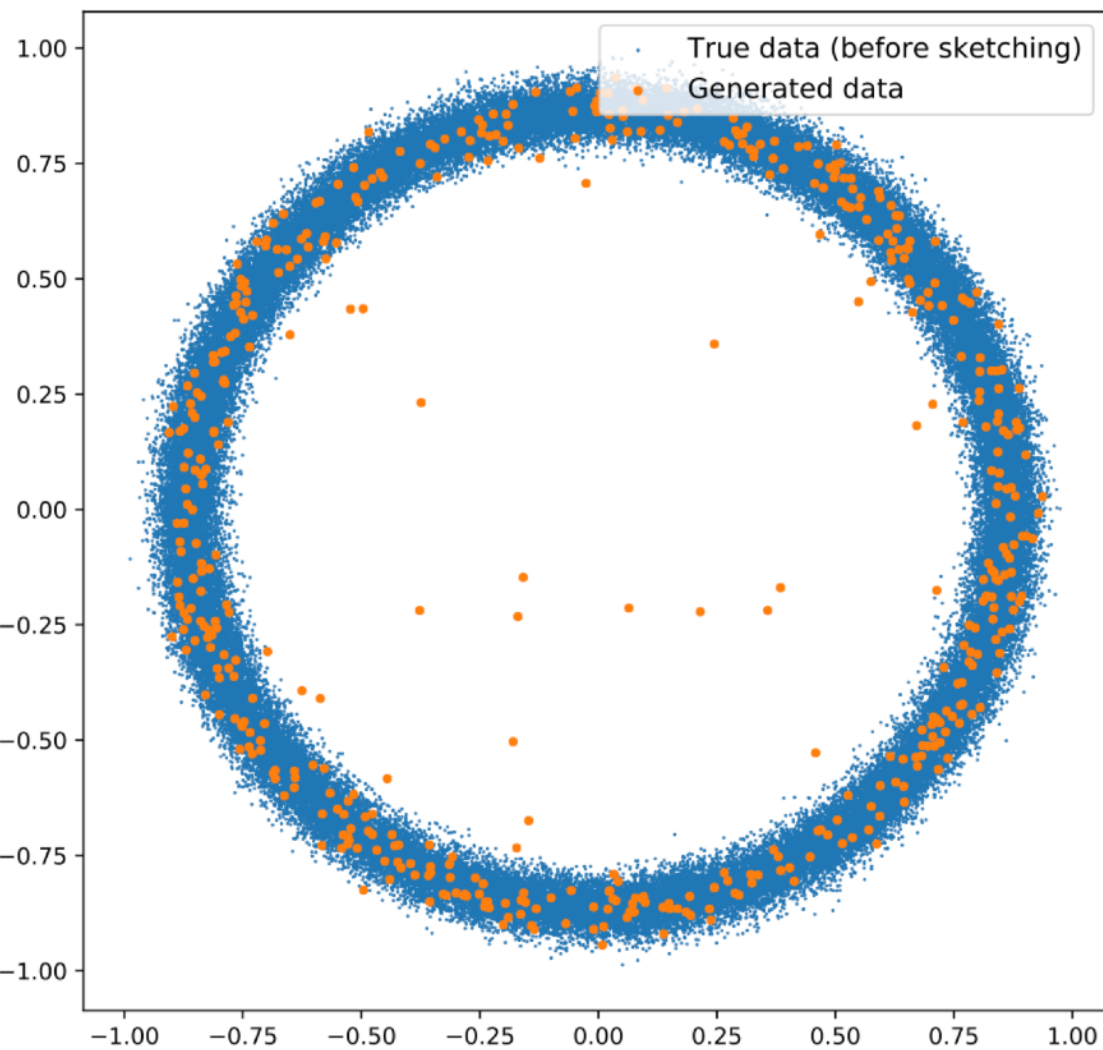
$$\mathcal{O}(N')$$

Results: preliminary toy example



Simple 2d signals...

...it works (in principle)!



Results (from other authors)

Differentially Private Mean Embeddings with Random Features for Synthetic Data Generation

Frederik Harder^{*,1,2}, Kamil Adamczewski^{*,1,3}, Mijung Park^{1,2}

^{*}Equal Contribution

¹Max Planck Institute for Intelligent Systems

²University of Tübingen

³ETH Zürich

{fharder|kadamczewski|mpark}@tue.mpg.de

Abstract

We present a differentially private data generation paradigm using random feature representations of kernel mean embeddings when comparing the distribution of true data with that of synthetic data. We exploit the random feature representations for two important benefits. First, we require a very low privacy cost for training deep generative models. This is because unlike kernel-based distance metrics that require computing the kernel matrix on all pairs of true and synthetic data points, we can detach the data-dependent term from the term solely dependent on synthetic data. Hence, we no and then use it until the end analytic sensitivity of the kernel bounded by construction. This for a clipping norm to handle provide several variants of our with random features (DP-M for datasets such as heterogeneous achieves better privacy-utility several datasets.

Submission history

From: Frederik Harder [[view email](#)]

[v1] Wed, 26 Feb 2020 16:41:41 UTC (1,629 KB)

- Use the same cost function
- Rely on privacy-friendly benefits of sketching
- Rely on pre-trained autoencoders to handle images

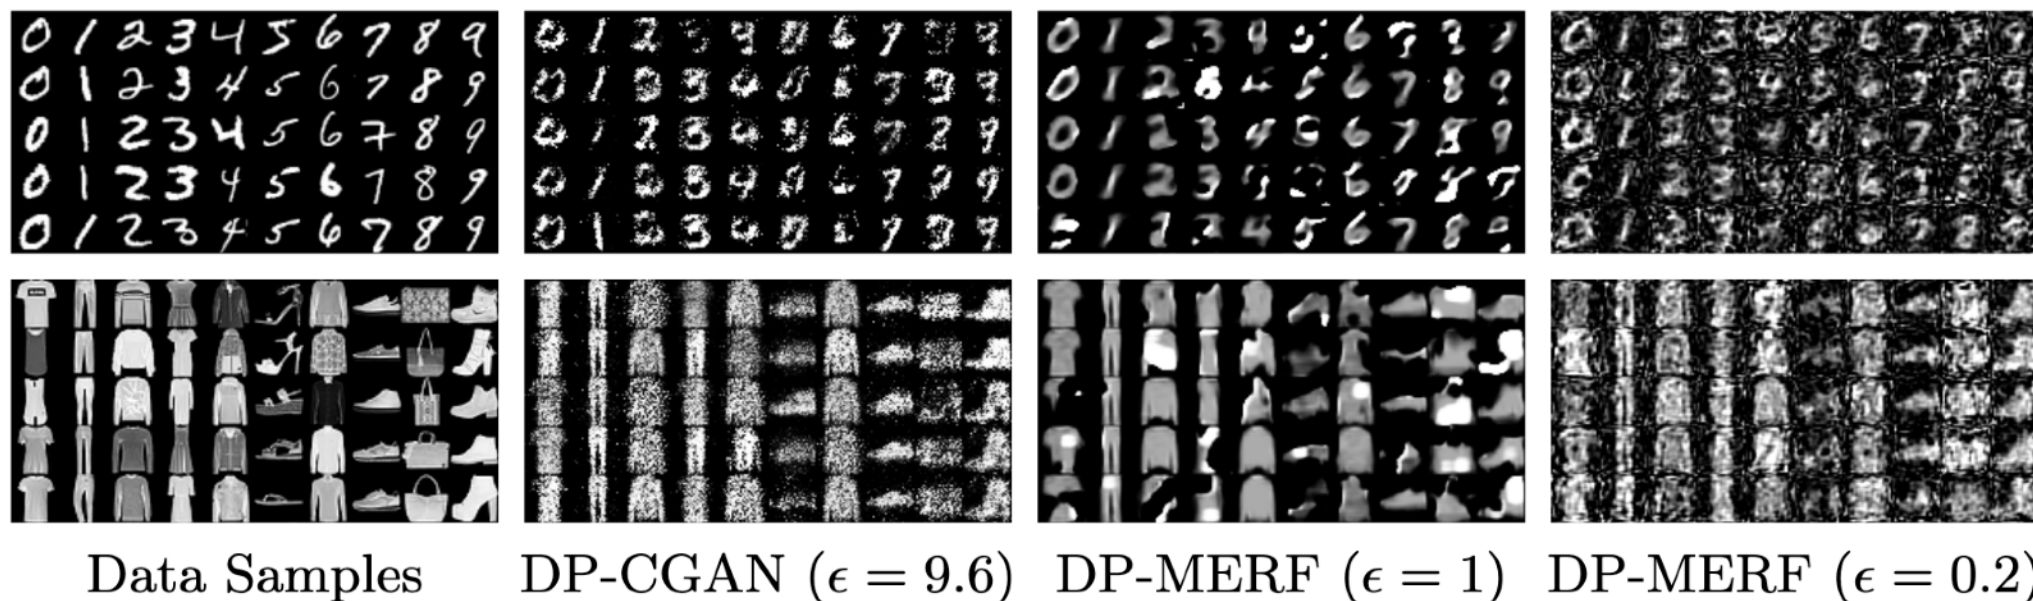


Figure 1: Generated samples with different levels of privacy

To conclude: open challenges

