Deep Neural Networks

Deep Generative Modelling

Hellenic Mediterranean University
Department of Electrical and Computer Engineering
Dr. Konstantinos Karampidis

Which face is real?

A

B

C

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Supervised vs unsupervised learning

Supervised Learning

Data: (x, y)
x is data, y is label

Goal: Learn function to map
$$x \rightarrow y$$

Examples: Classification, regression, object detection, semantic segmentation, etc.

Unsupervised Learning

Data: x
x is data, no labels!

Goal: Learn some *hidden* or *underlying structure* of the data

Examples: Clustering, feature or dimensionality reduction, etc.

MIT 6.S191

Supervised vs unsupervised learning

Supervised Learning

Data: (x, y)
x is data, y is label

Goal: Learn function to map
$$x \rightarrow y$$

Examples: Classification, regression, object detection, semantic segmentation, etc.

Unsupervised Learning

Data: x
x is data, no labels!

Goal: Learn the *hidden* or *underlying structure* of the data

Examples: Clustering, feature or dimensionality reduction, etc.

MIT 6.S191

Generative modeling

Goal: Take as input training samples from some distribution and learn a model that represents that distribution

Density Estimation

Sample Generation

Input samples

Training data ~ $P_{data}(x)$

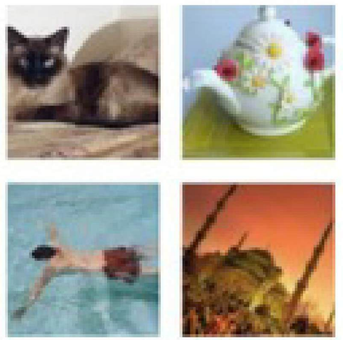

Generated samples

Generated ~ $P_{model}(x)$

How can we learn $P_{model}(x)$ similar to $P_{data}(x)$?

MIT 6.S191

Why generative models? Debiasing

Capable of uncovering **underlying features** in a dataset

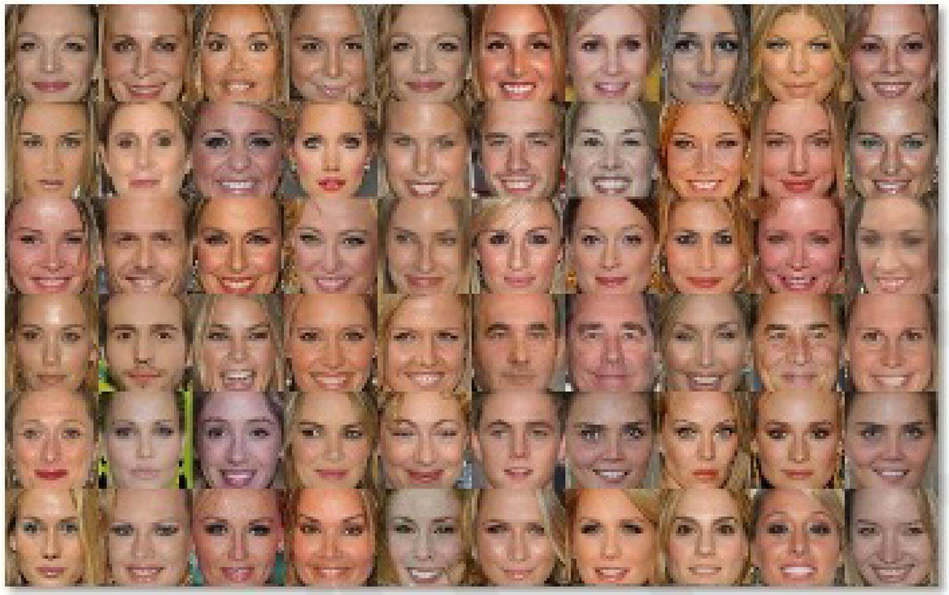

Homogeneous skin color, pose

VS

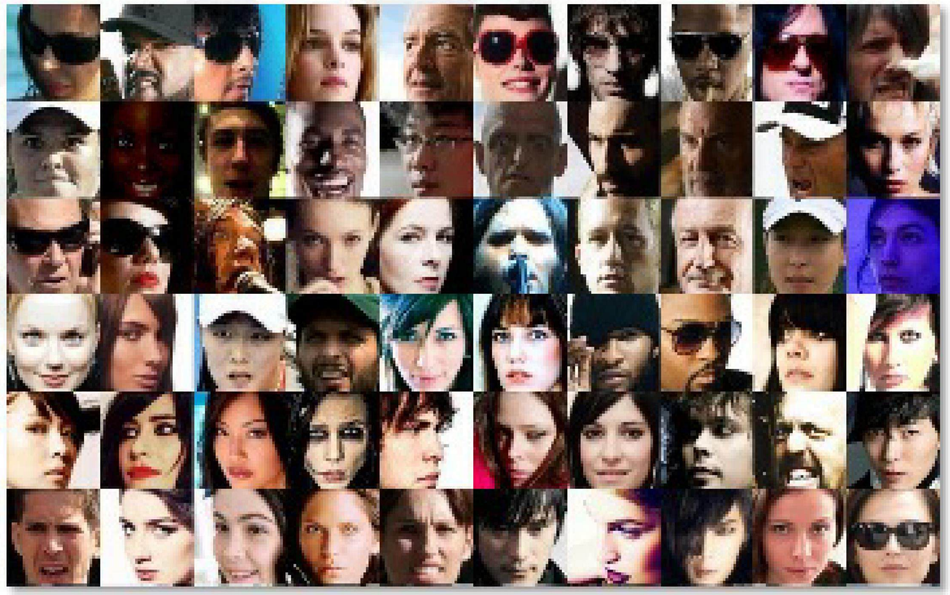

Diverse skin color, pose, illumination

How can we use this information to create fair and representative datasets?

Software Lab!

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Why generative models? Outlier detection

- **Problem:** How can we detect when we encounter something new or rare?
- **Strategy:** Leverage generative models, detect outliers in the distribution
- Use outliers during training to improve even more!

95% of Driving Data:
(1) sunny, (2) highway, (3) straight road

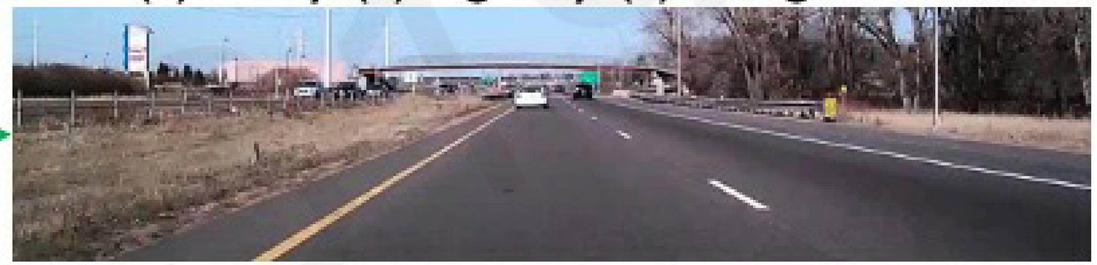

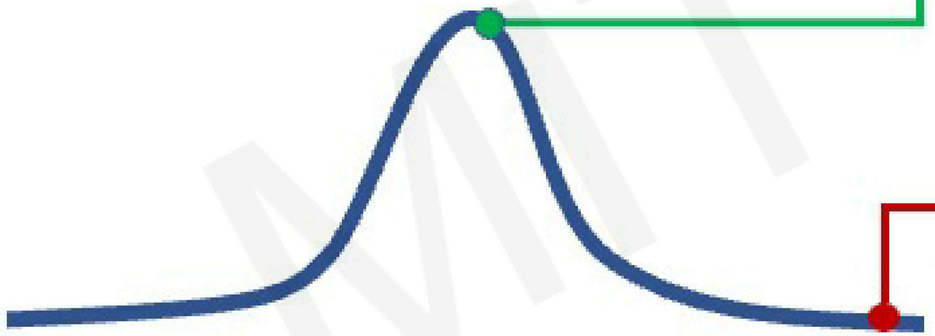

Detect outliers to avoid unpredictable behavior when training

Edge Cases

Harsh Weather

Pedestrians

MIT 6.S191

Why generative models? Sample generation

- **Generative models learn probability distributions**
- **Sampling** from that distribution → new data instances
- **Backbone of Generative AI:** generate language, images, and more

Natural Language

6.S191 Guest Lectures!

Images & Videos

Biology

6.S191 Guest Lecture!

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Latent variable models

Autoencoders and Variational Autoencoders (VAEs)

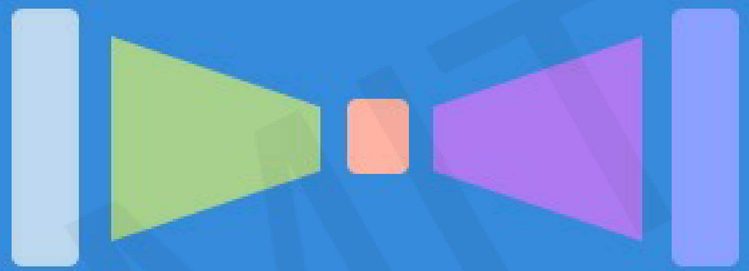

Generative Adversarial Networks (GANs)

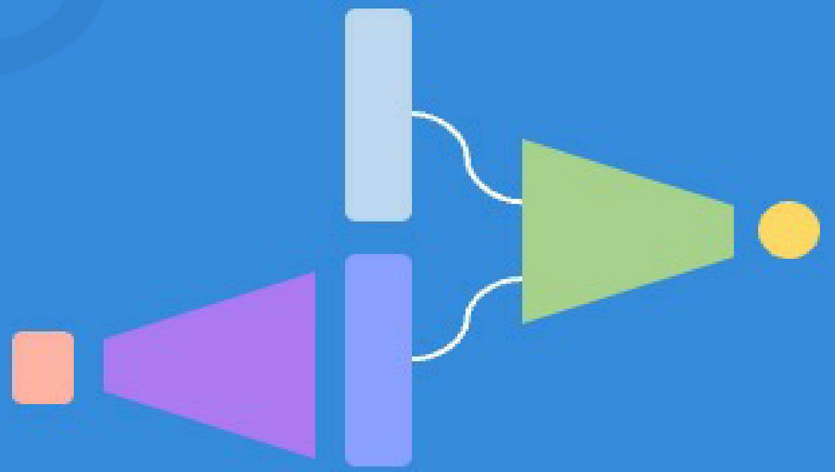

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

What is a latent variable?

Myth of the Cave

What is a latent variable?

Can we learn the **true explanatory factors**, e.g. latent variables, from only observed data?

Autoencoders

MIT 6.S191

Autoencoders: background

Unsupervised approach for learning a **lower-dimensional** feature representation from unlabeled training data

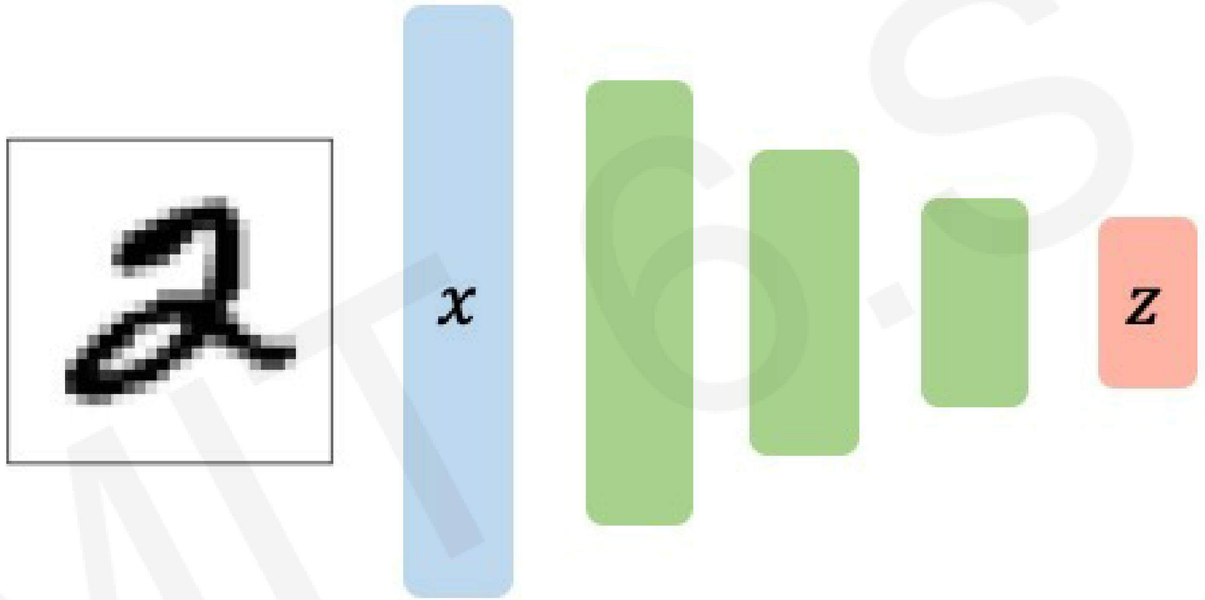

Why do we care about a low-dimensional z?

"Encoder" learns mapping from the data, x, to a low-dimensional latent space, z

MIT 6.S191

Autoencoders: background

How can we learn this latent space?
Train the model to use these features to **reconstruct the original data**

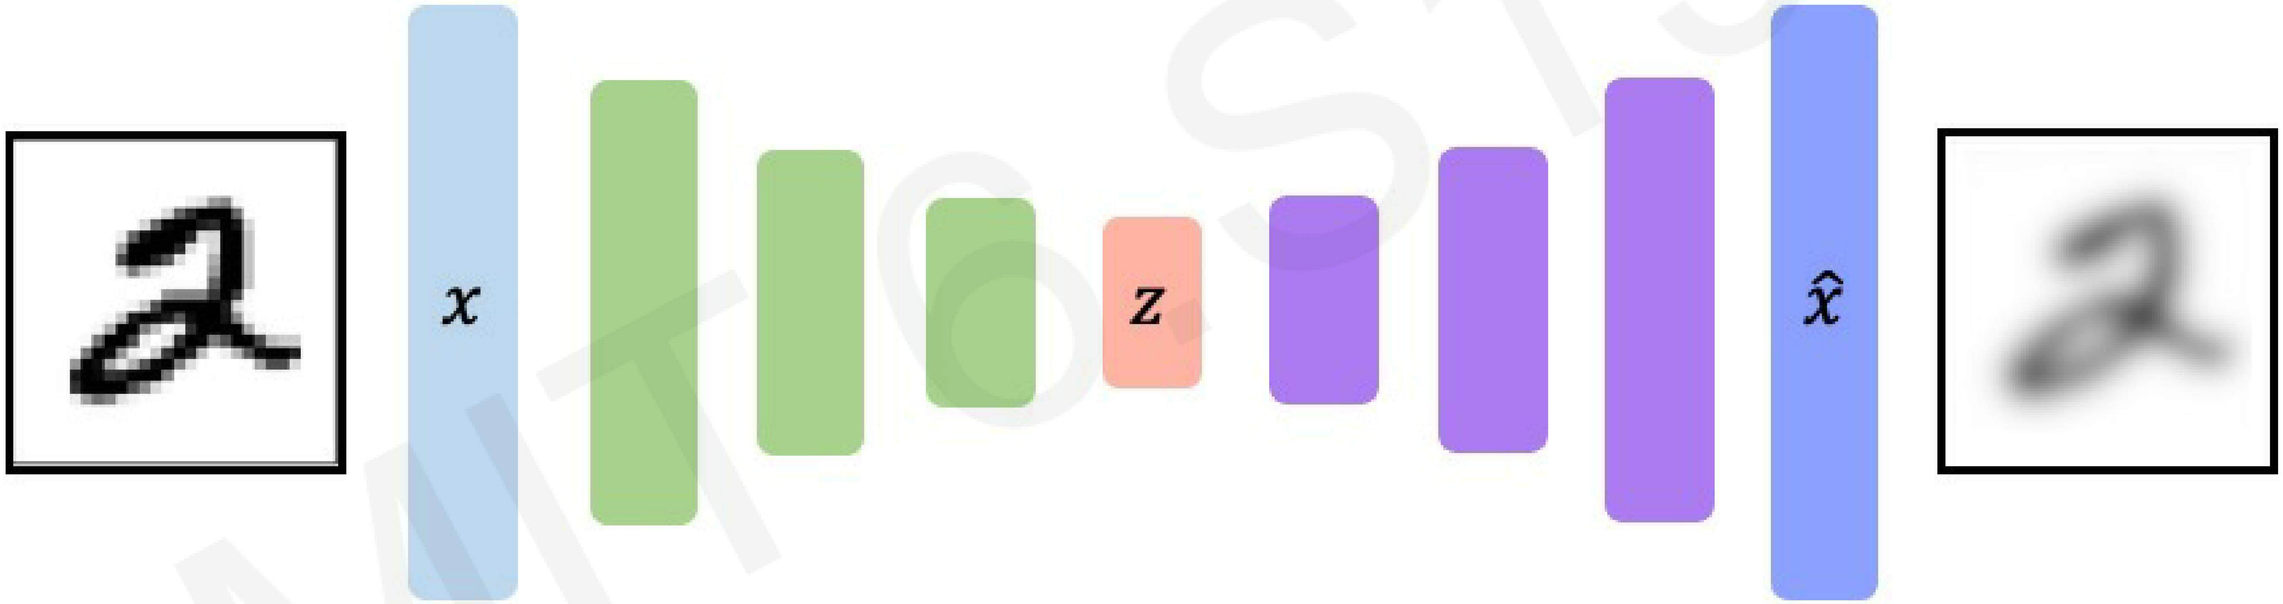

"Decoder" learns mapping back from latent space, z,
to a reconstructed observation, $\hat{x}$

MIT6.S191

Autoencoders: background

How can we learn this latent space?

Train the model to use these features to **reconstruct the original data**

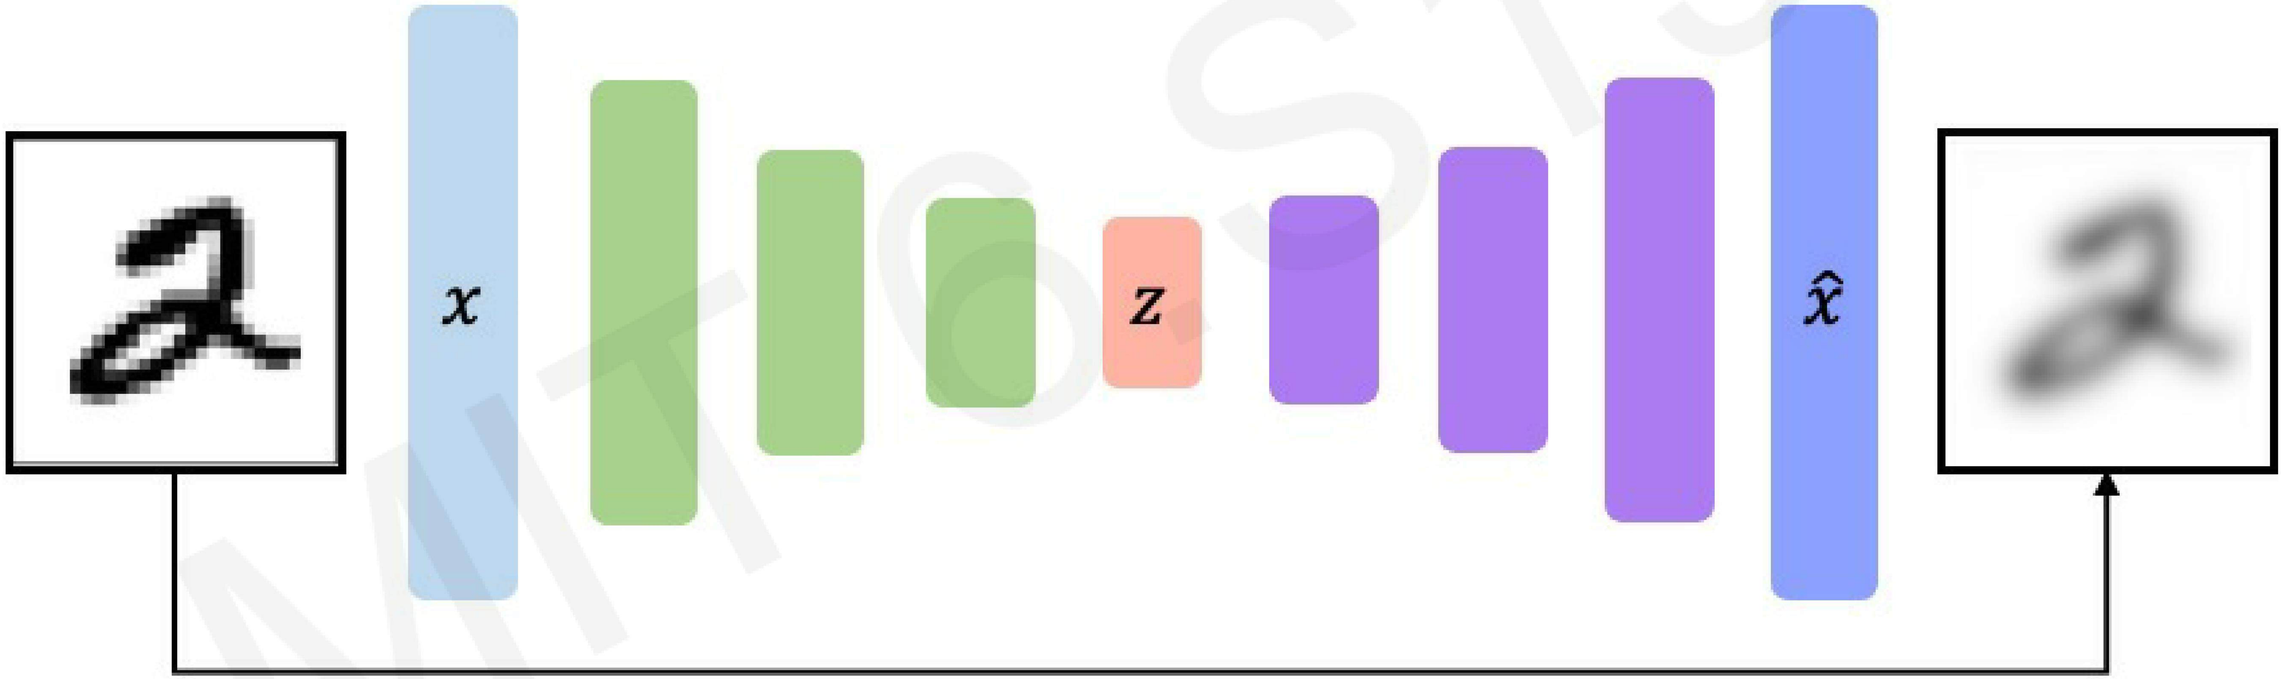

$$\mathcal{L}(x, \hat{x}) = \|x - \hat{x}\|^2$$

Loss function doesn't use any labels!!

MIT6.S191

Autoencoders: background

How can we learn this latent space?

Train the model to use these features to **reconstruct the original data**

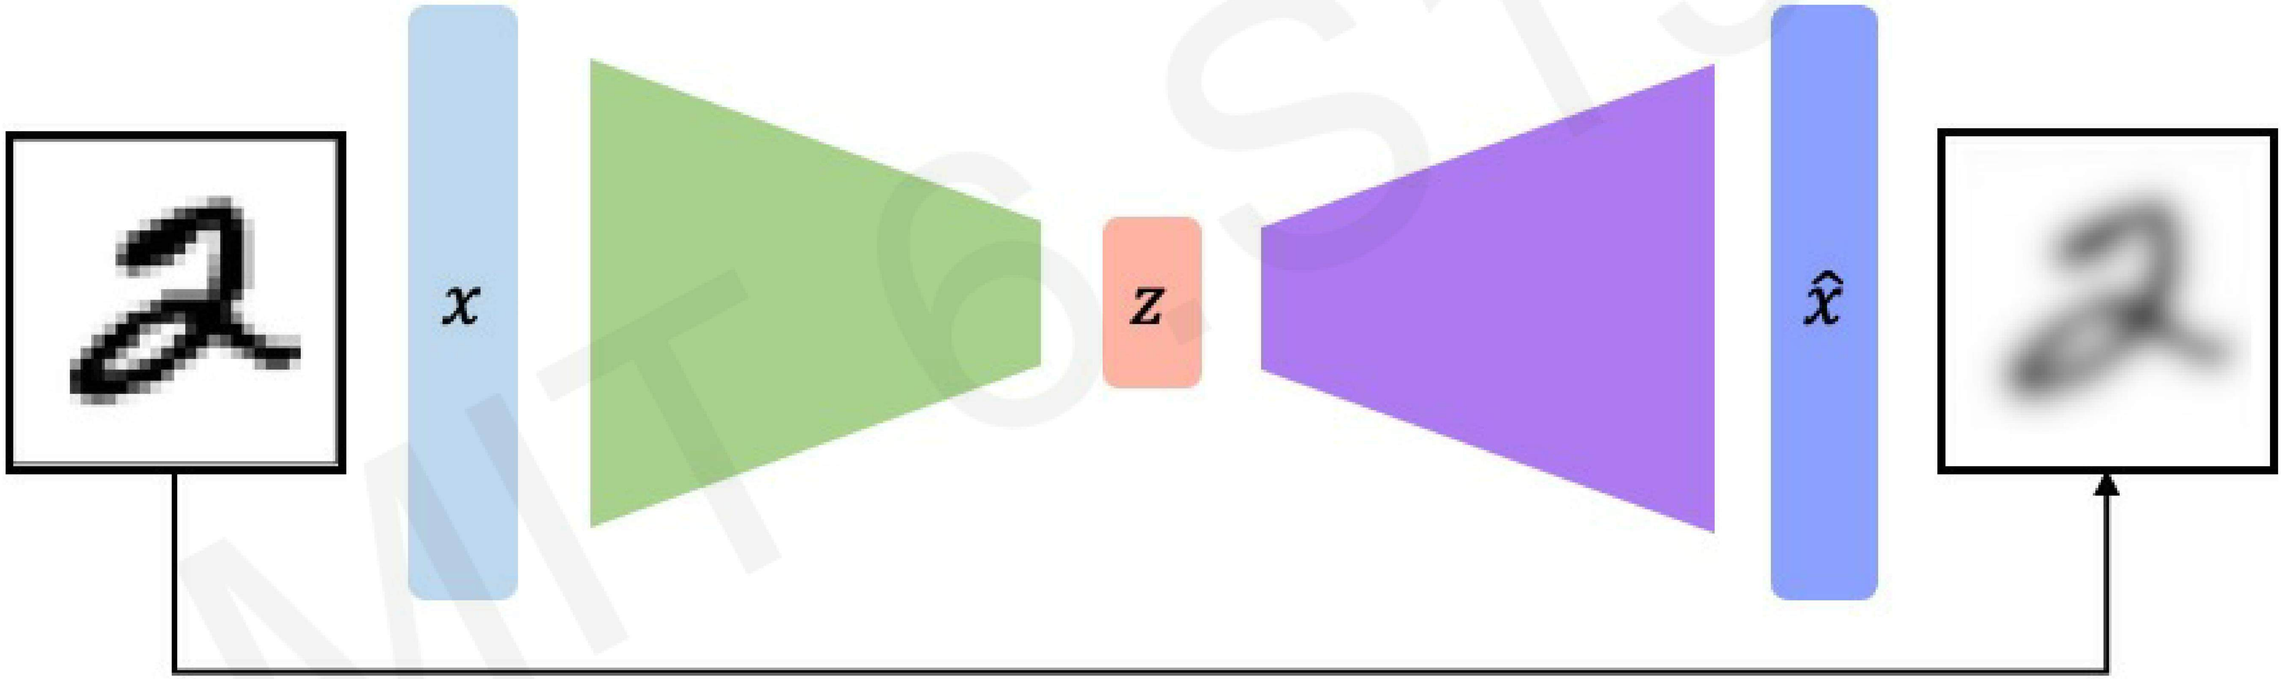

$$\mathcal{L}(x, \hat{x}) = \|x - \hat{x}\|^2$$

Loss function doesn't use any labels!!

MIT 6.S191

Dimensionality of latent space → reconstruction quality

Autoencoding is a form of compression!
Smaller latent space will force a larger training bottleneck

2D latent space

5D latent space

Ground Truth

Autoencoders for representation learning

Bottleneck hidden layer forces network to learn a compressed latent representation

Reconstruction loss forces the latent representation to capture (or encode) as much "information" about the data as possible

Autoencoding = **Auto**matically **encoding** data; "Auto" = **self**-encoding

MIT 6.S191

Variational Autoencoders (VAEs)

MIT 6.S191

Traditional autoencoders

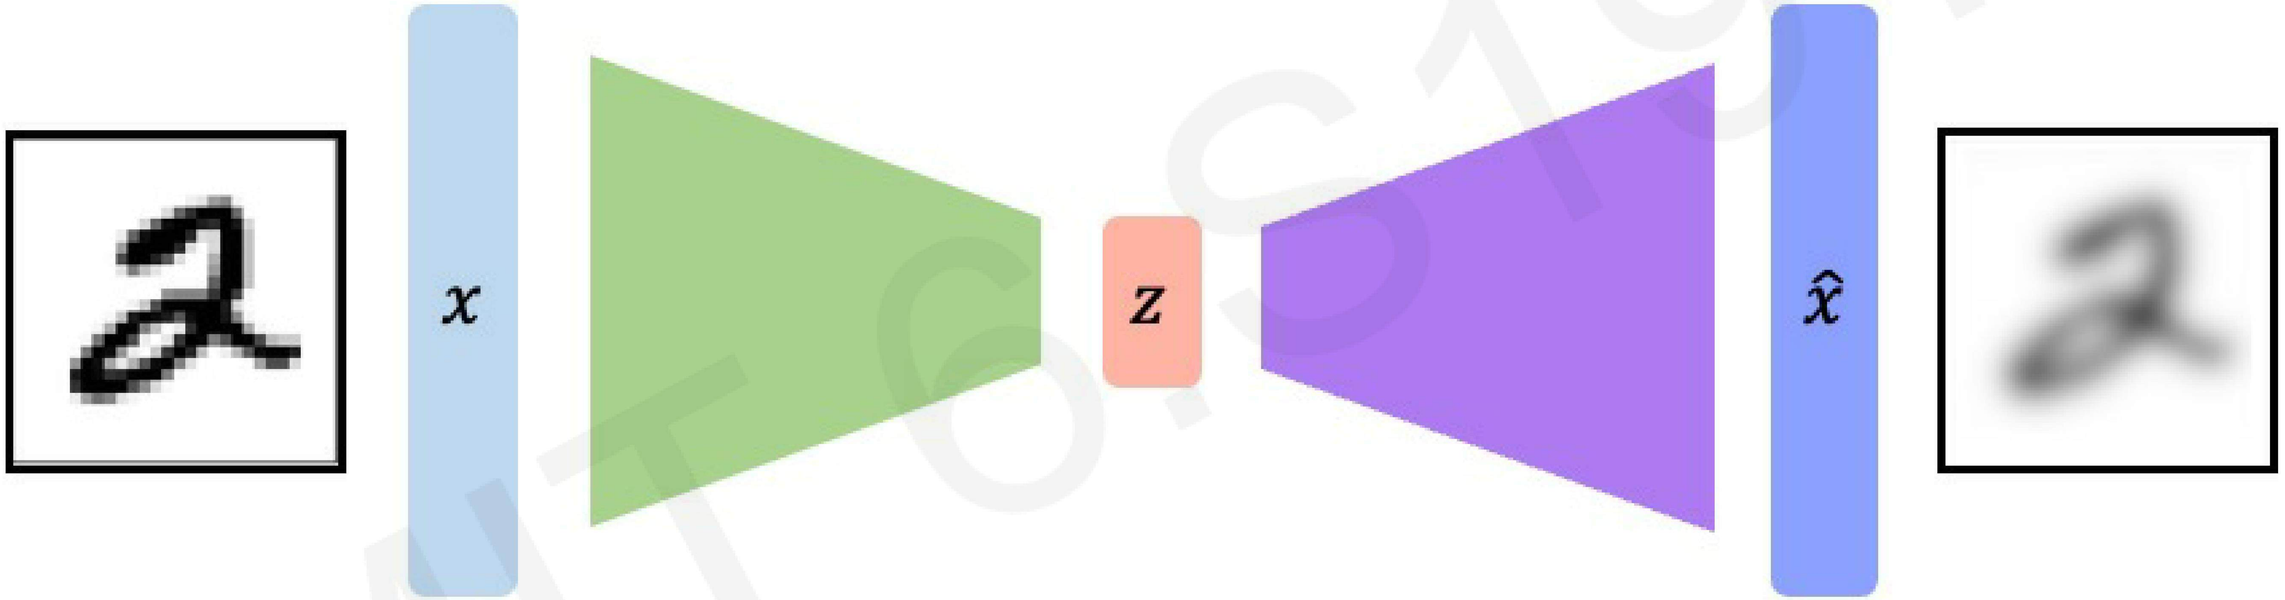

MIT 6.S191

VAEs: key difference with traditional autoencoder

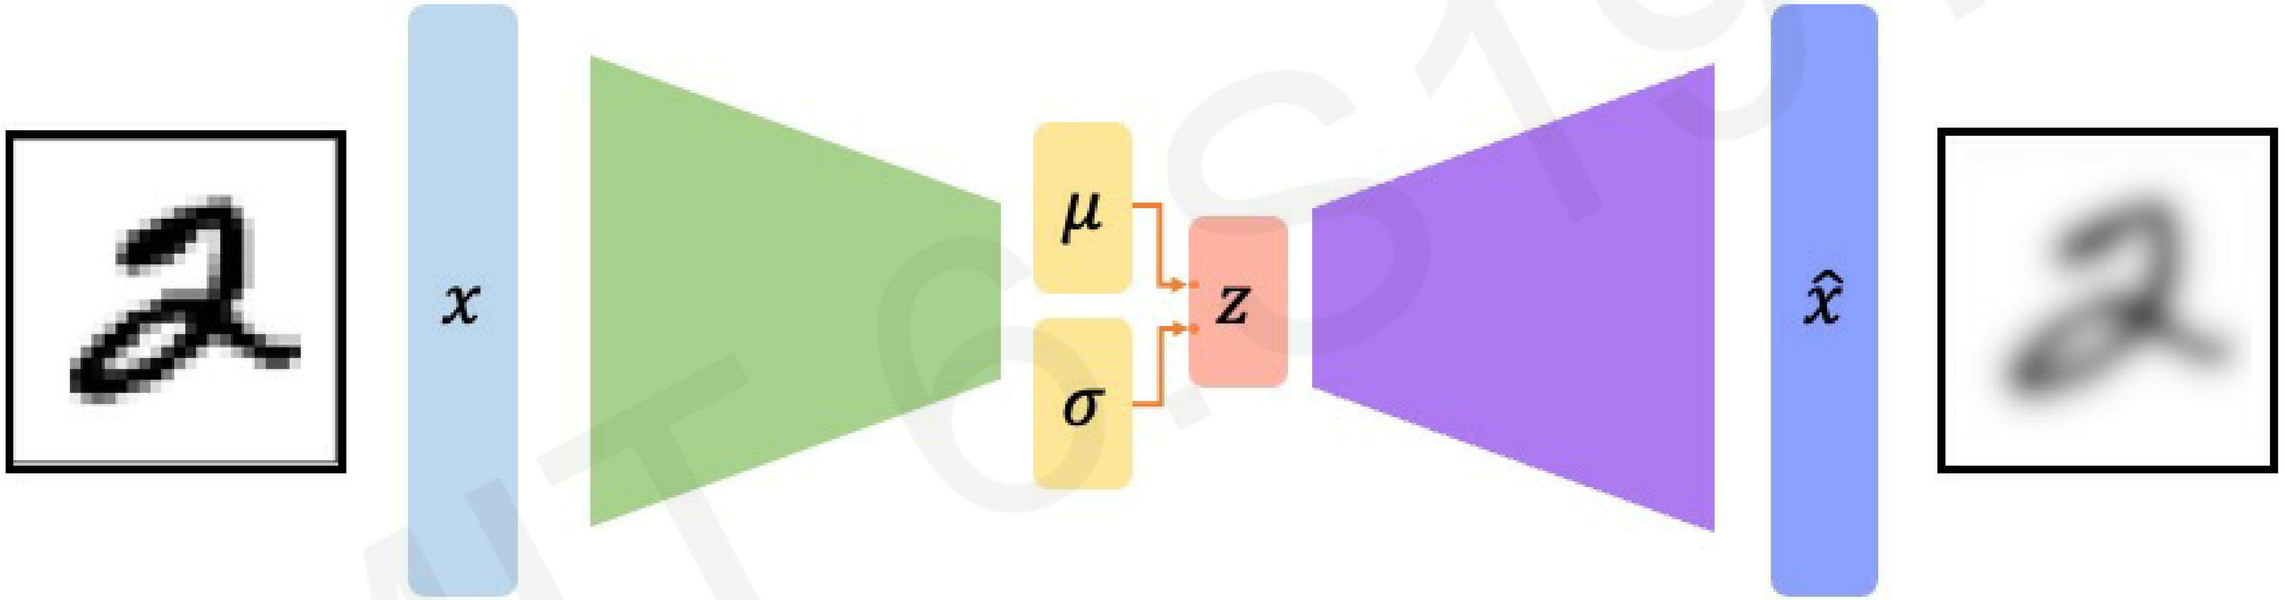

MIT 6.S191

VAEs: key difference with traditional autoencoder

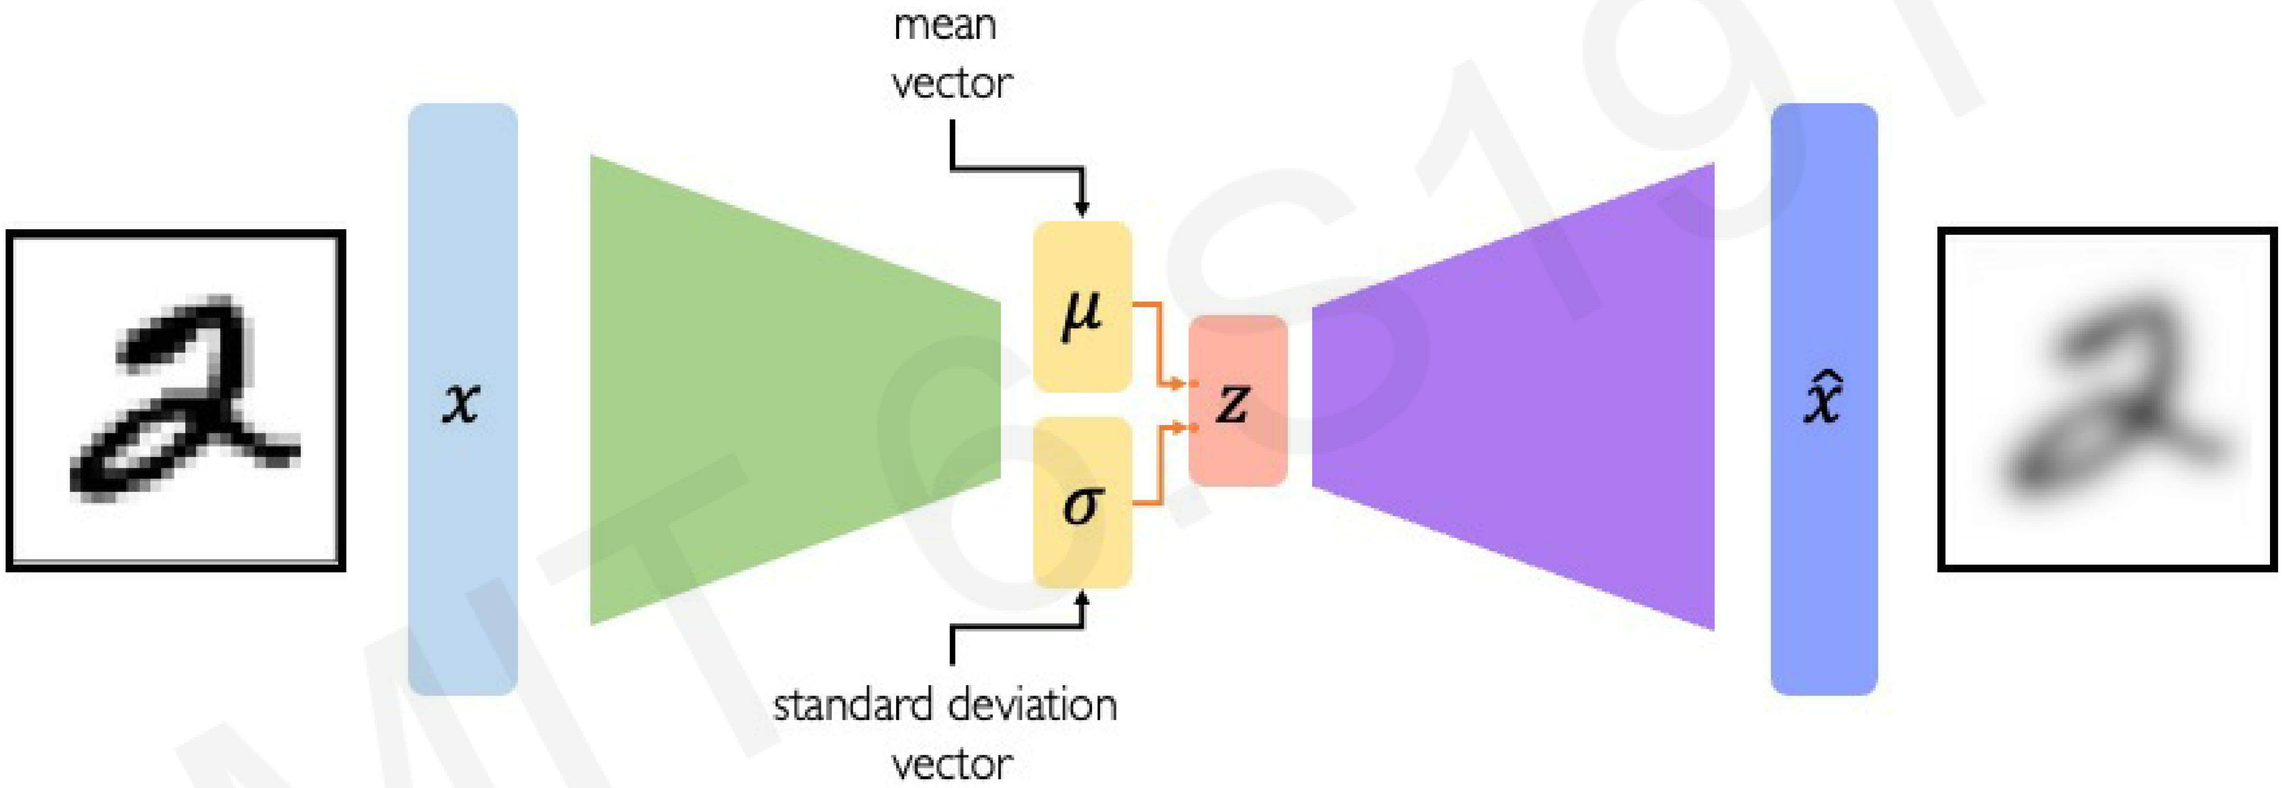

Variational autoencoders are a probabilistic twist on autoencoders!

Sample from the mean and standard deviation to compute latent sample

MIT 6.S191

VAE optimization

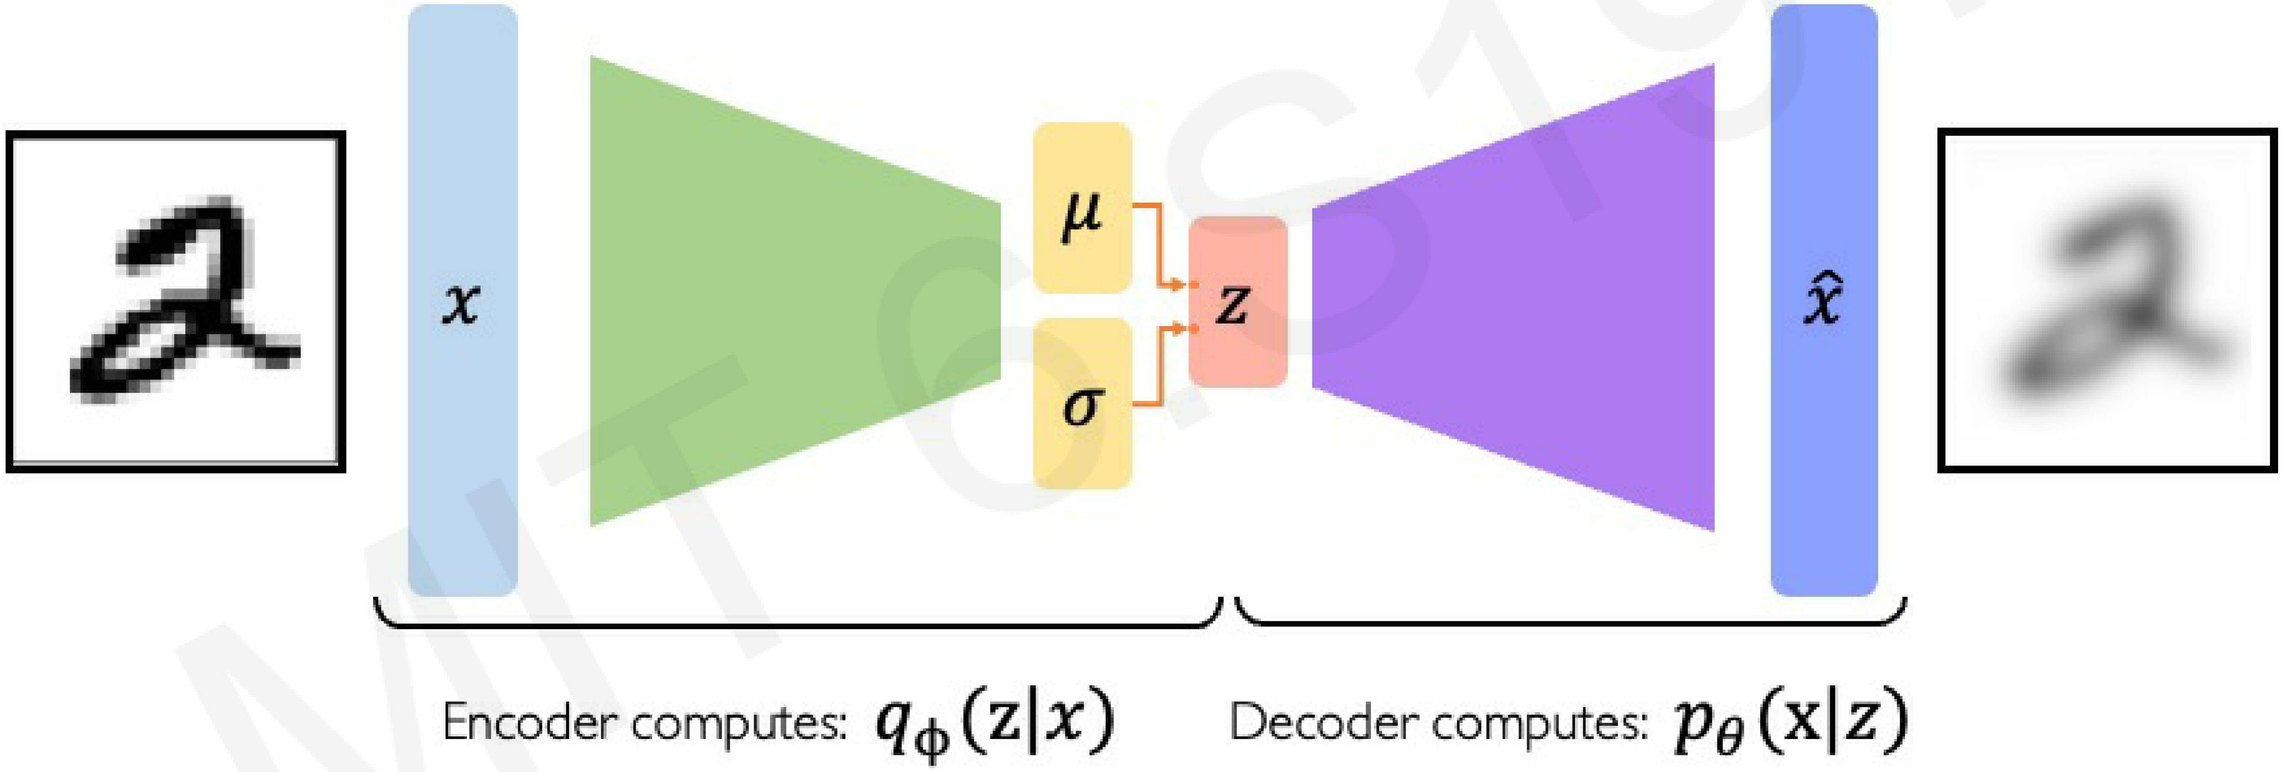

MIT 6.S191

VAE optimization

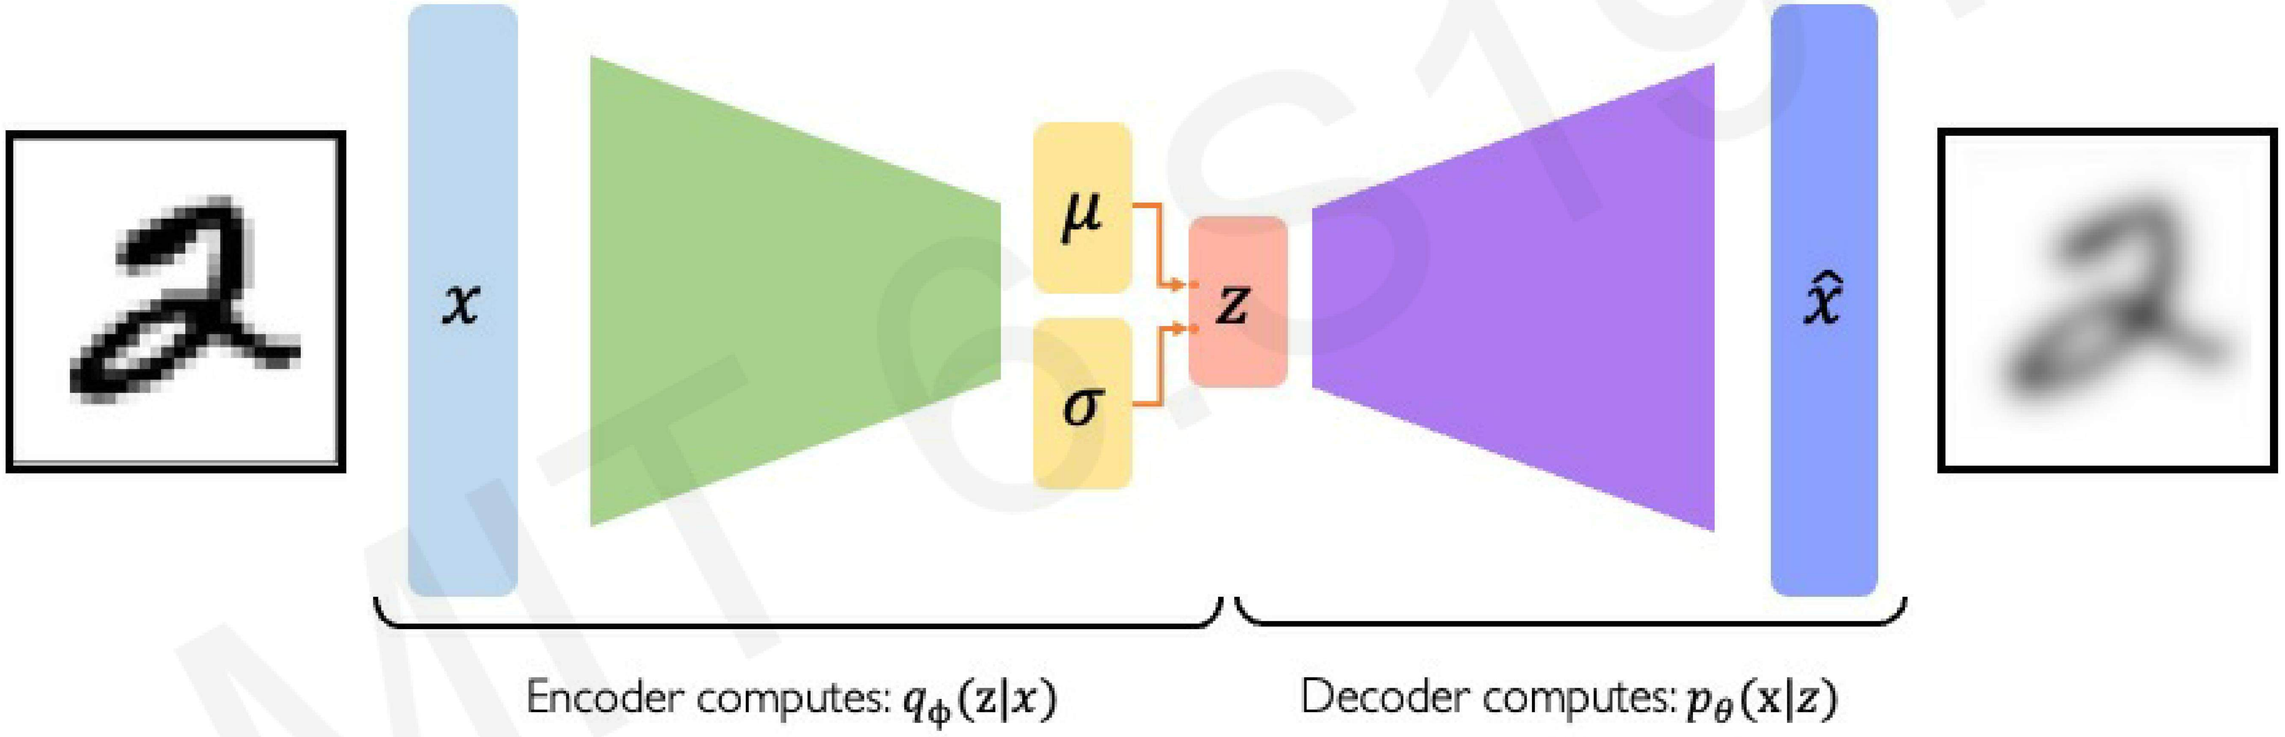

$$\mathcal{L}(\phi, \theta, x) = \text{(reconstruction loss)} + \text{(regularization term)}$$

MIT 6.S191

VAE optimization

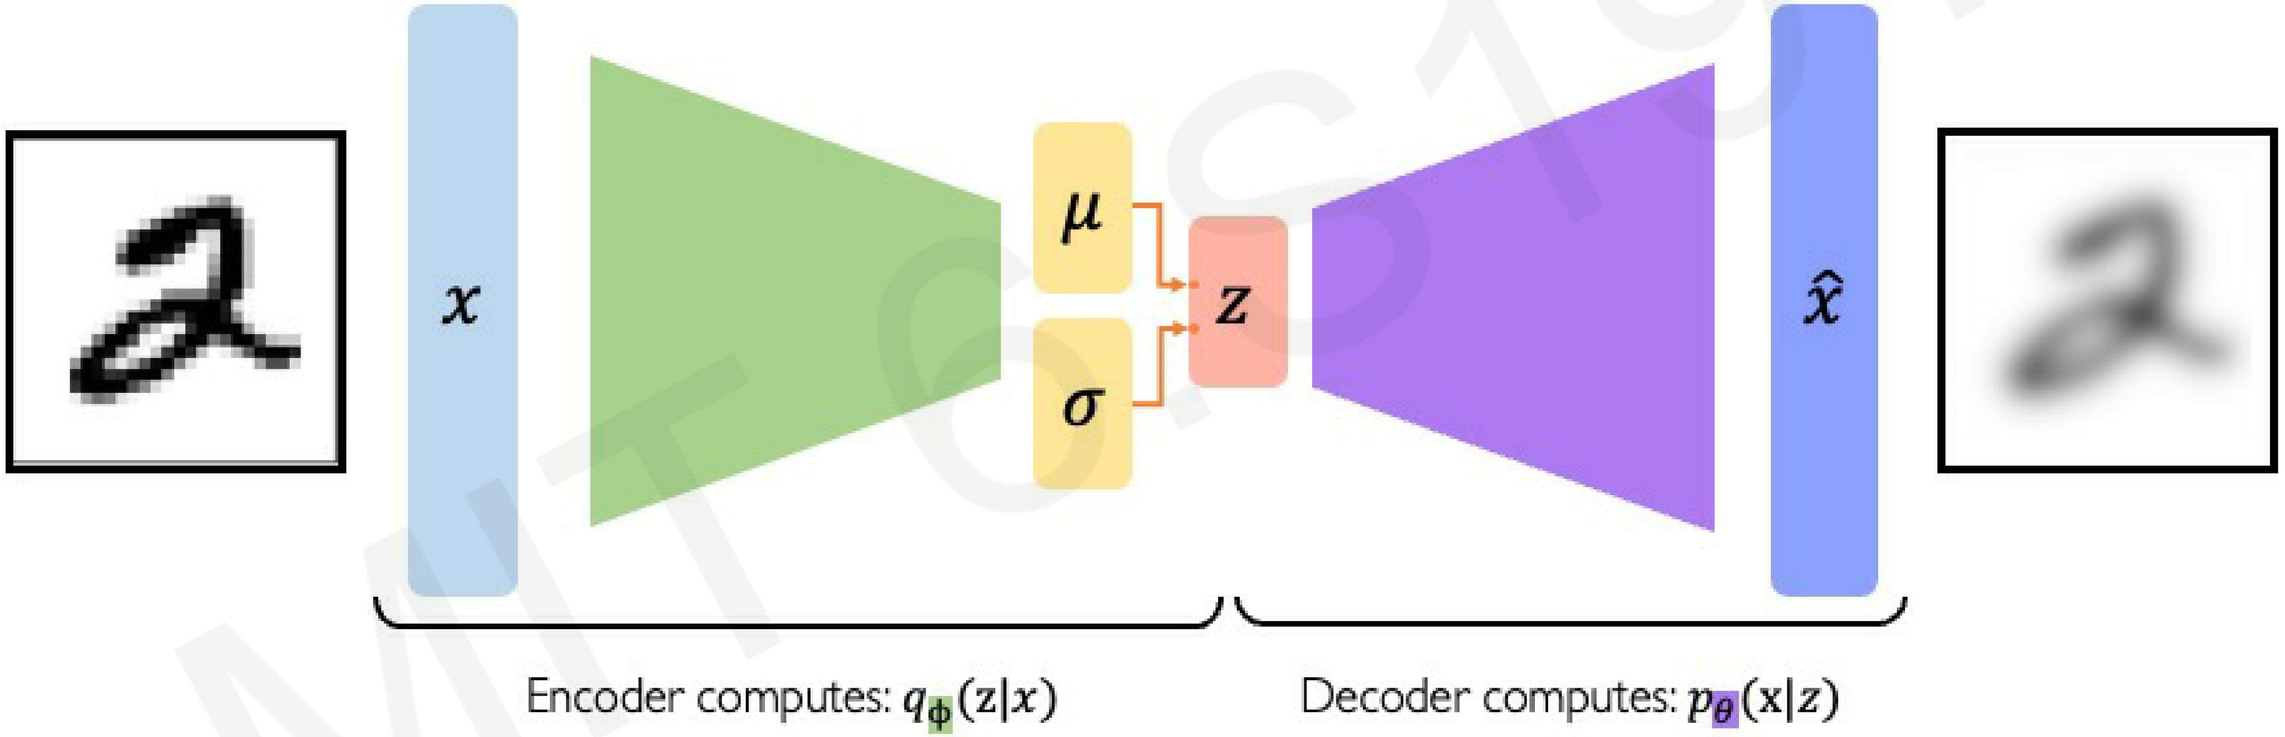

$$\mathcal{L}(\phi, \theta, x) = \text{(reconstruction loss)} + \text{(regularization term)}$$

MIT 6.S191

VAE optimization

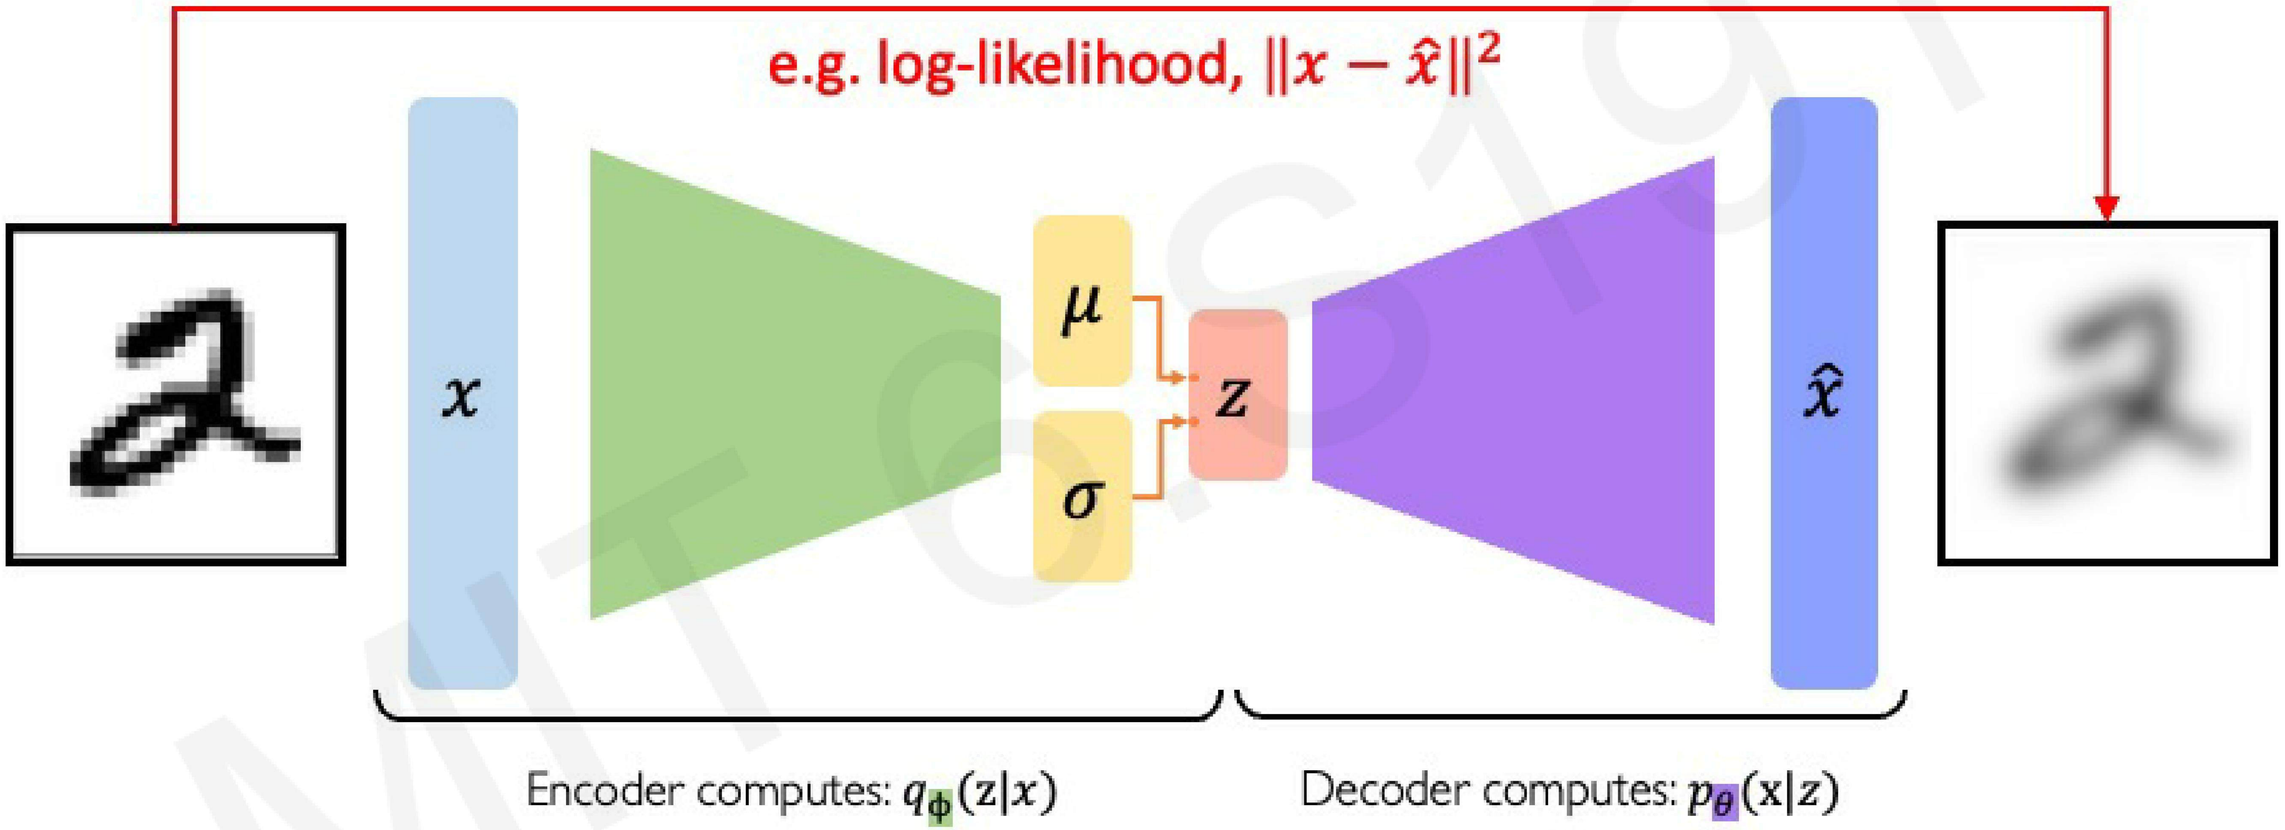

$$\mathcal{L}(\phi, \theta, x) = \boxed{\text{(reconstruction loss)}} + \text{(regularization term)}$$

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

VAE optimization

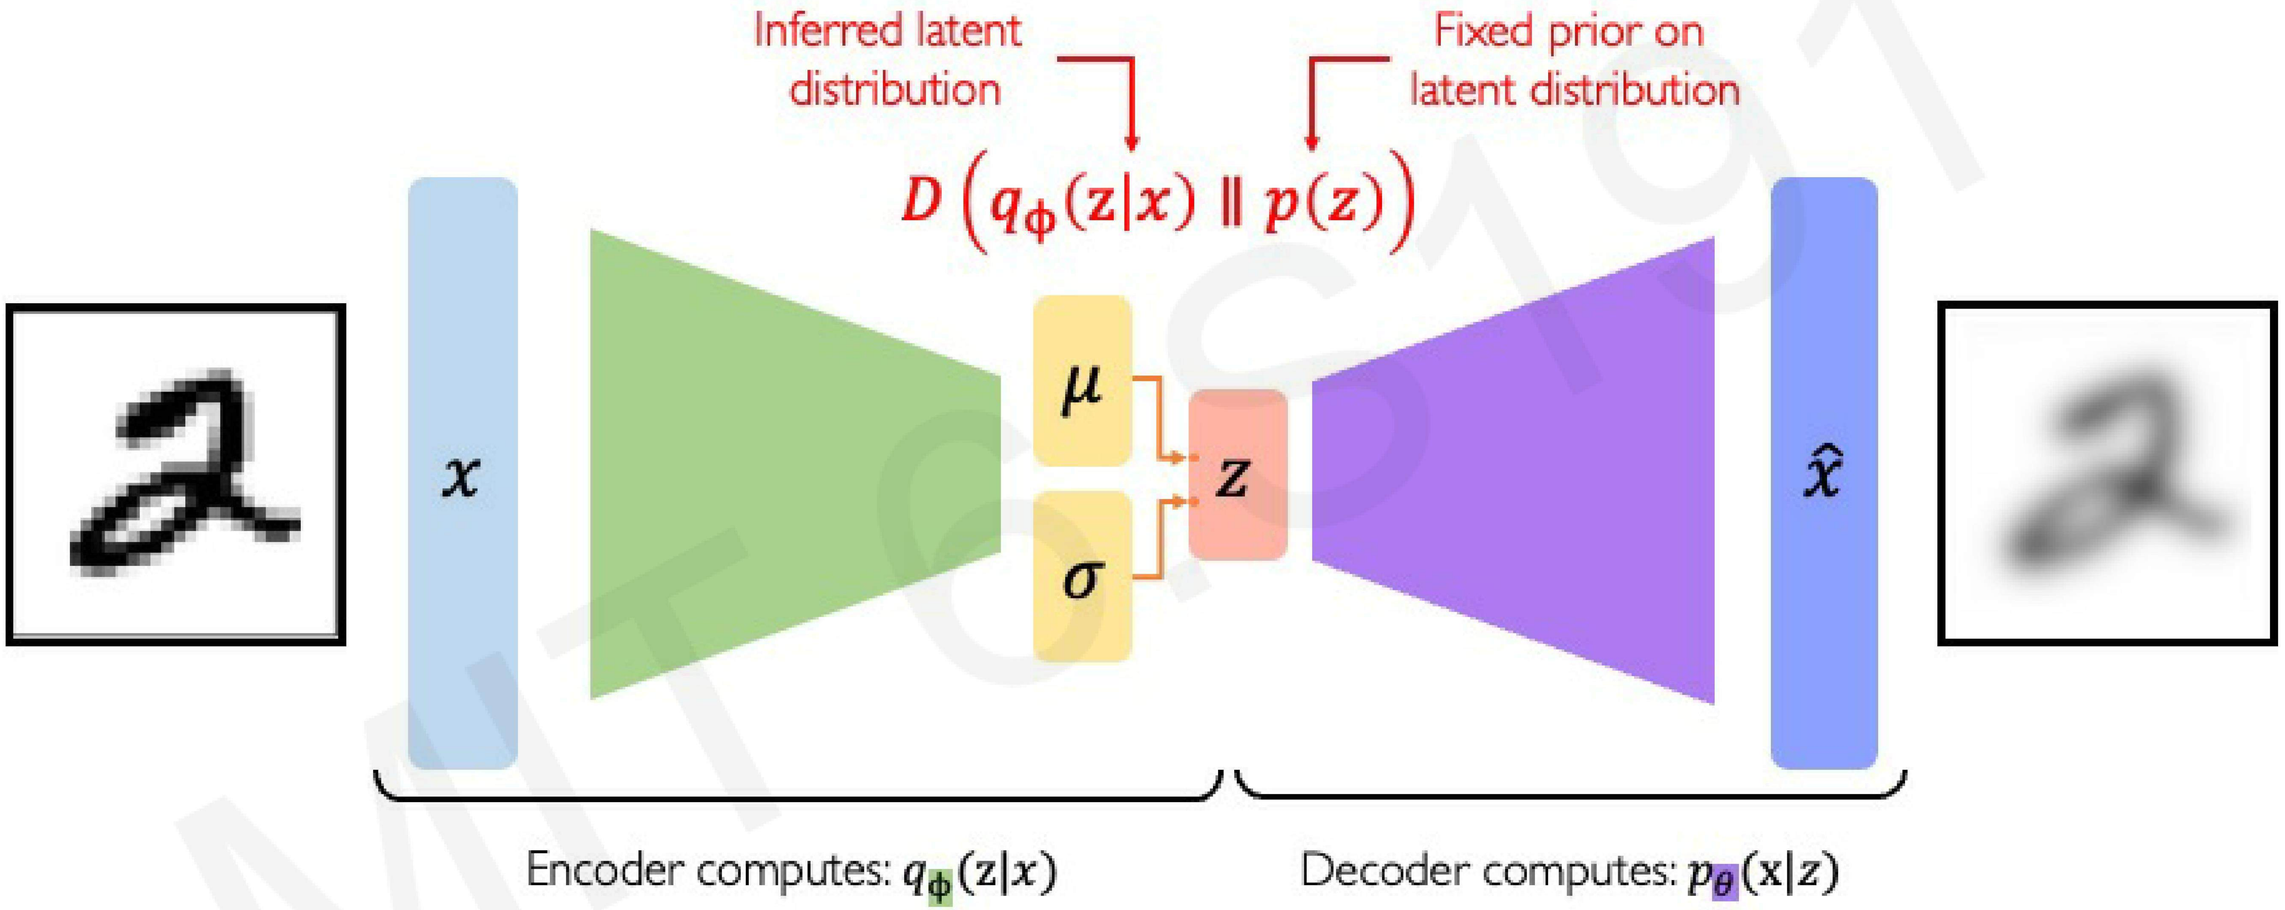

$$\mathcal{L}(\phi, \theta, x) = \text{(reconstruction loss)} + \boxed{\text{(regularization term)}}$$

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Priors on the latent distribution

$$D\left(q_\phi(\mathbf{z}|x) \parallel p(z)\right)$$

Inferred latent distribution

Fixed prior on latent distribution

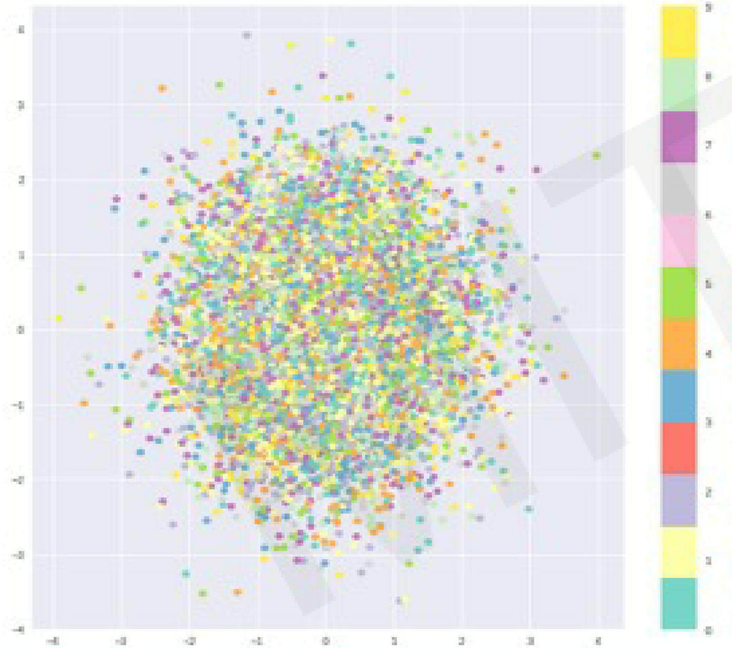

Common choice of prior – Normal Gaussian:

$$p(z) = \mathcal{N}(\mu = 0, \sigma^2 = 1)$$

- Encourages encodings to distribute encodings evenly around the center of the latent space
- Penalize the network when it tries to "cheat" by clustering points in specific regions (i.e., by memorizing the data)

6.S191

Priors on the latent distribution

$$D\left(q_\phi(\mathrm{z}|x) \parallel p(z)\right)$$

$$= -\frac{1}{2}\sum_{j=0}^{k-1}\left(\sigma_j + \mu_j^2 - 1 - \log\sigma_j\right)$$

KL-divergence between the two distributions

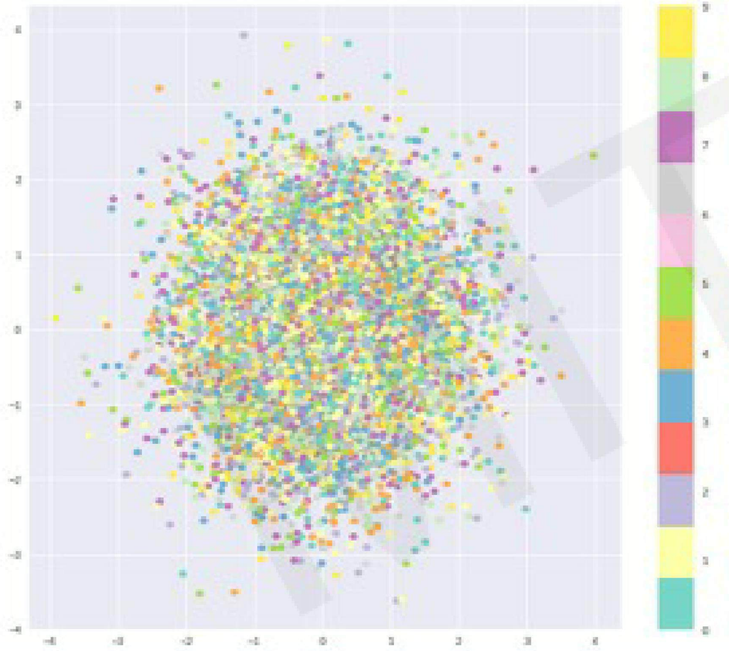

Common choice of prior – Normal Gaussian:

$$p(z) = \mathcal{N}(\mu = 0, \sigma^2 = 1)$$

- Encourages encodings to distribute encodings evenly around the center of the latent space
- Penalize the network when it tries to "cheat" by clustering points in specific regions (i.e., by memorizing the data)

6.S191

Intuition on regularization and the Normal prior

What properties do we want to achieve from regularization? 🤔

1. **Continuity:** points that are close in latent space → similar content after decoding
2. **Completeness:** sampling from latent space → "meaningful" content after decoding

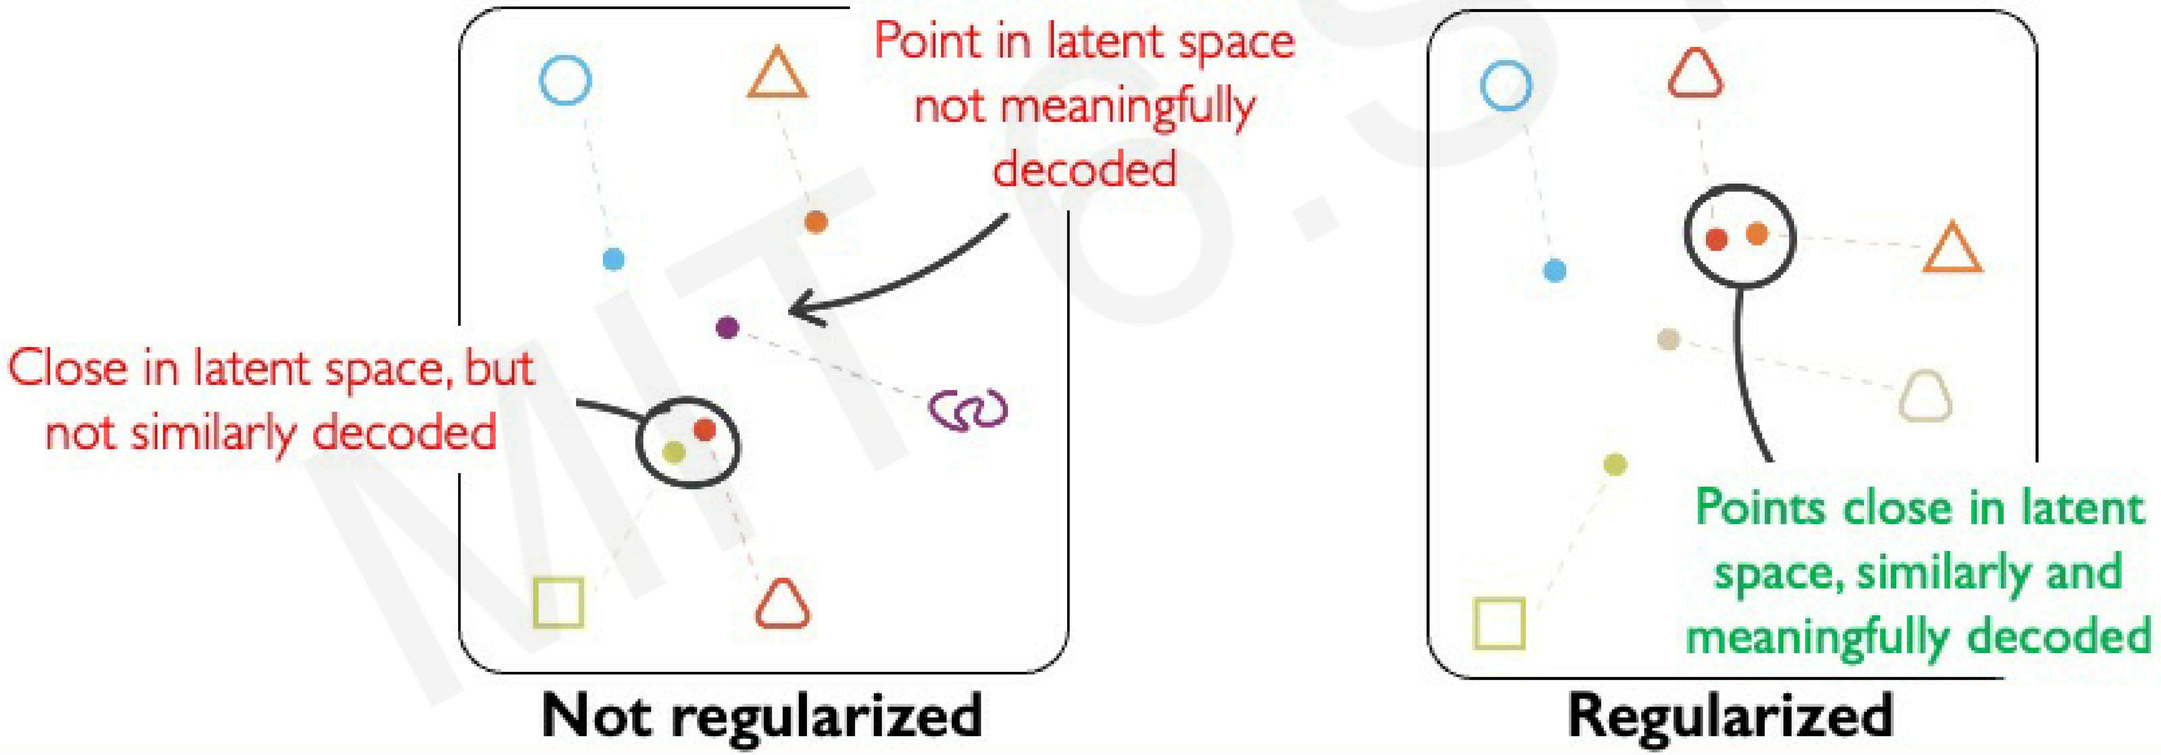

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Intuition on regularization and the Normal prior

1. **Continuity**: points that are close in latent space → similar content after decoding
2. **Completeness:** sampling from latent space → "meaningful" content after decoding

Encoding as a distribution does not guarantee these properties!

Normal prior →
continuity + completeness

Small variances →
Pointed distributions

Different means →
Discontinuities

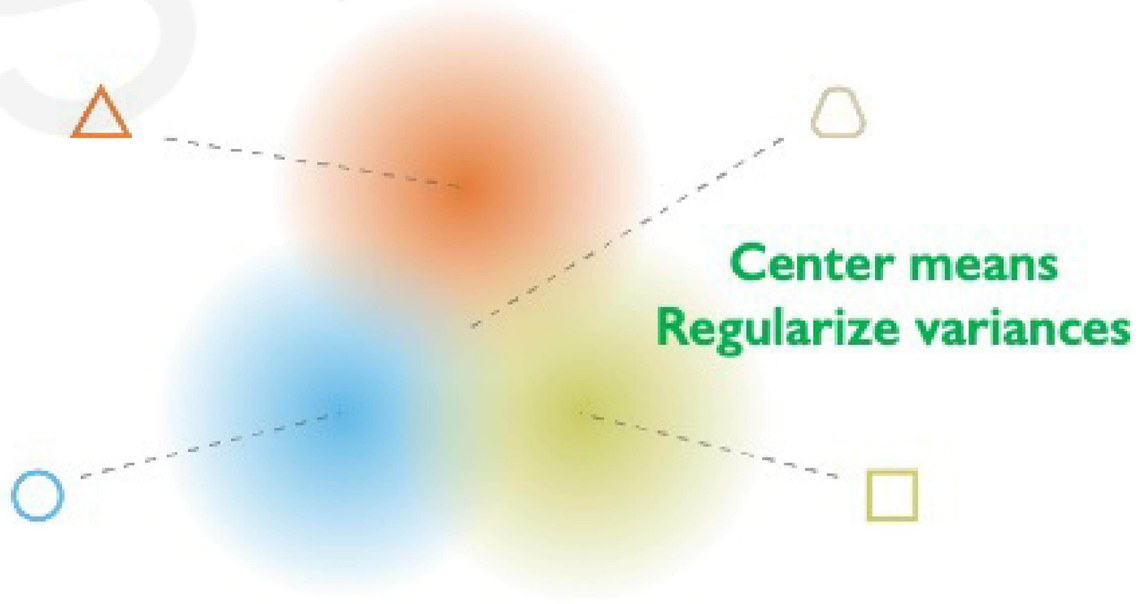

Not regularized

Regularized

MIT 6.S191

Intuition on regularization and the Normal prior

1. **Continuity**: points that are close in latent space → similar content after decoding
2. **Completeness:** sampling from latent space → "meaningful" content after decoding

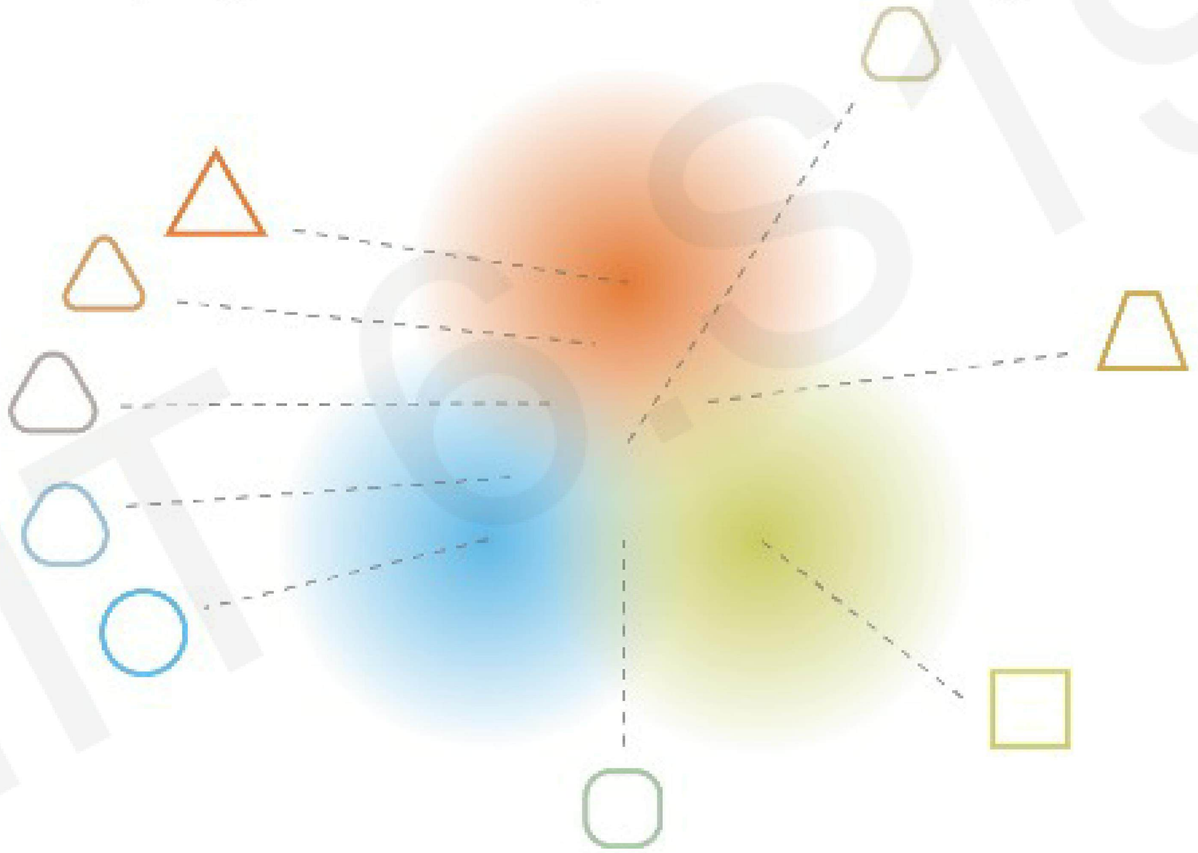

Regularization with Normal prior helps enforce **continuity & completeness** in the latent space.

MIT 6.S191

VAE computation graph

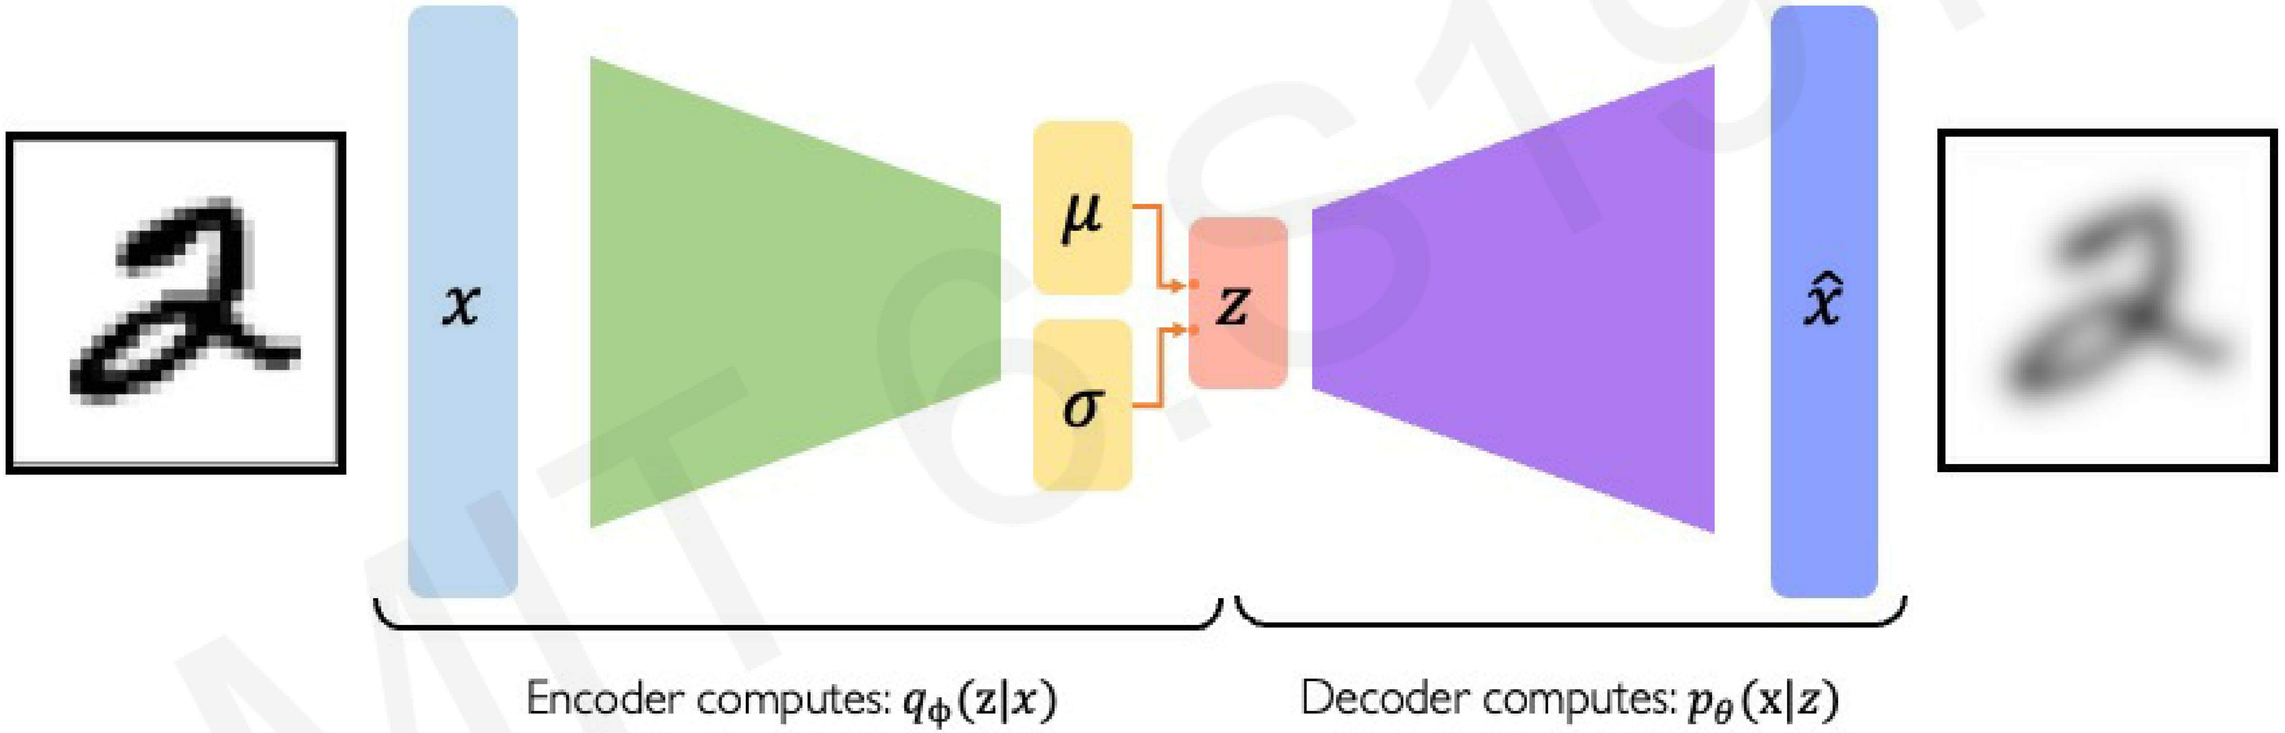

$$\mathcal{L}(\phi, \theta, x) = \text{(reconstruction loss)} + \text{(regularization term)}$$

MIT 6.S191

VAE computation graph

Problem: We cannot backpropagate gradients through sampling layers!

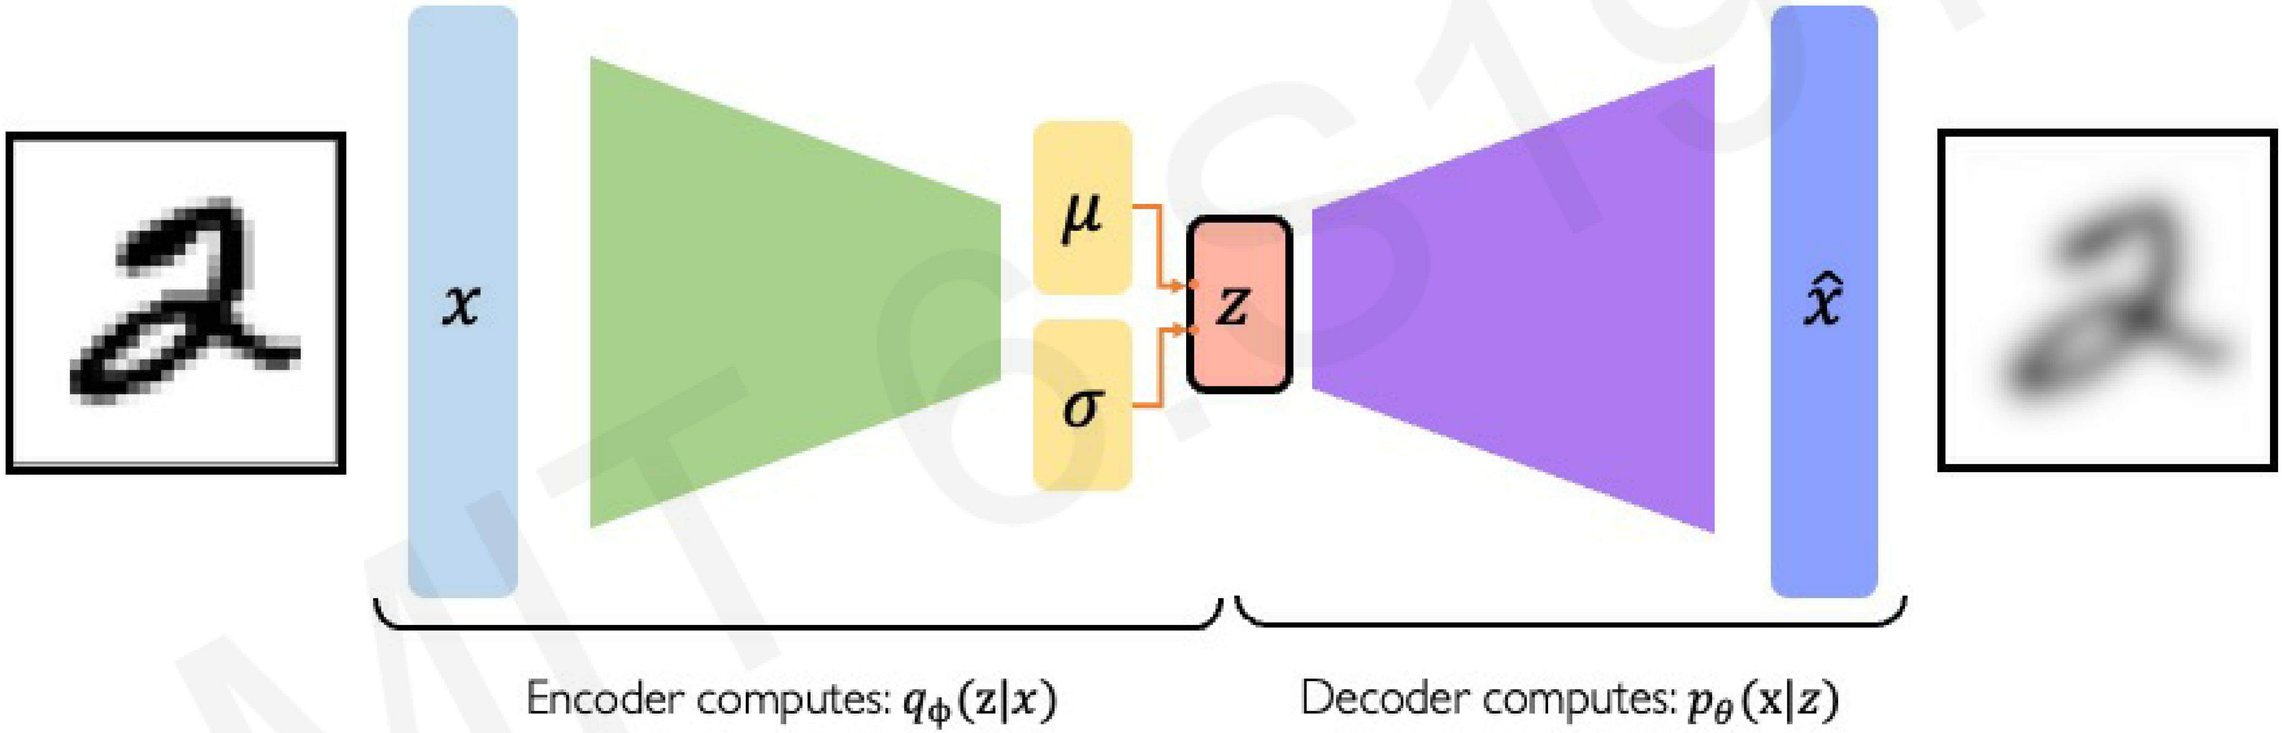

$$\mathcal{L}(\phi, \theta, x) = \text{(reconstruction loss)} + \text{(regularization term)}$$

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Reparametrizing the sampling layer

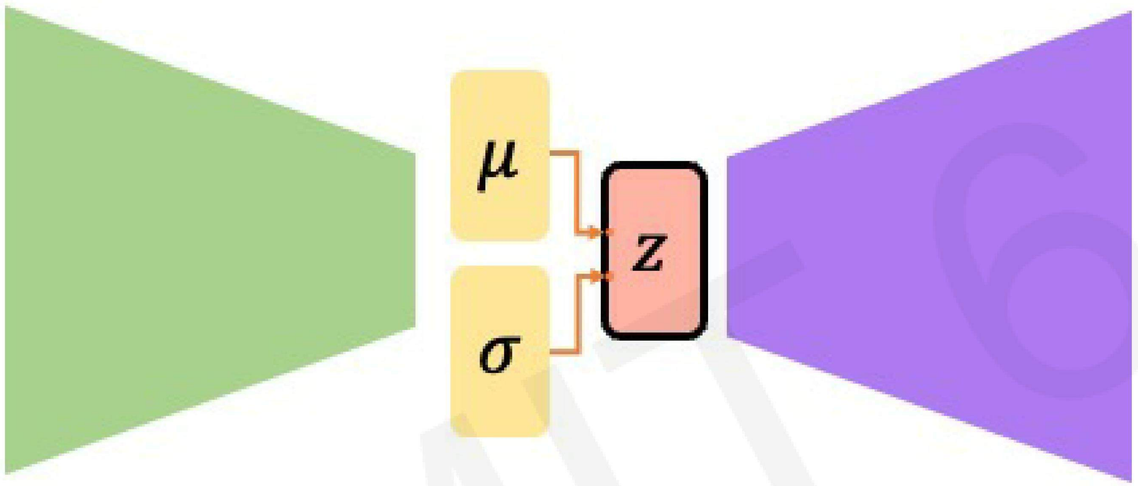

Key Idea:

~~$z \sim \mathcal{N}(\mu, \sigma^2)$~~

Consider the sampled latent vector z as a sum of

- a fixed μ vector,
- and fixed σ vector, scaled by random constants drawn from the prior distribution

$$\Rightarrow z = \mu + \sigma \odot \varepsilon$$

where $\varepsilon \sim \mathcal{N}(0,1)$

MIT6.S191

Reparametrizing the sampling layer

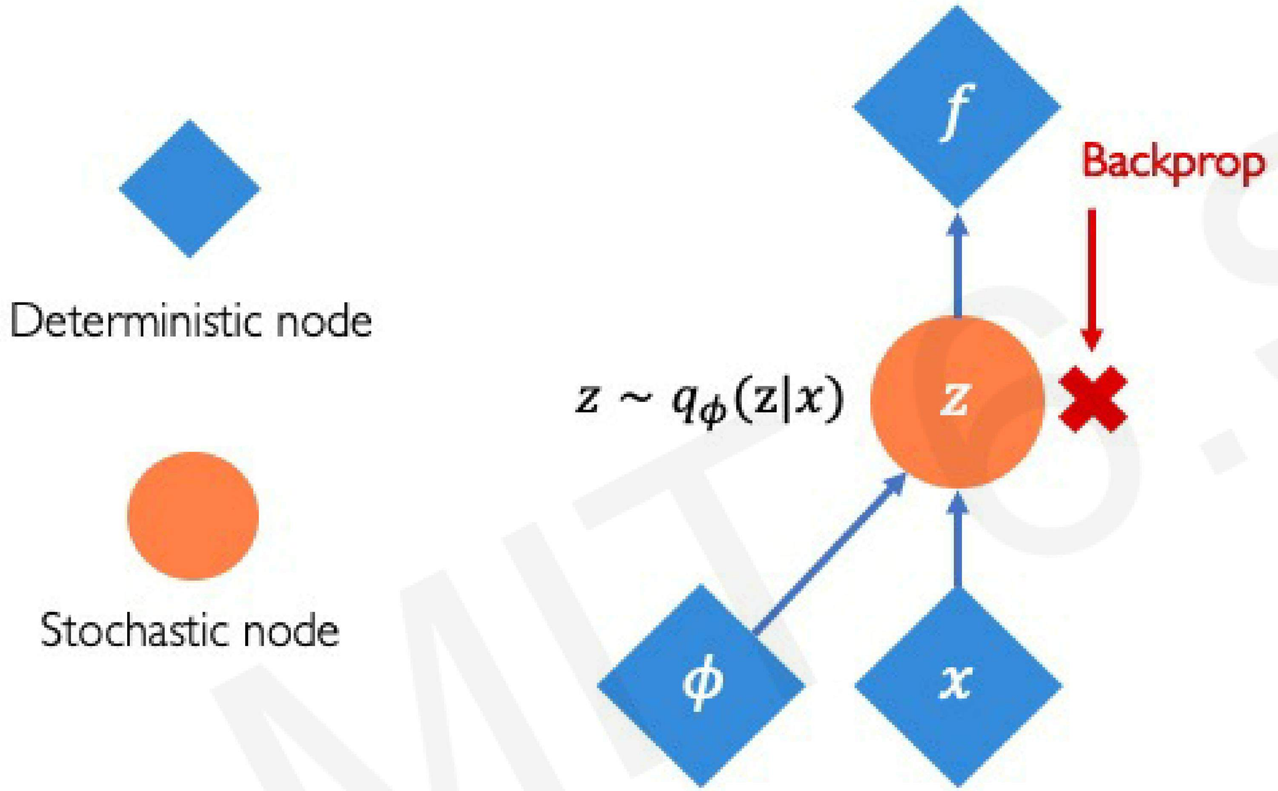

Original form

MIT 6.S191

Reparametrizing the sampling layer

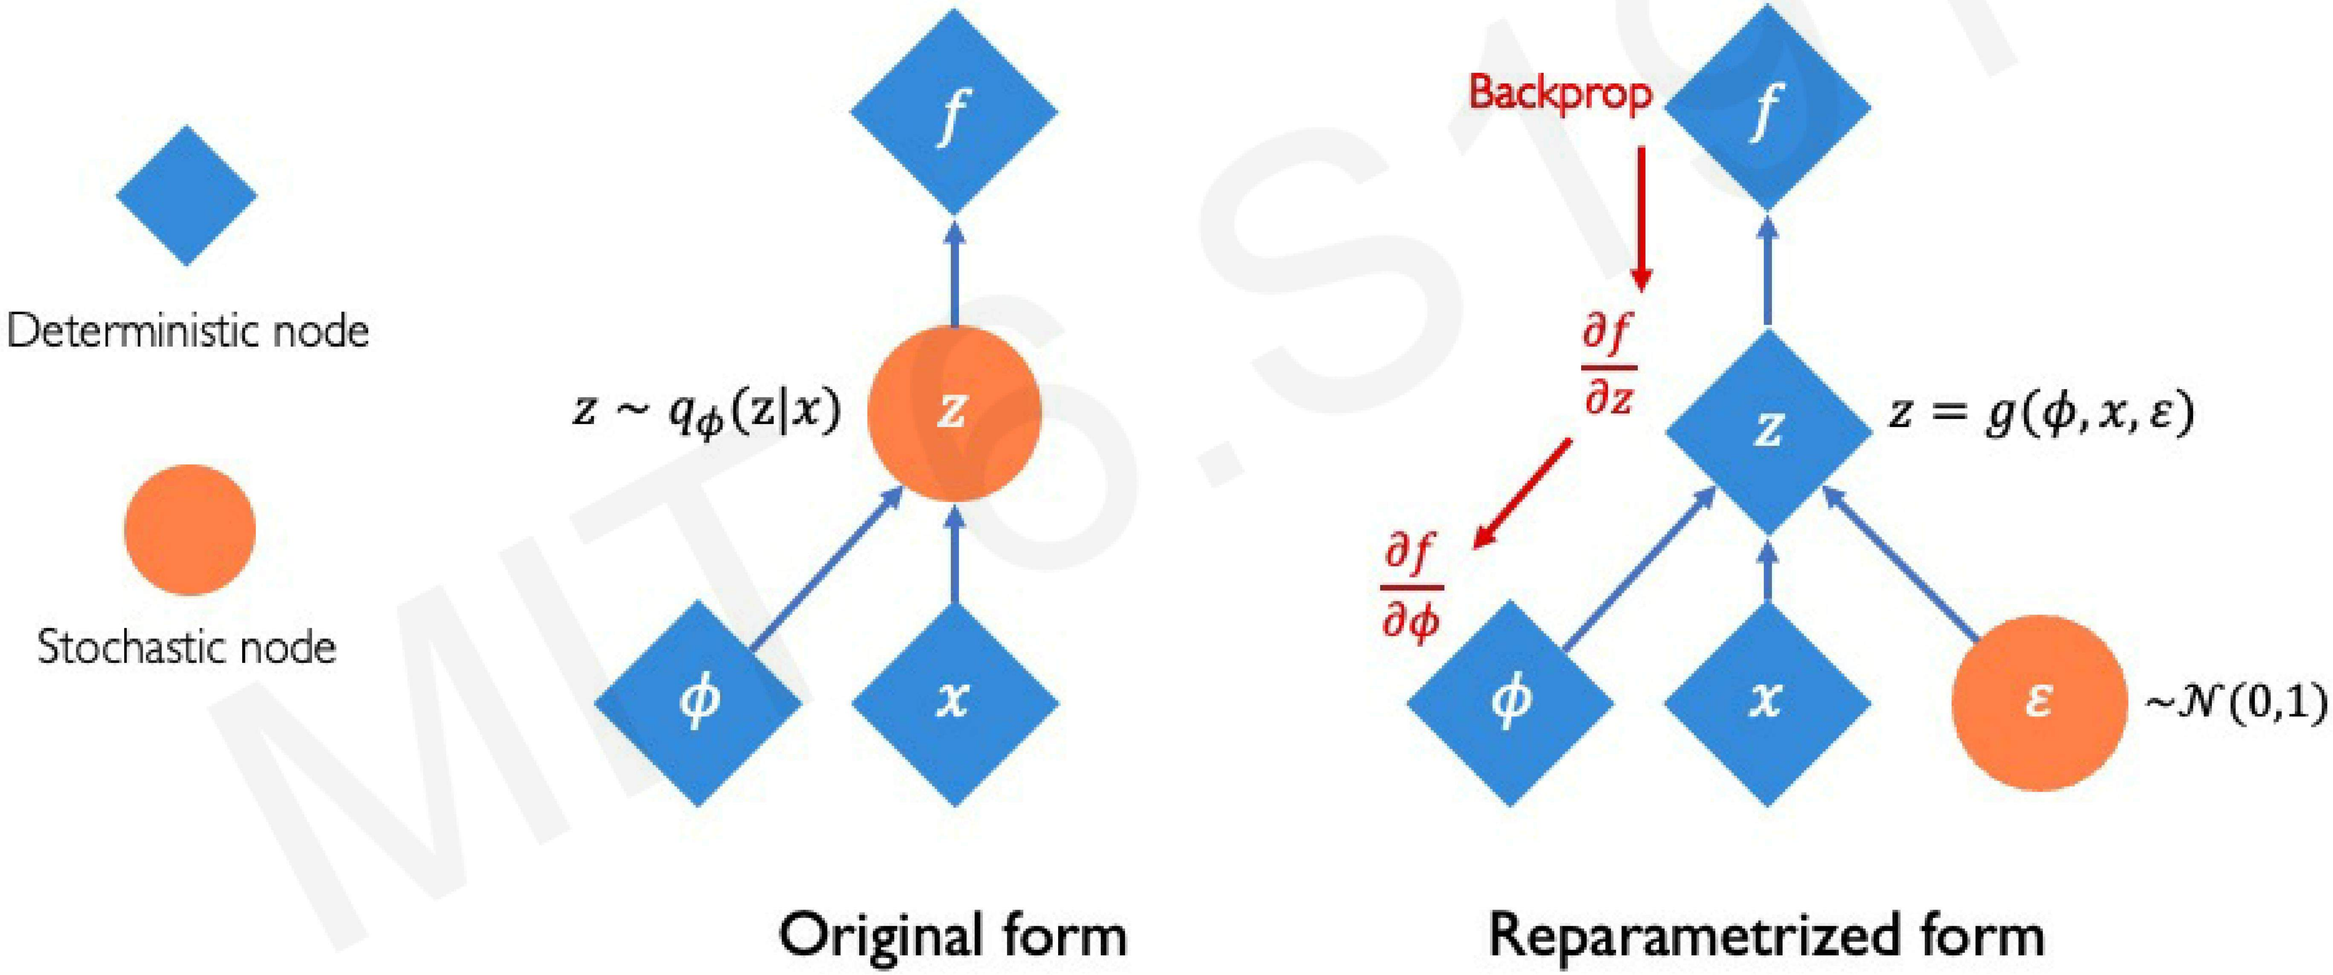

MIT 6.S191

VAEs: Latent perturbation

Slowly increase or decrease a **single latent variable**
Keep all other variables fixed

Head pose

Different dimensions of z encodes **different interpretable latent features**

MIT 6.S191

VAEs: Latent perturbation

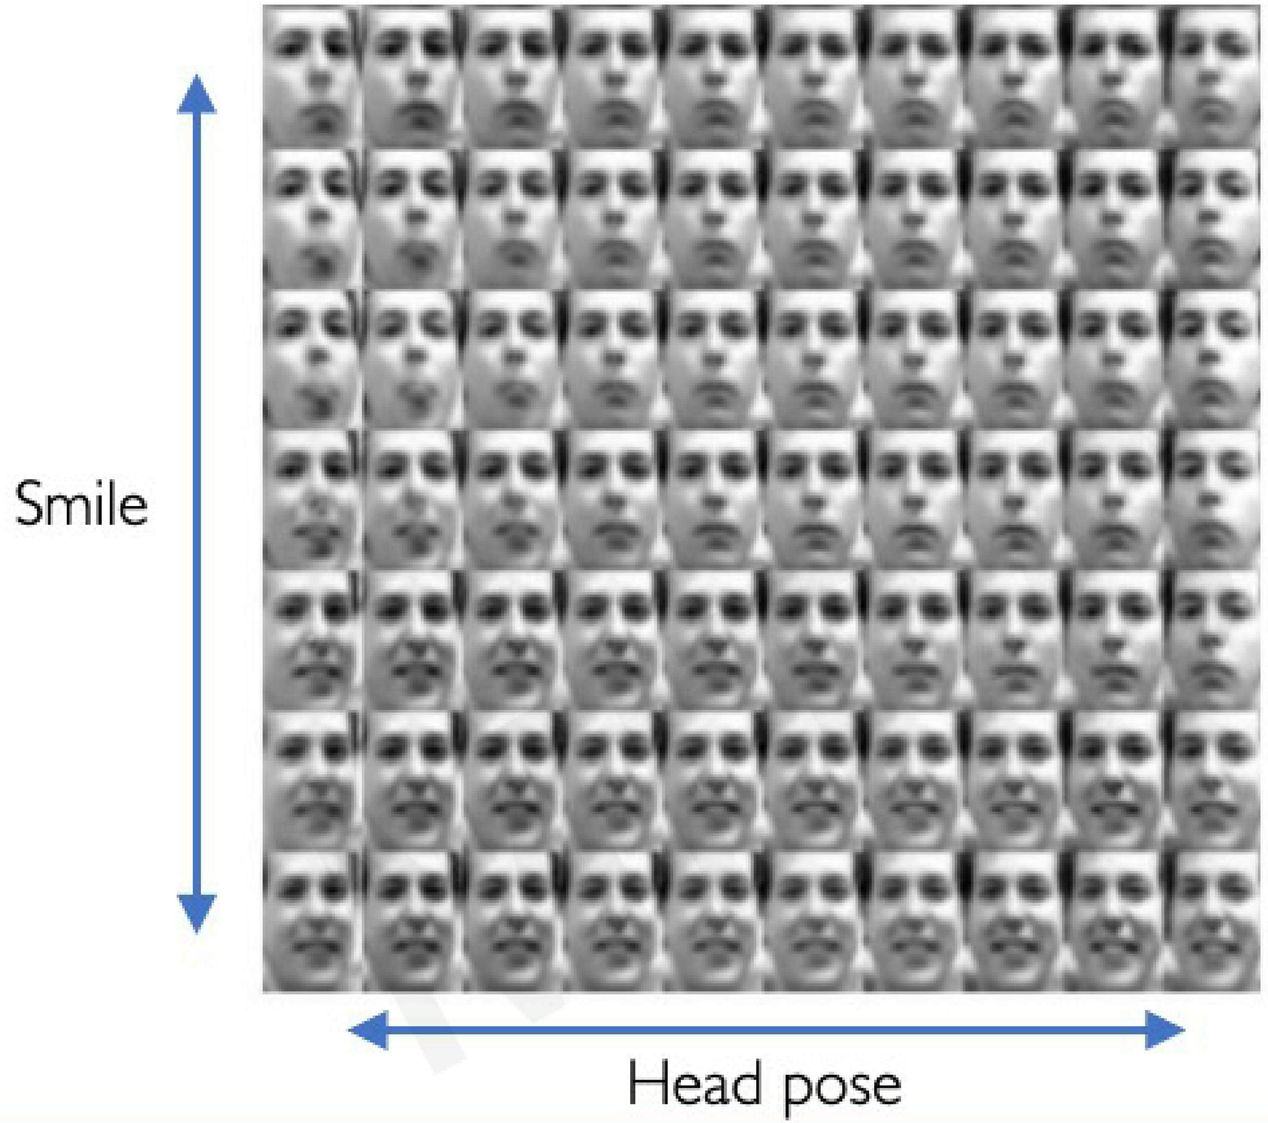

Ideally, we want latent variables that are uncorrelated with each other

Enforce diagonal prior on the latent variables to encourage independence

Disentanglement

Latent space disentanglement with β-VAEs

β-VAE loss:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{z}, \beta) = \mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \beta D_{KL}(q_\phi(\mathbf{z}|\mathbf{x}) \parallel p(\mathbf{z}))$$

Reconstruction term **Regularization term**

β > 1: constrain latent bottleneck, encourage efficient latent encoding → disentanglement

Head rotation (azimuth)

Smile also changing!

Standard VAE (β = 1)

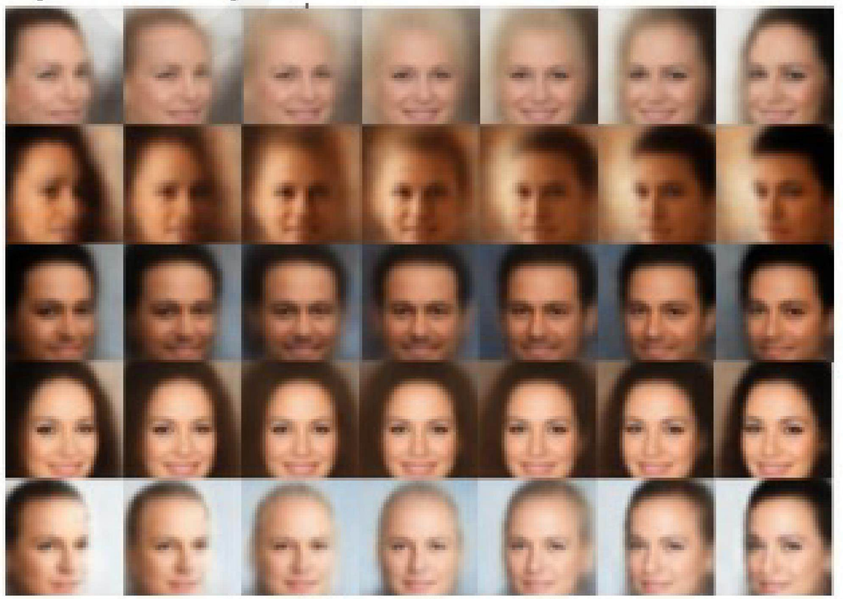

Smile relatively constant!

β-VAE (β = 250)

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Why latent variable models? Debiasing

Capable of uncovering **underlying latent variables** in a dataset

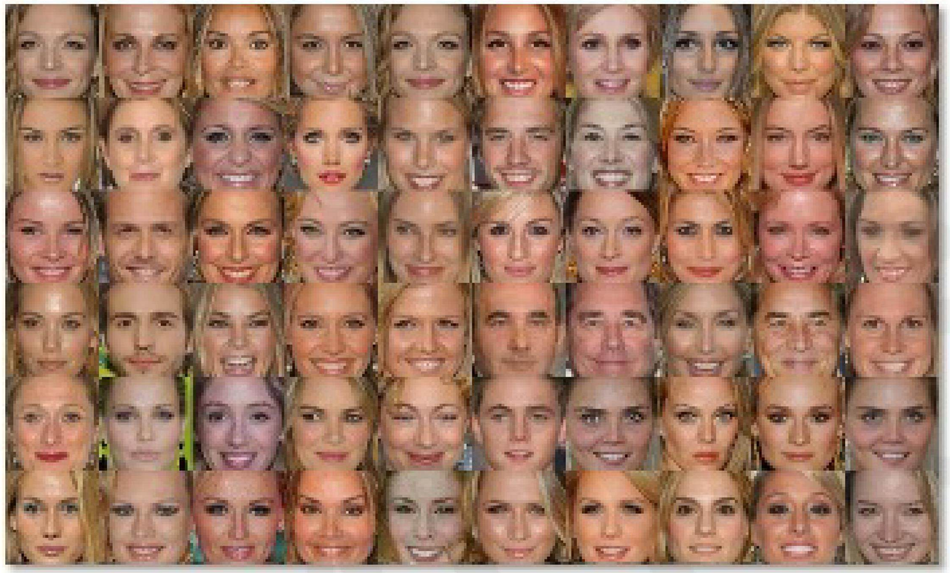

Homogeneous skin color, pose

VS

Diverse skin color, pose, illumination

How can we use latent distributions to create fair and representative datasets?

Software Lab!

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

VAE summary

1. Compress representation of world to something we can use to learn

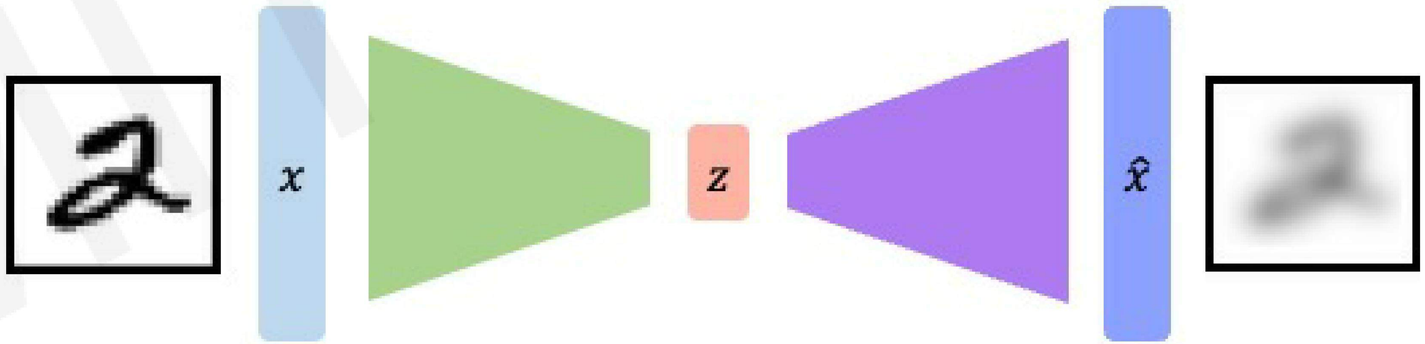

MIT 6.S191

VAE summary

1. Compress representation of world to something we can use to learn
2. Reconstruction allows for unsupervised learning (no labels!)

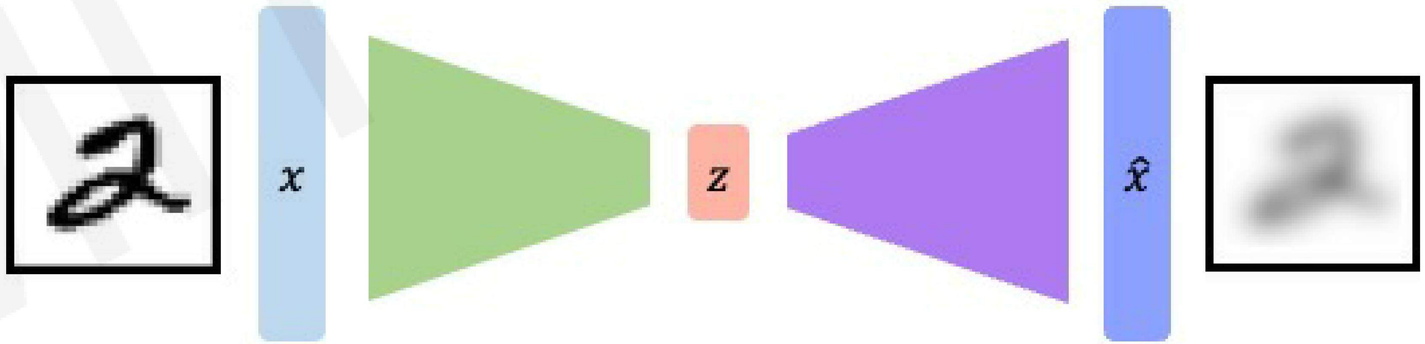

MIT 6.S191

VAE summary

1. Compress representation of world to something we can use to learn
2. Reconstruction allows for unsupervised learning (no labels!)
3. Reparameterization trick to train end-to-end

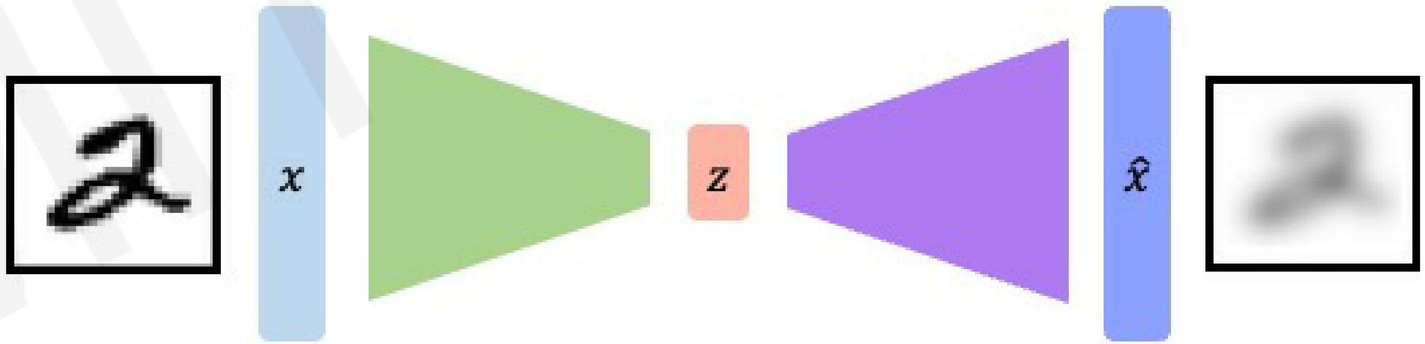

MIT 6.S191

VAE summary

1. Compress representation of world to something we can use to learn
2. Reconstruction allows for unsupervised learning (no labels!)
3. Reparameterization trick to train end-to-end
4. Interpret hidden latent variables using perturbation

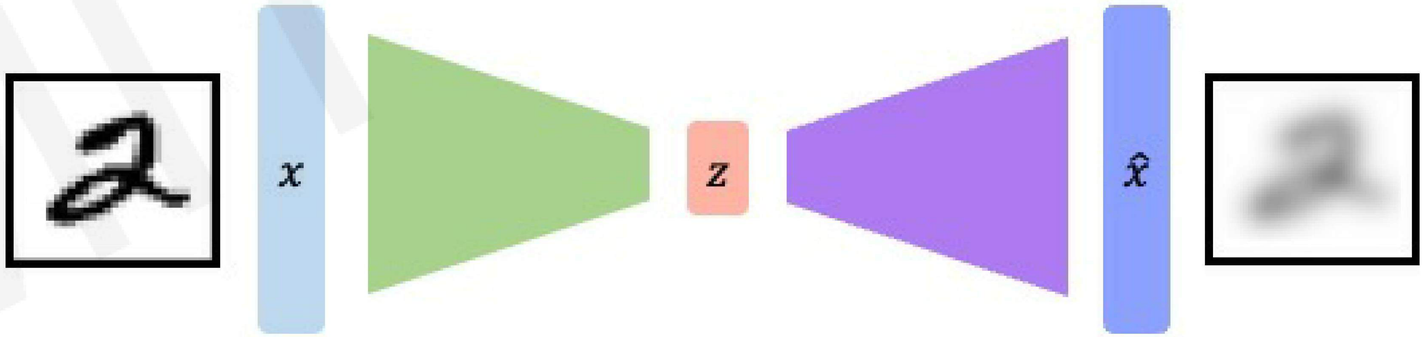

MIT 6.S191

VAE summary

1. Compress representation of world to something we can use to learn
2. Reconstruction allows for unsupervised learning (no labels!)
3. Reparameterization trick to train end-to-end
4. Interpret hidden latent variables using perturbation
5. Generating new examples

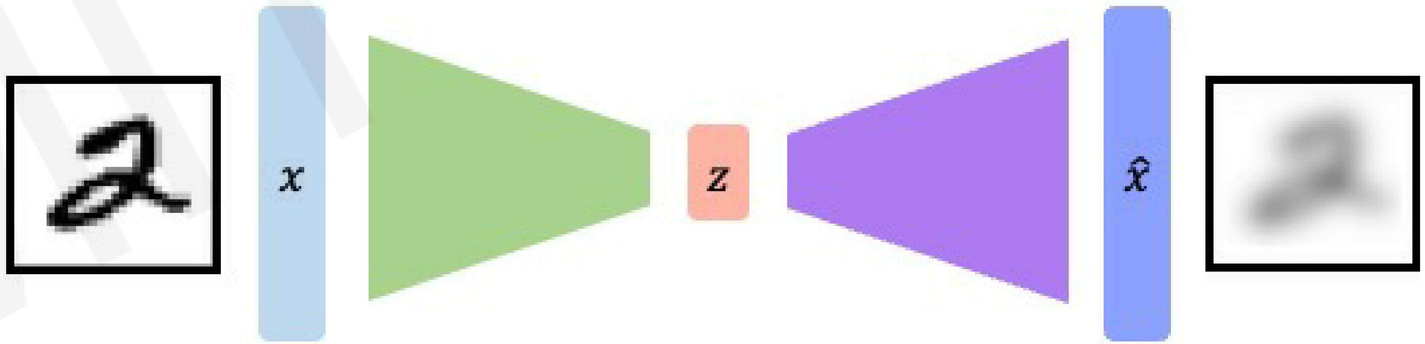

MIT 6.S191

Generative Adversarial Networks (GANs)

MIT 6.S191

What if we just want to sample?

Idea: don't explicitly model density, and instead just sample to generate new instances.

Problem: want to sample from complex distribution – can't do this directly!

Solution: sample from something simple (e.g., noise), learn a transformation to the data distribution.

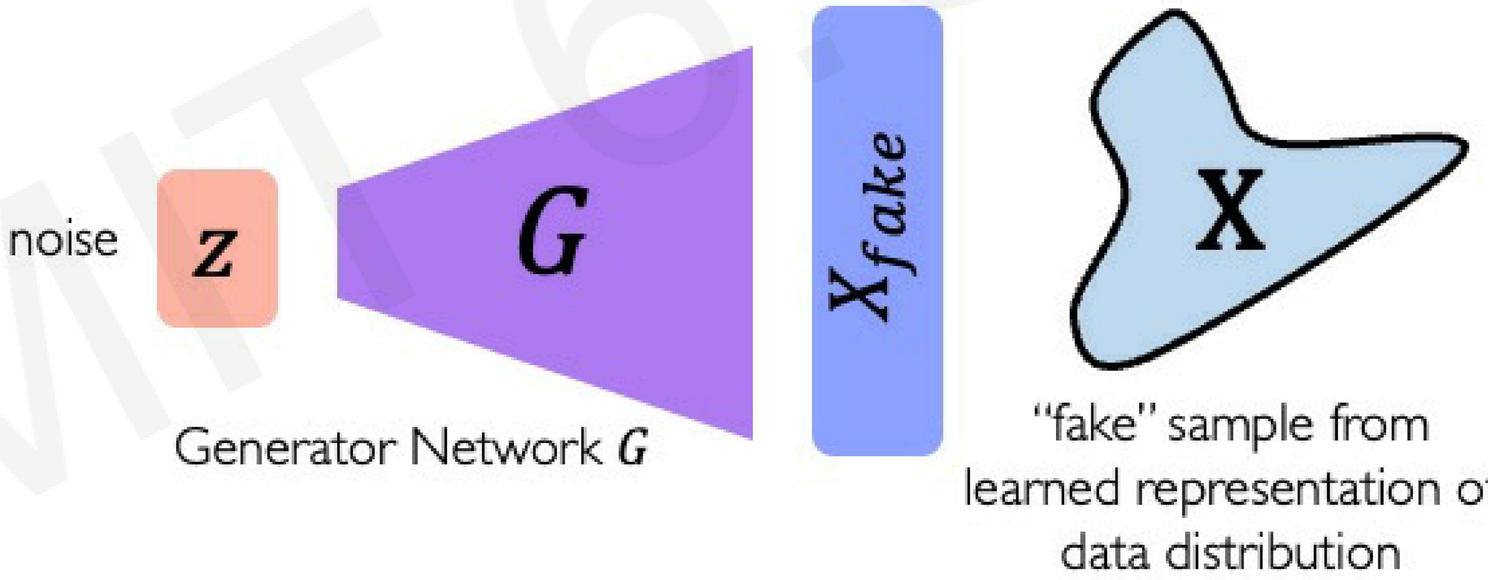

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Generative Adversarial Networks (GANs)

Generative Adversarial Networks (GANs) are a way to make a generative model by having two neural networks compete with each other.

The **discriminator** tries to identify real data from fakes created by the generator.

The **generator** turns noise into an imitation of the data to try to trick the discriminator.

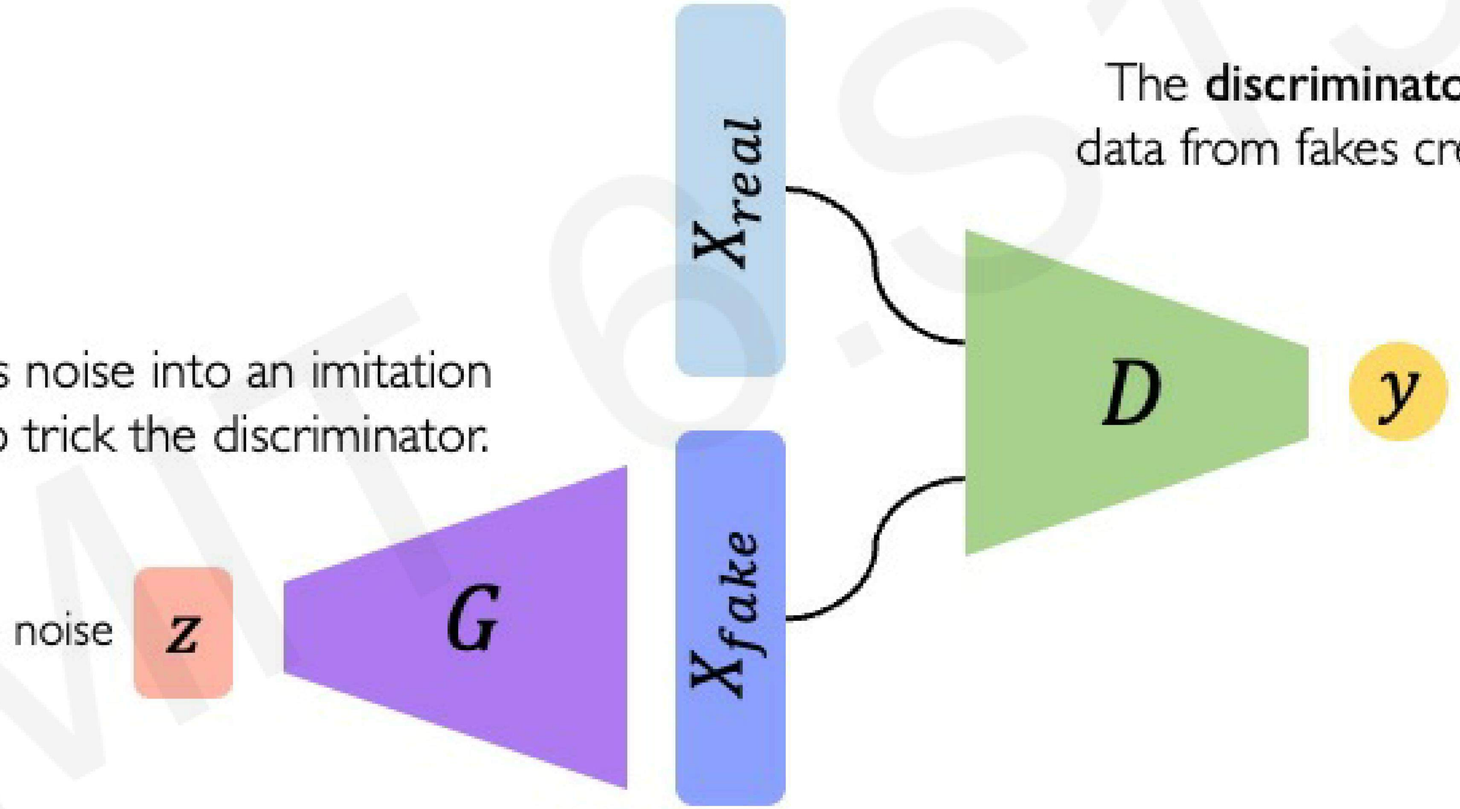

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Intuition behind GANs

Generator starts from noise to try to create an imitation of the data.

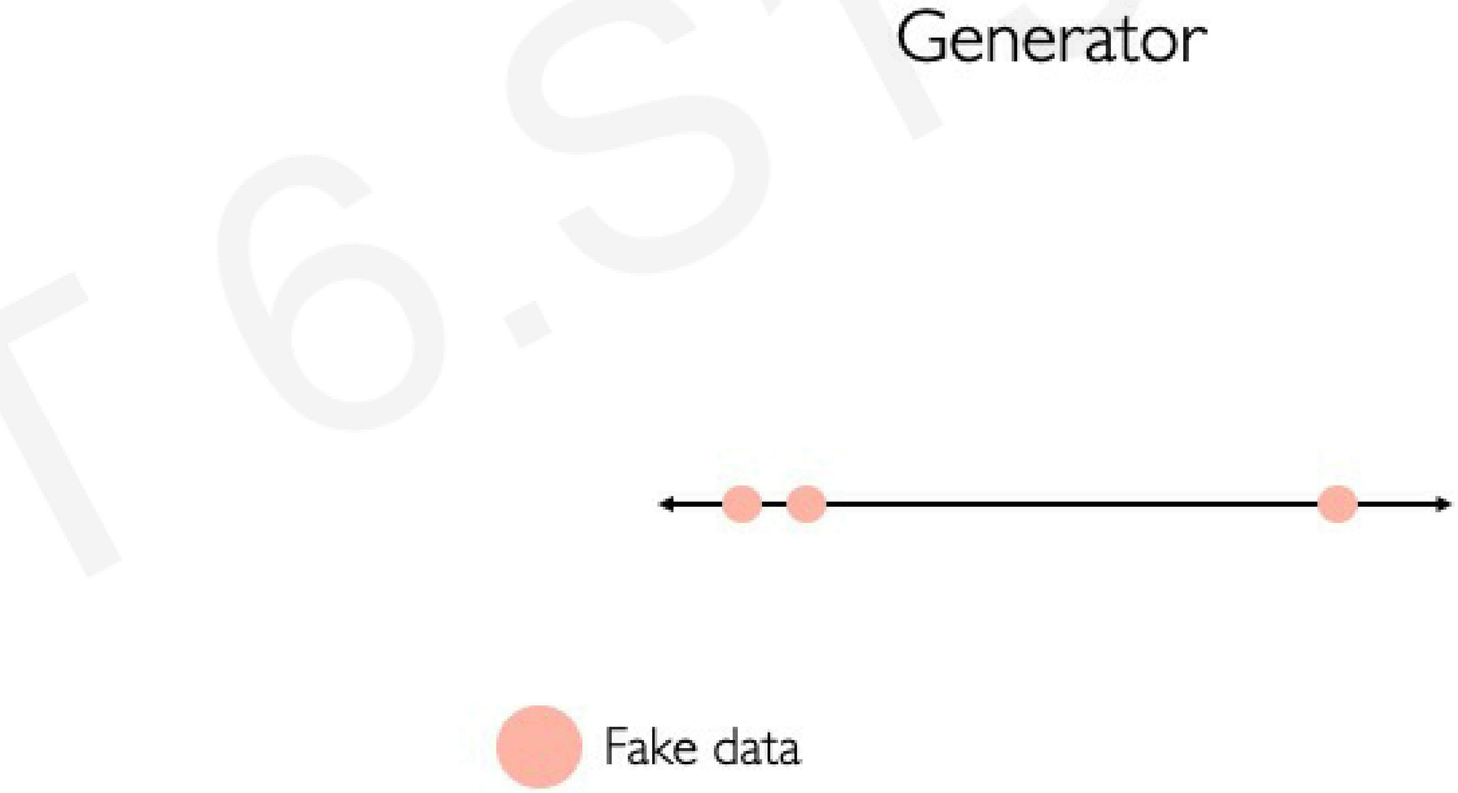

MIT 6.S191

Intuition behind GANs

Discriminator looks at both real data and fake data created by the generator.

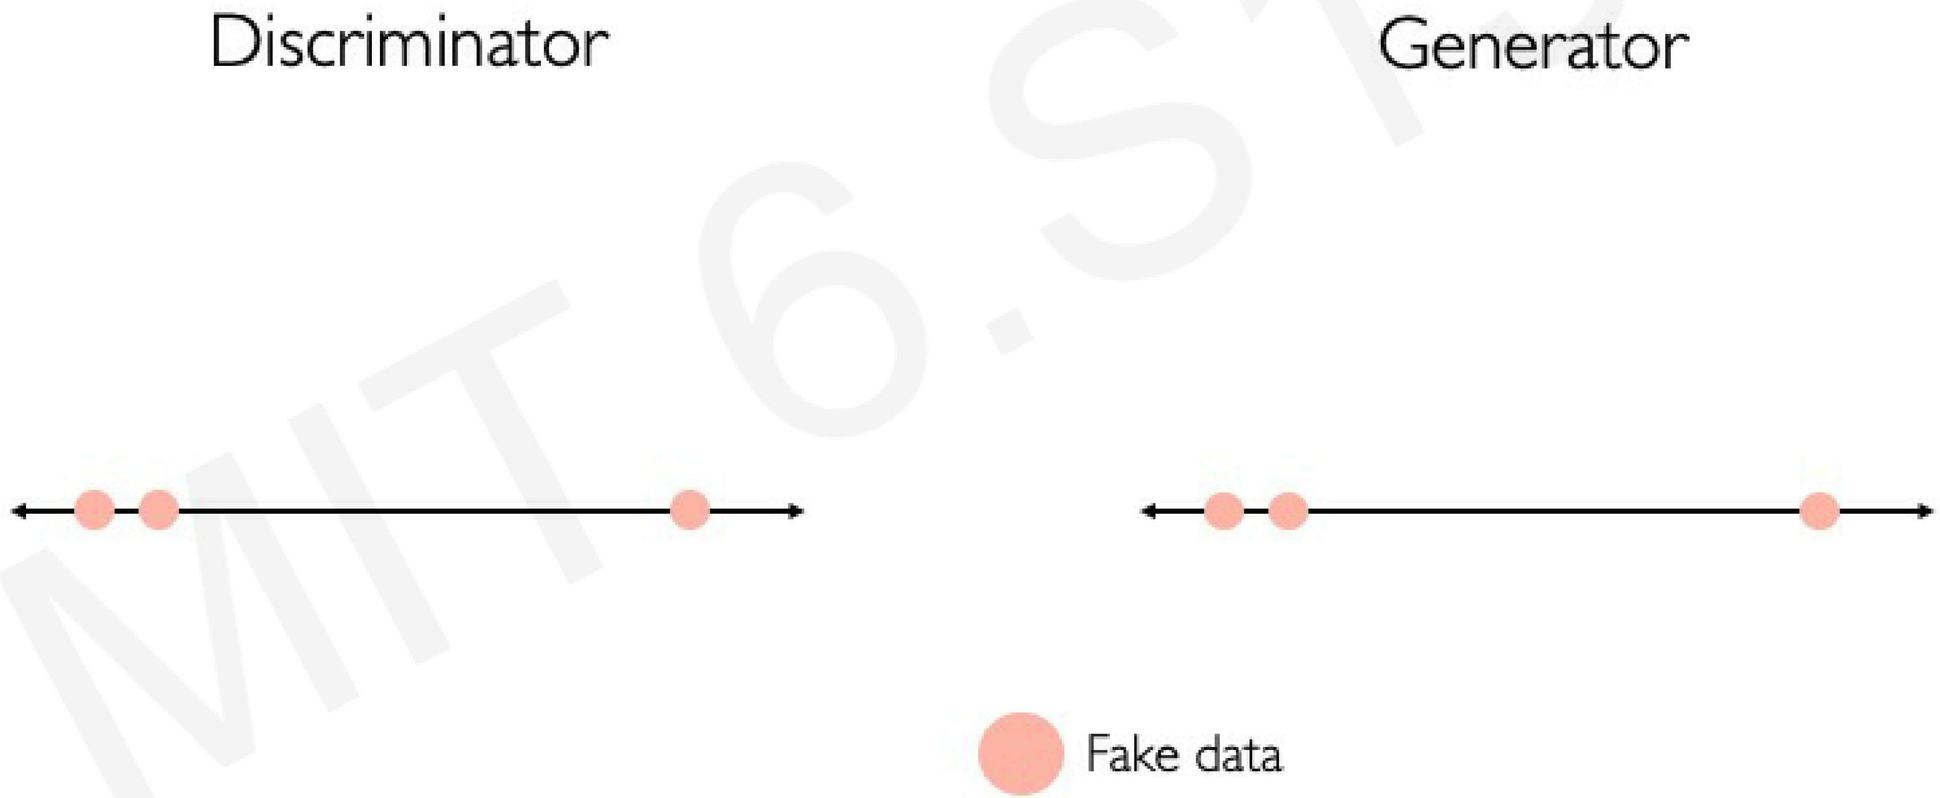

MIT 6.S191

Intuition behind GANs

Discriminator looks at both real data and fake data created by the generator.

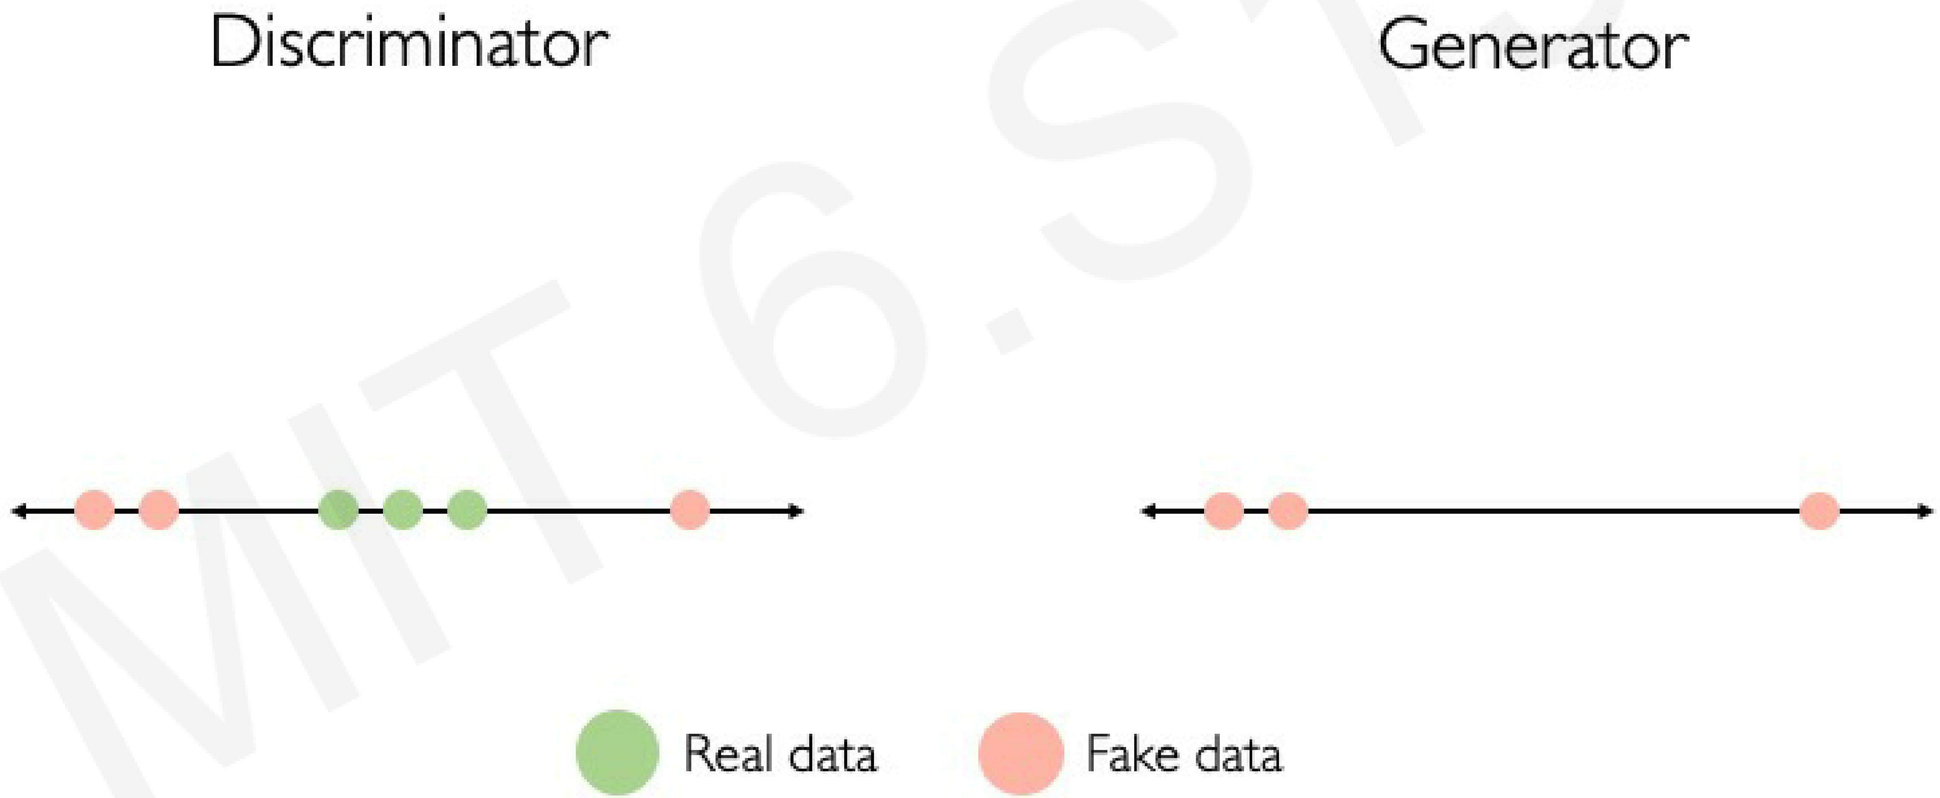

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

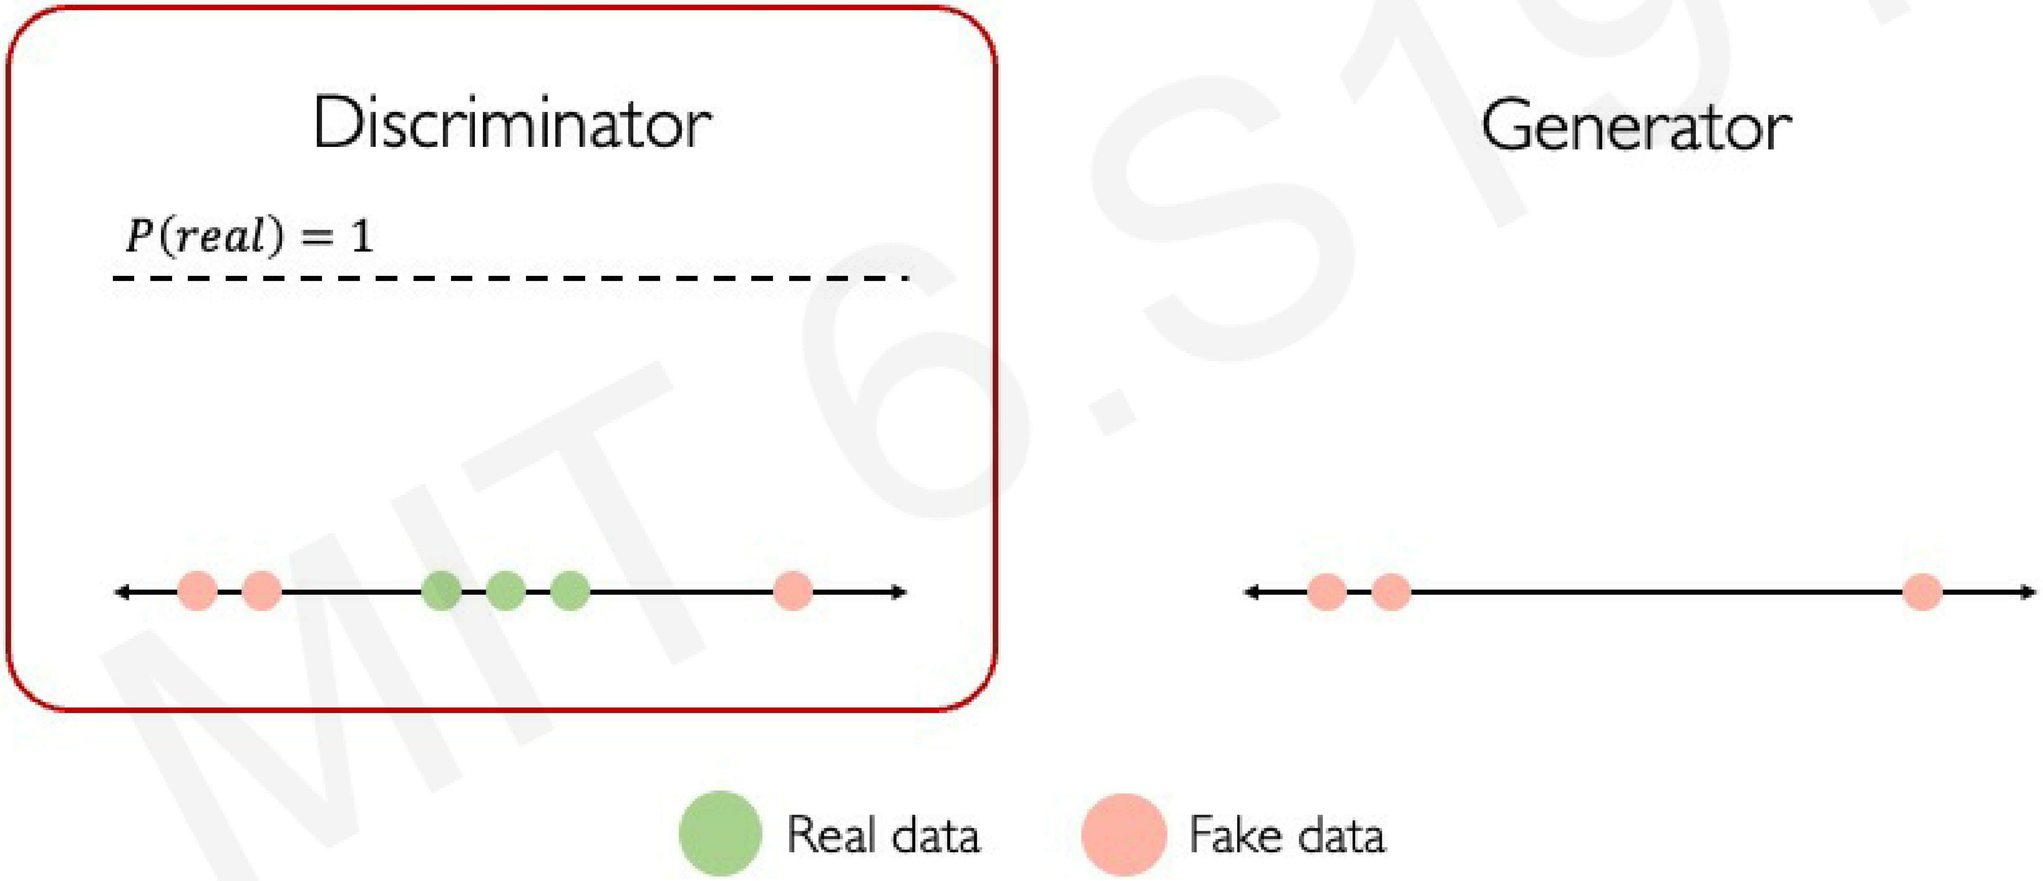

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

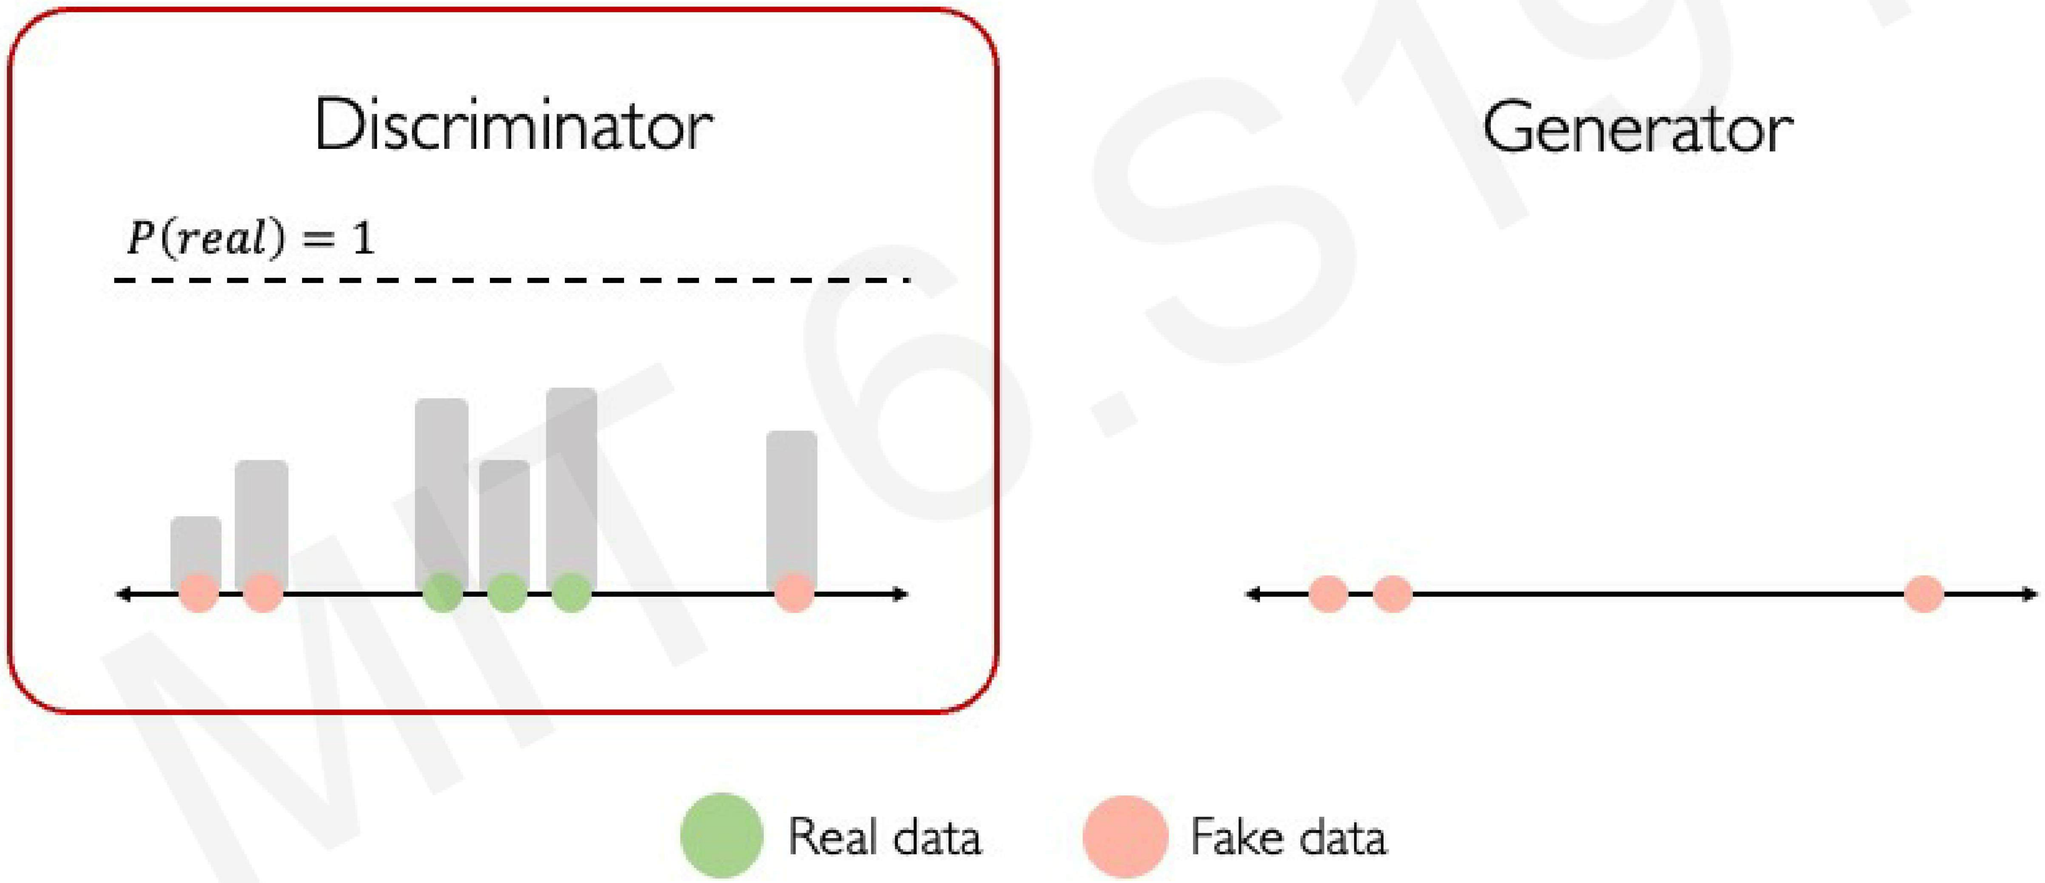

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

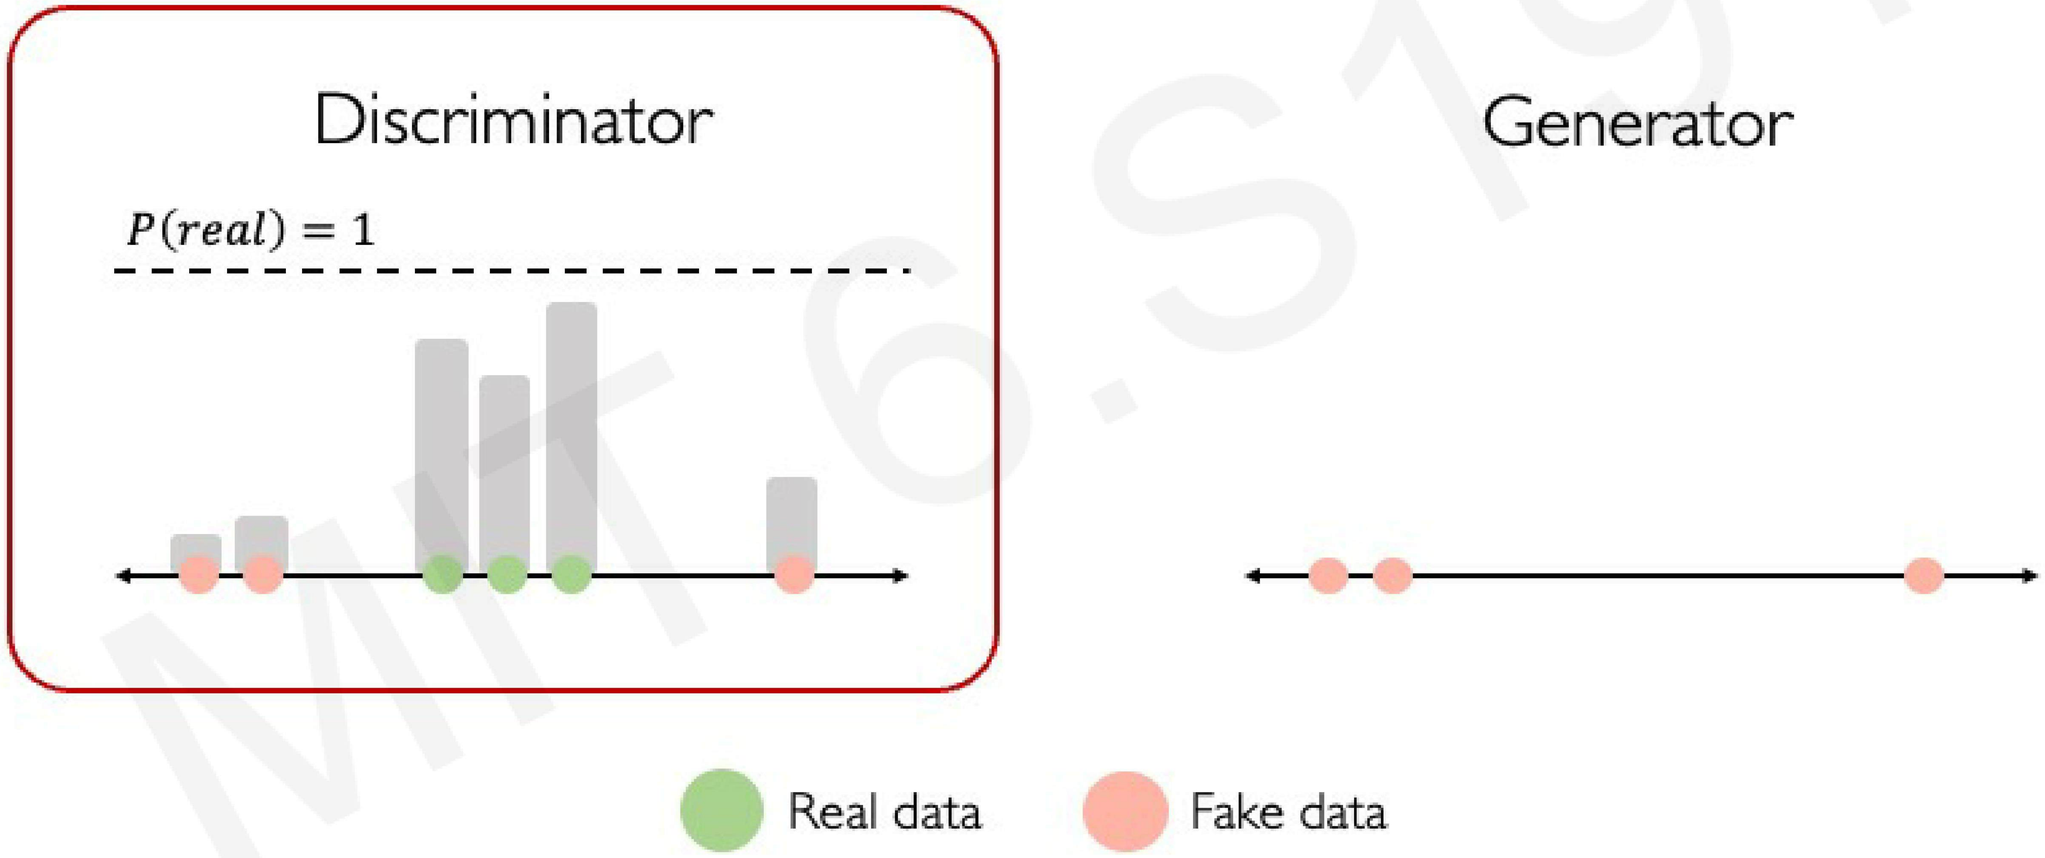

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

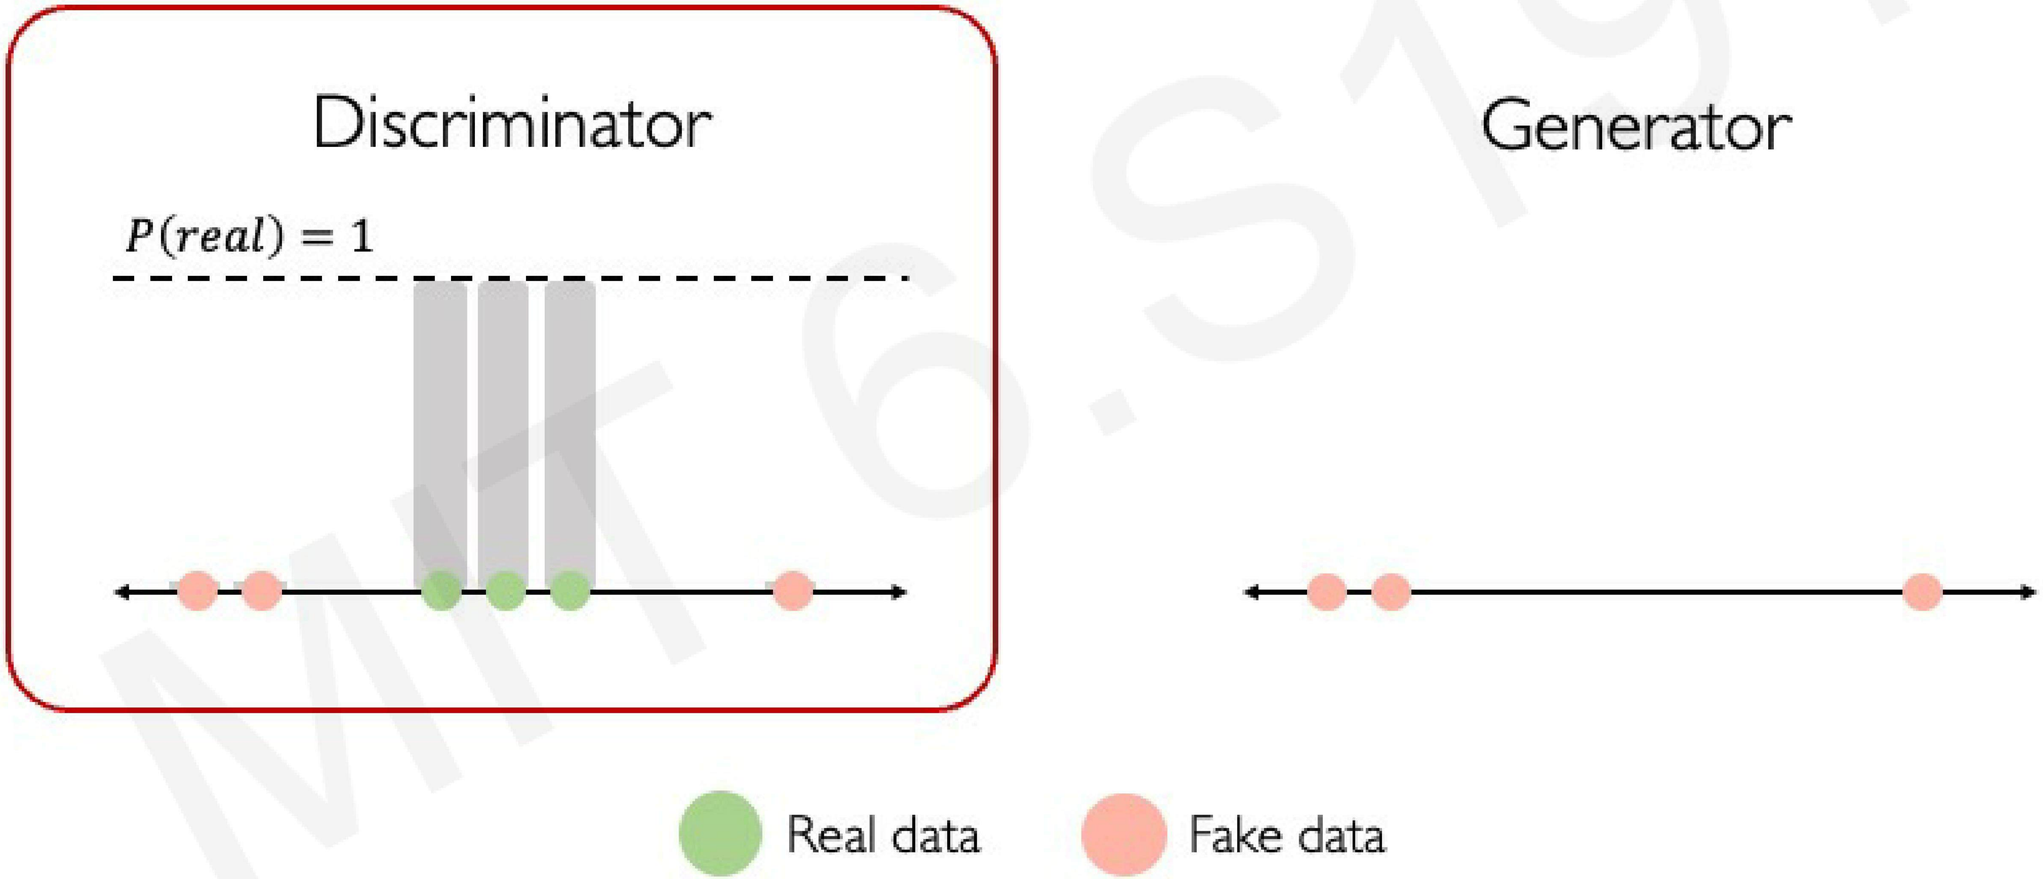

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

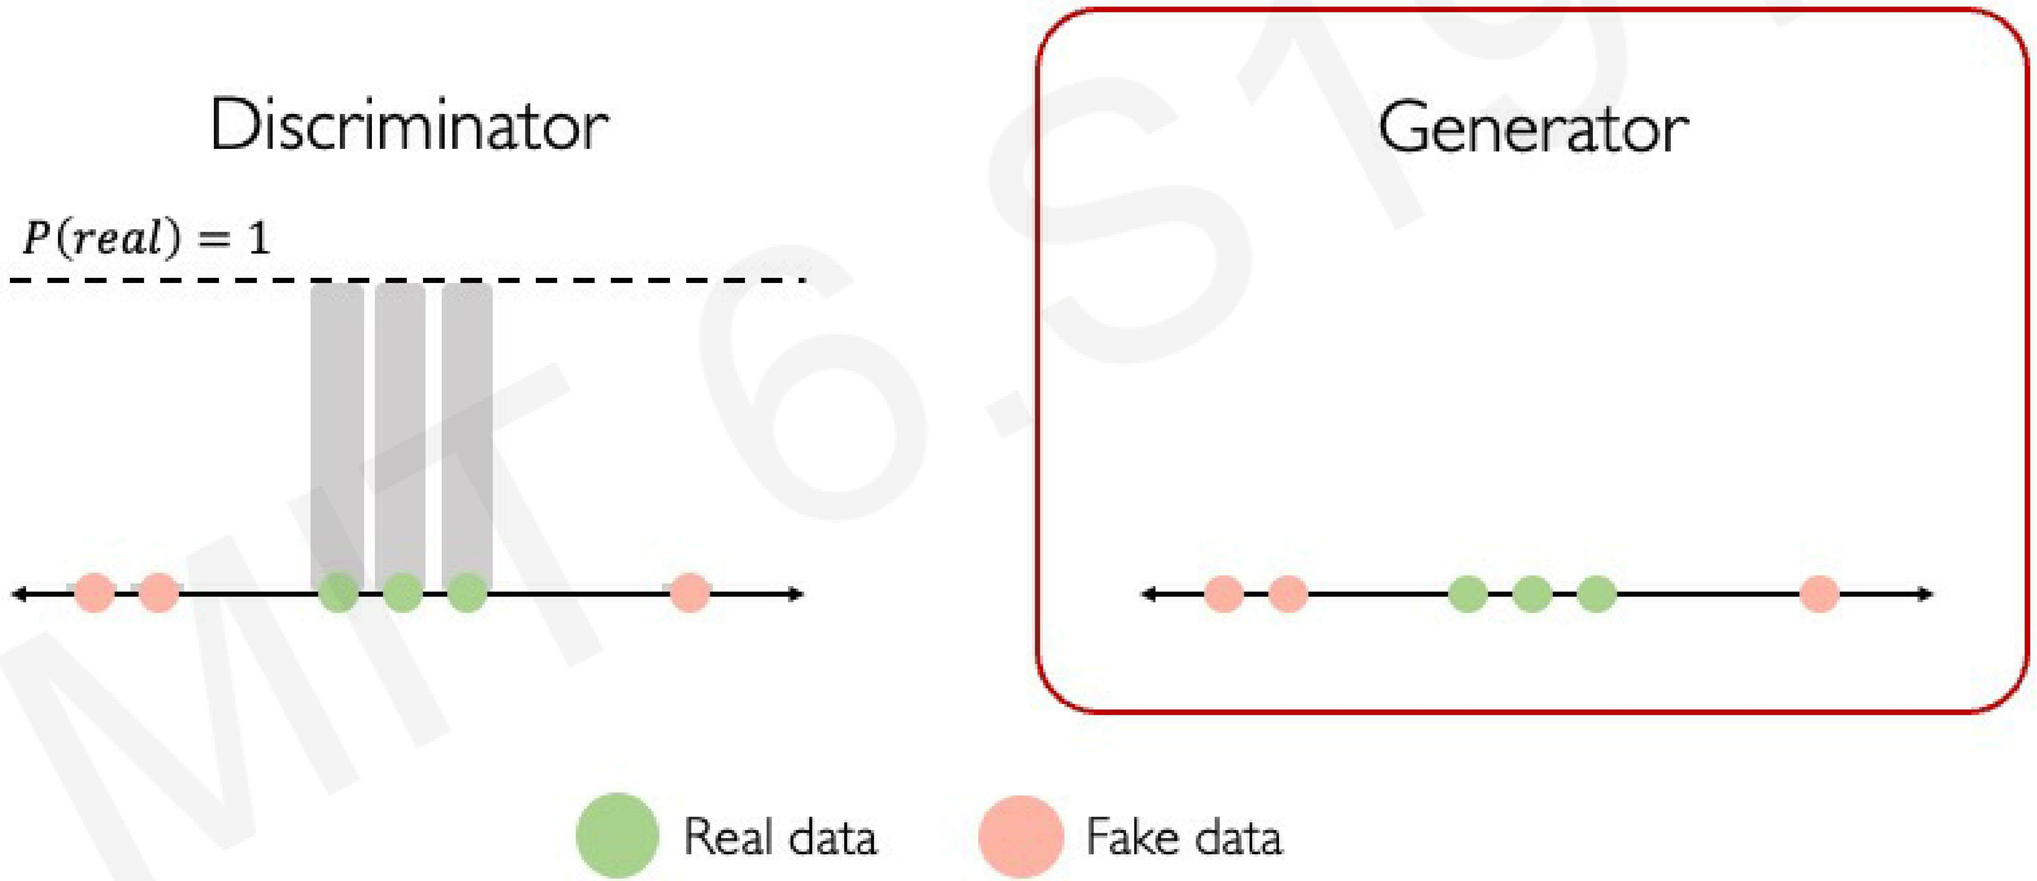

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

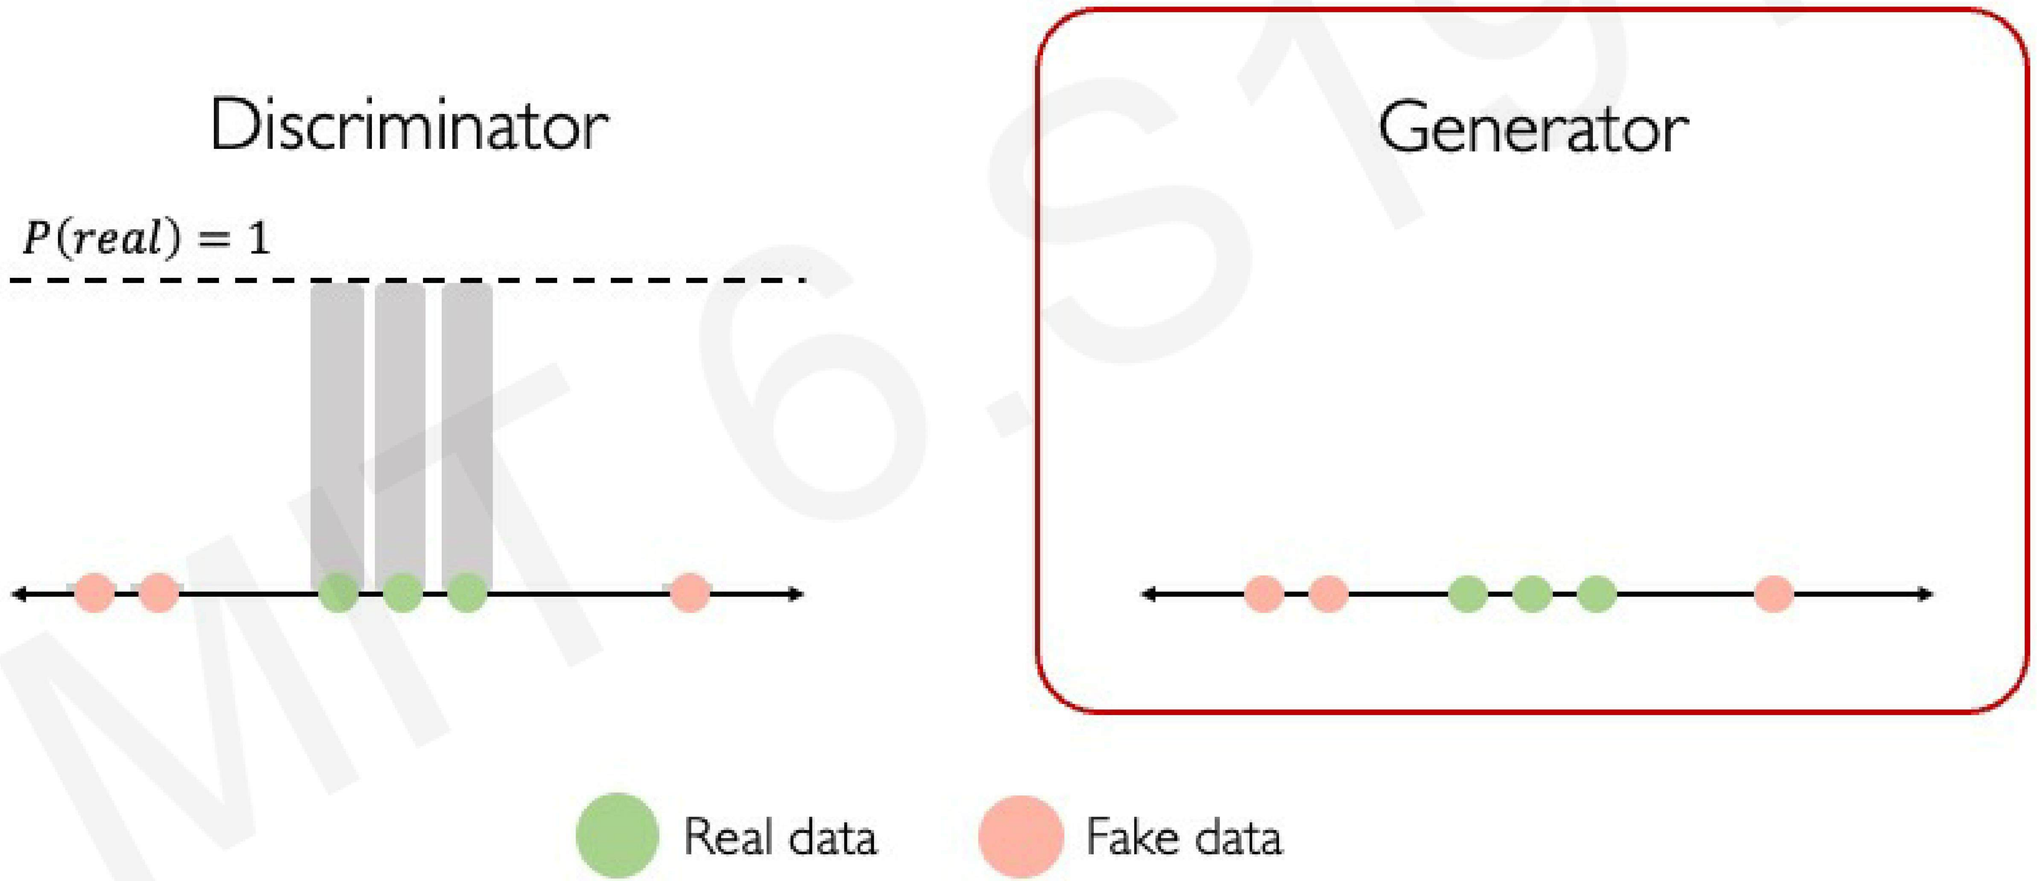

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

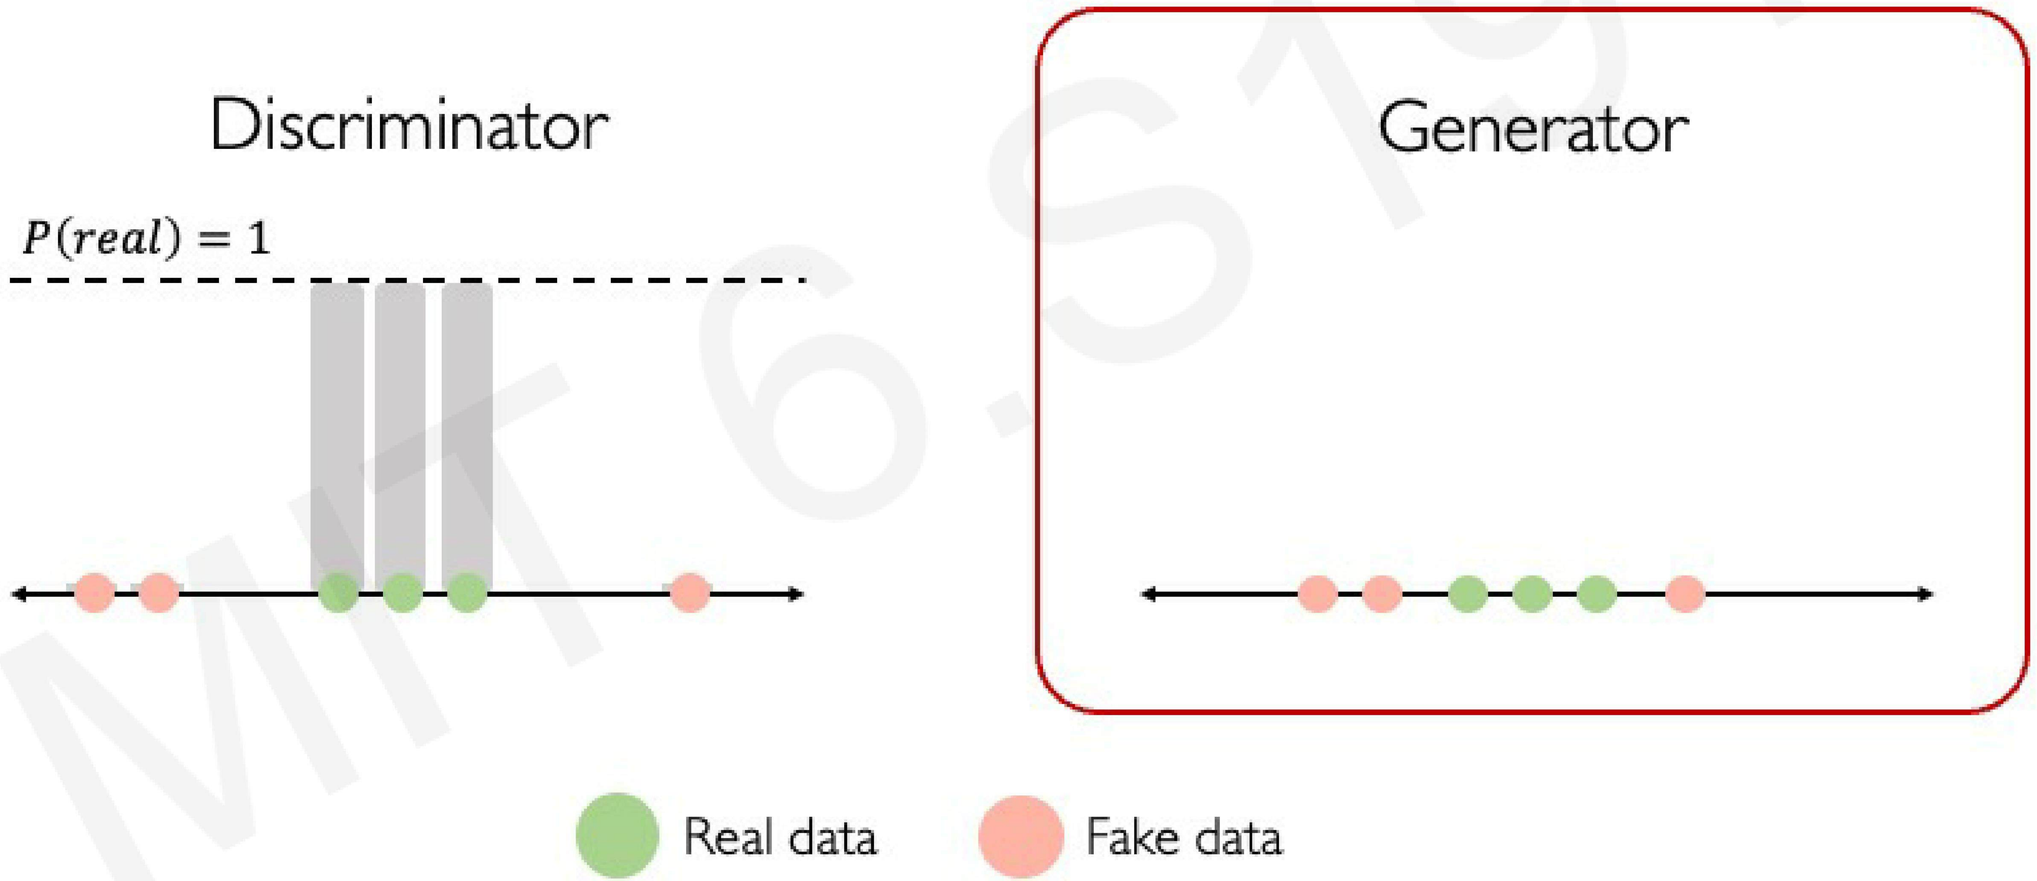

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

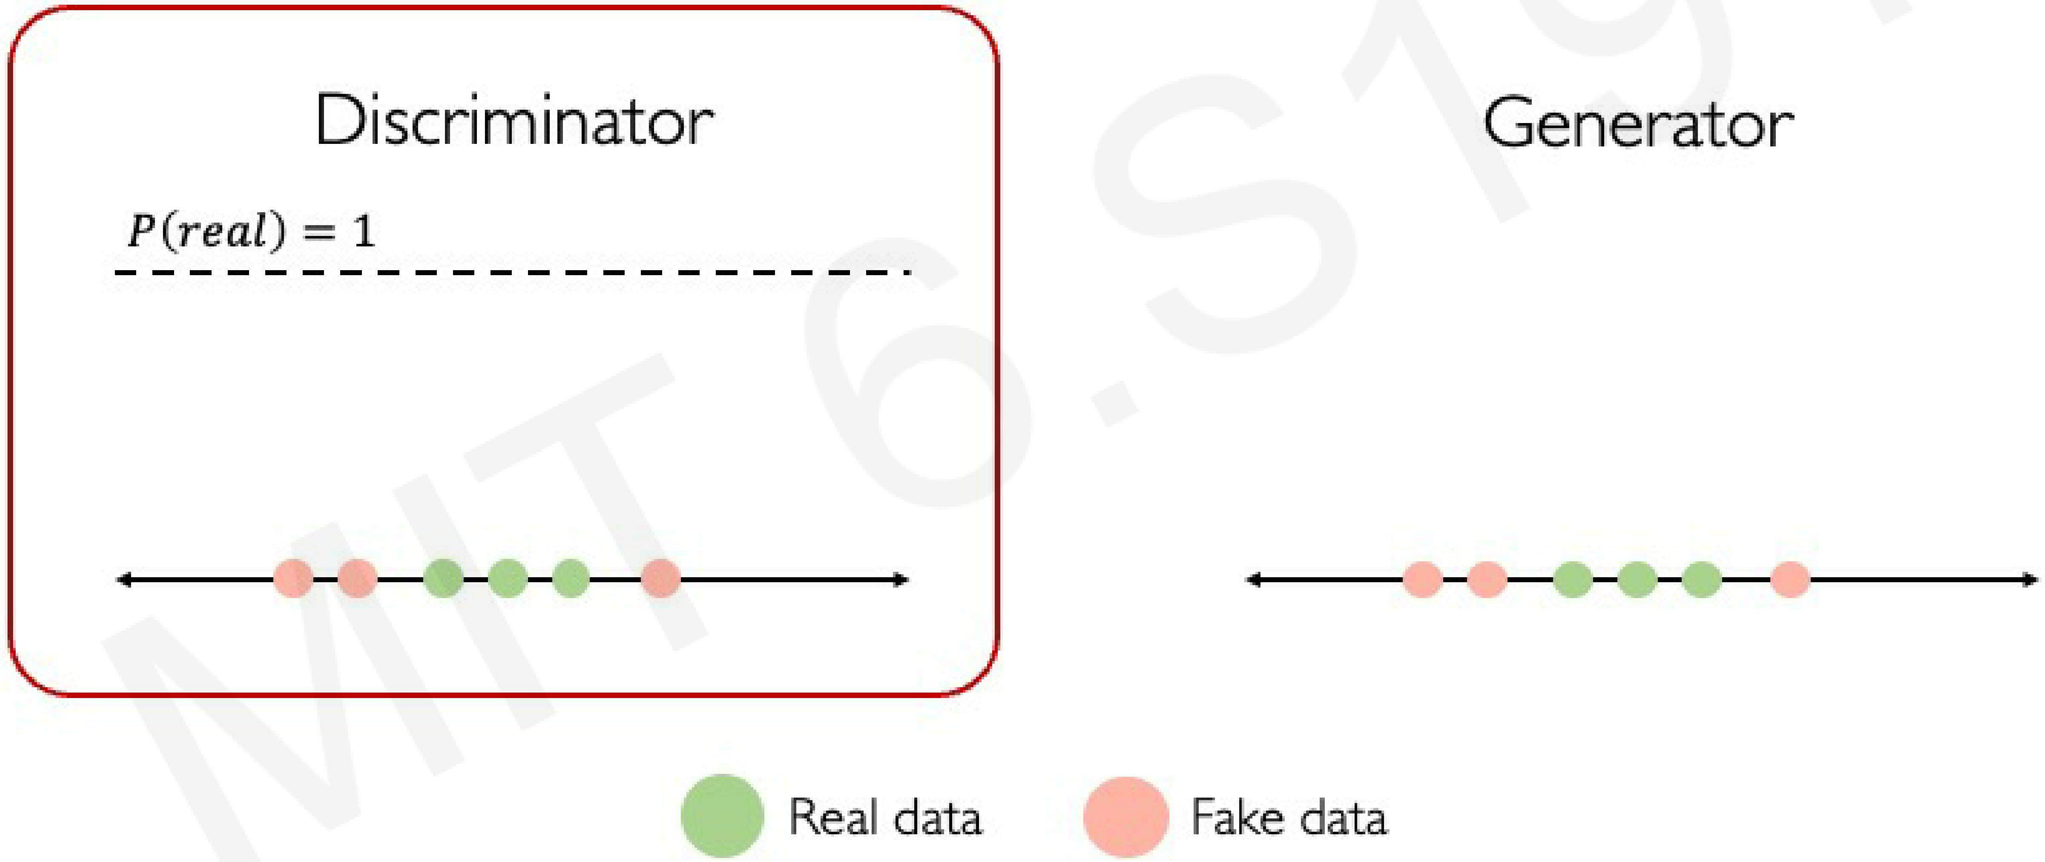

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

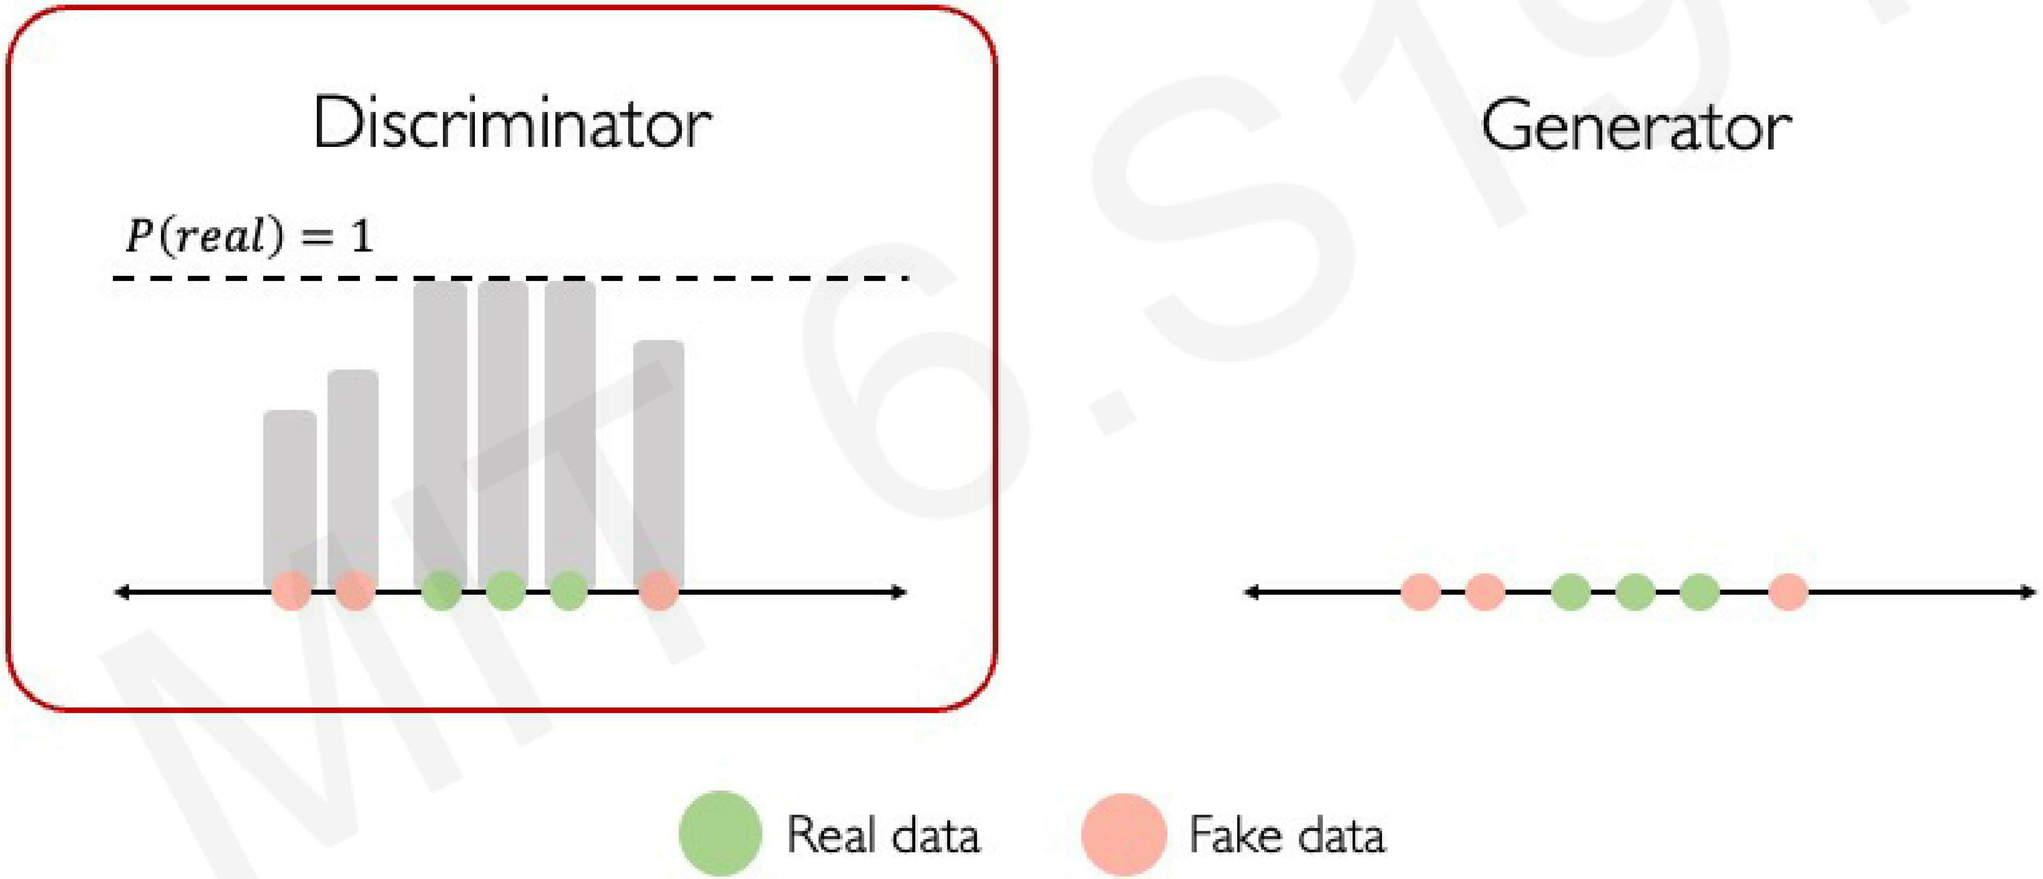

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

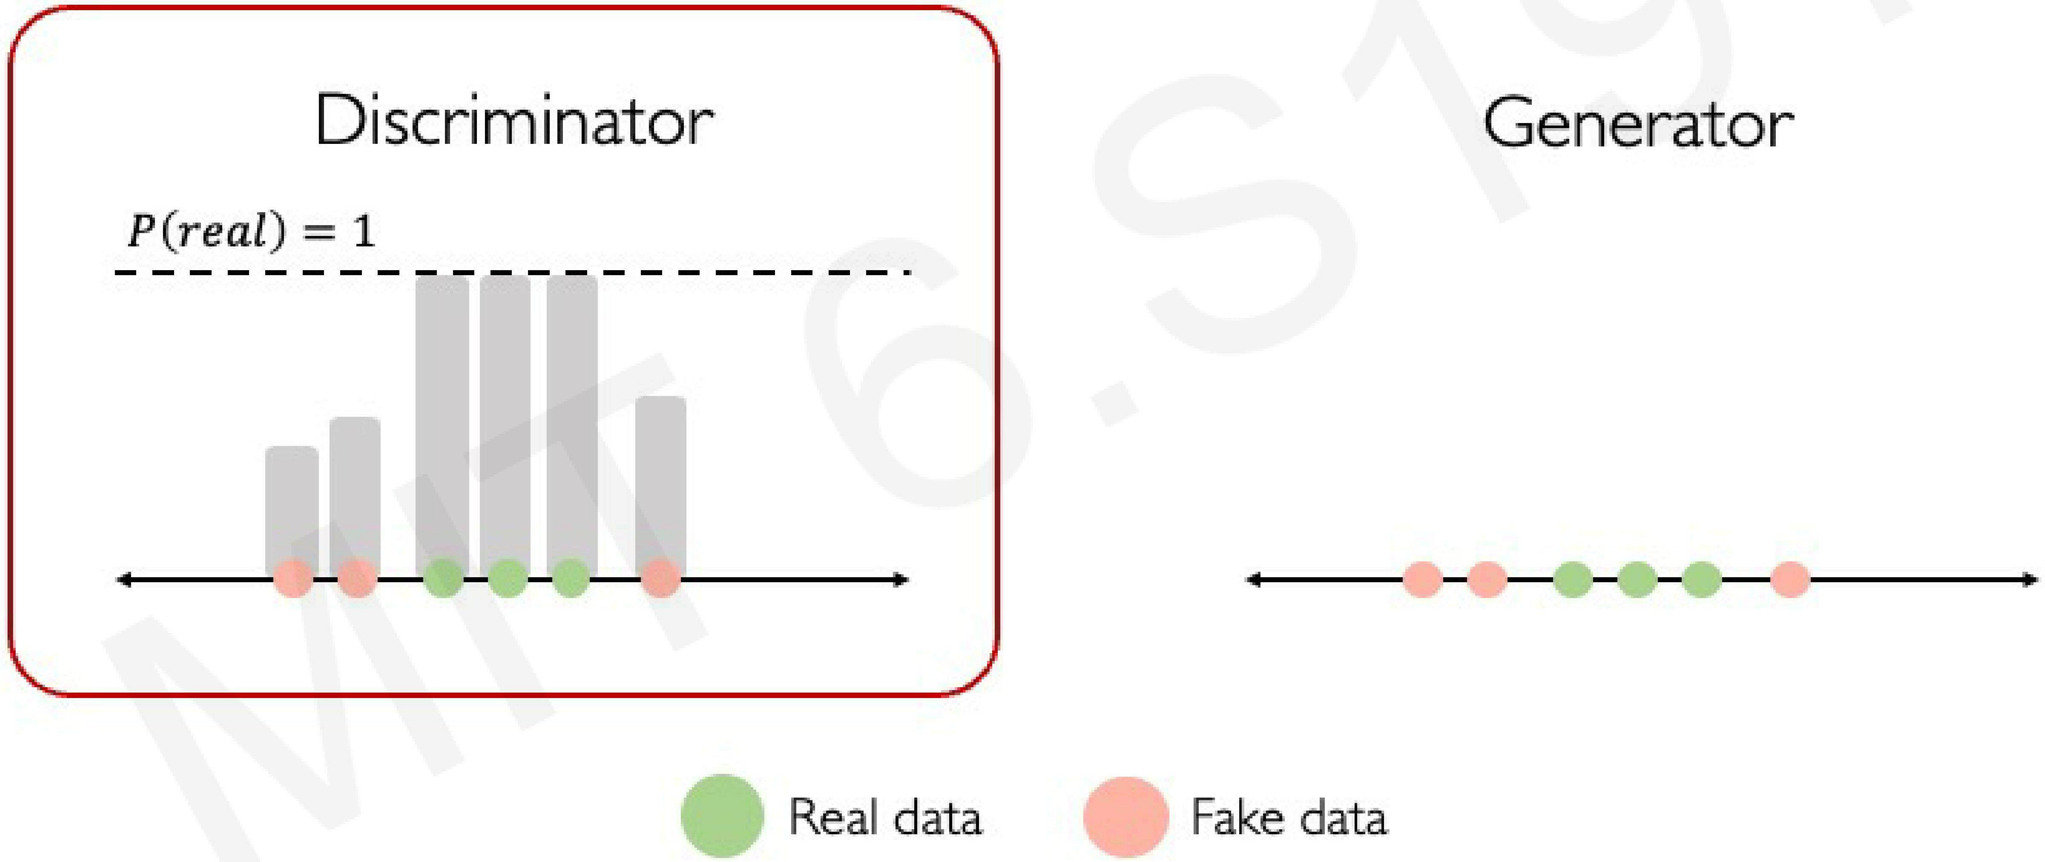

MIT 6.S191

Intuition behind GANs

Discriminator tries to predict what's real and what's fake.

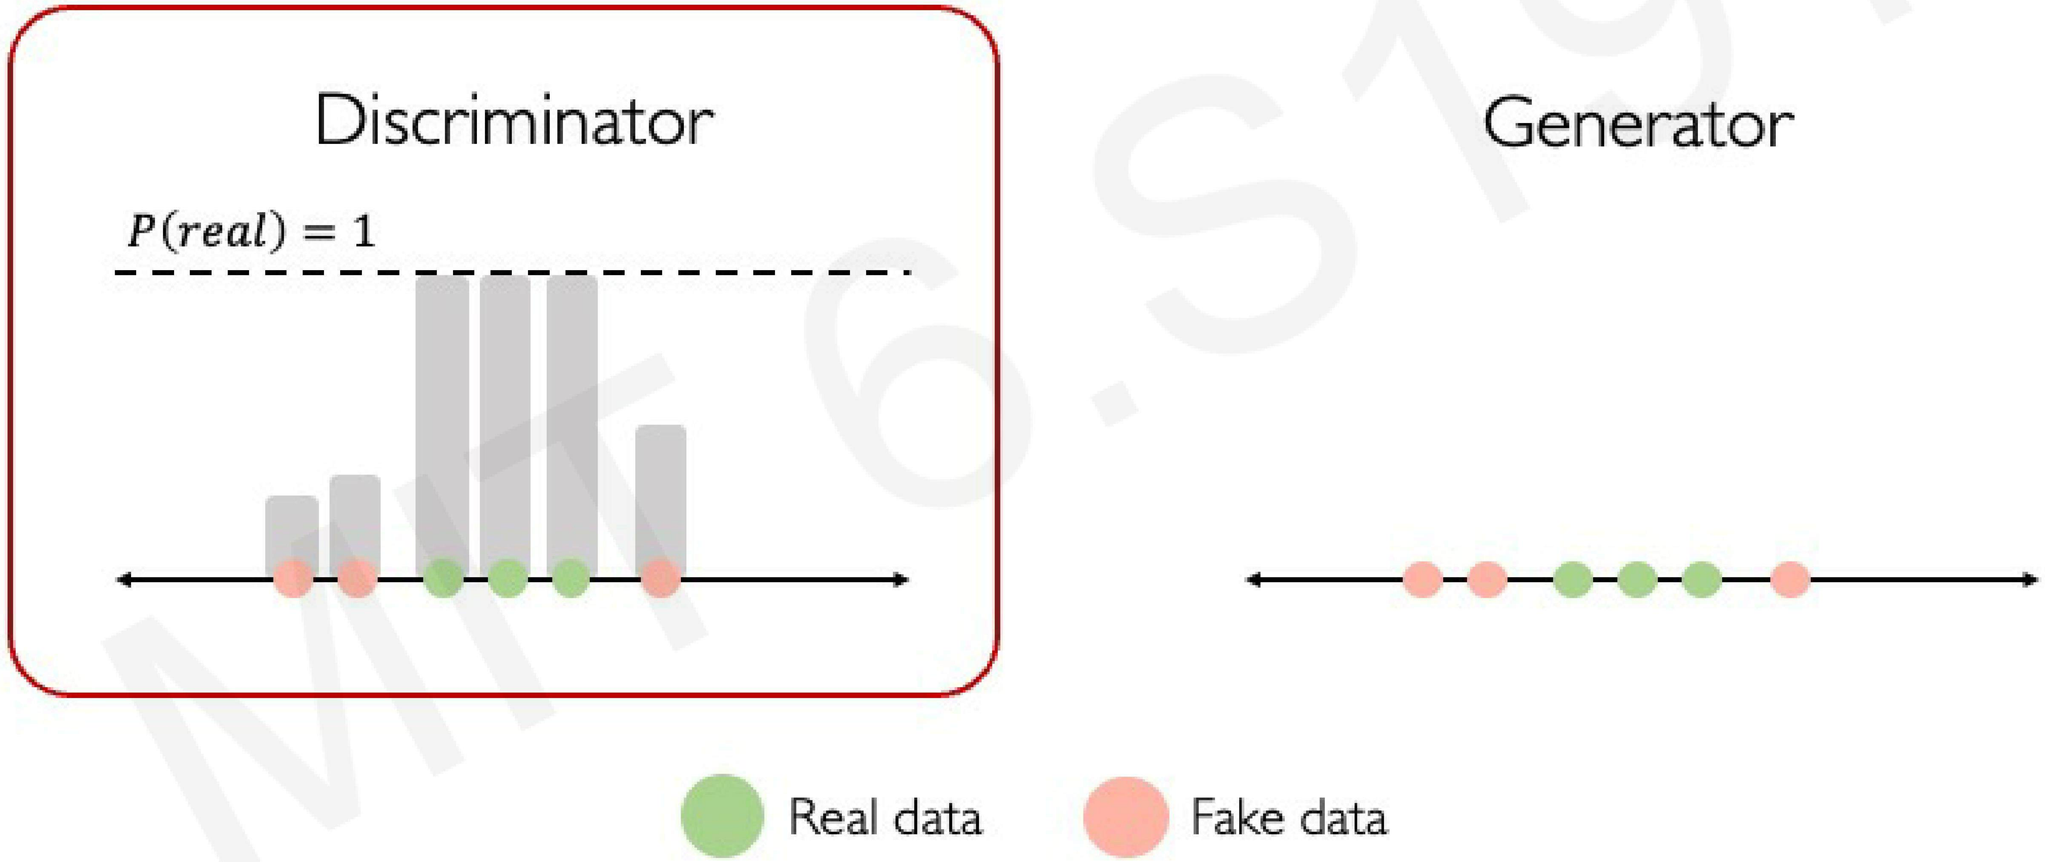

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

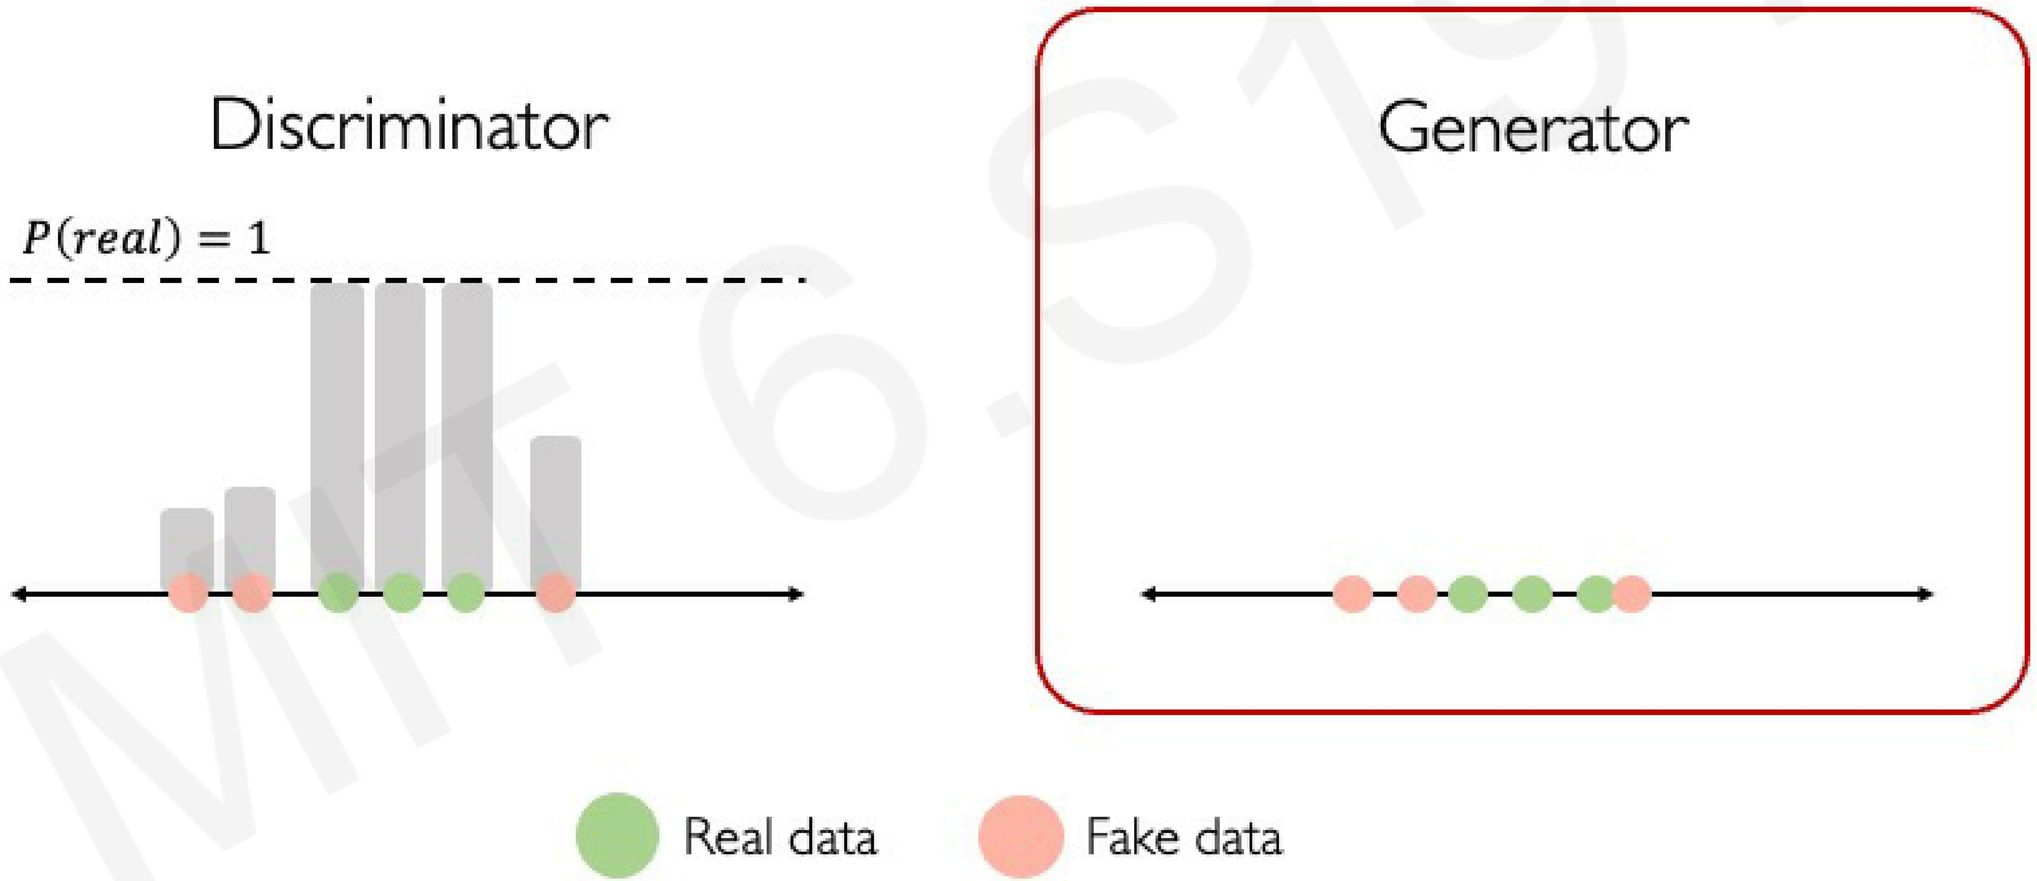

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

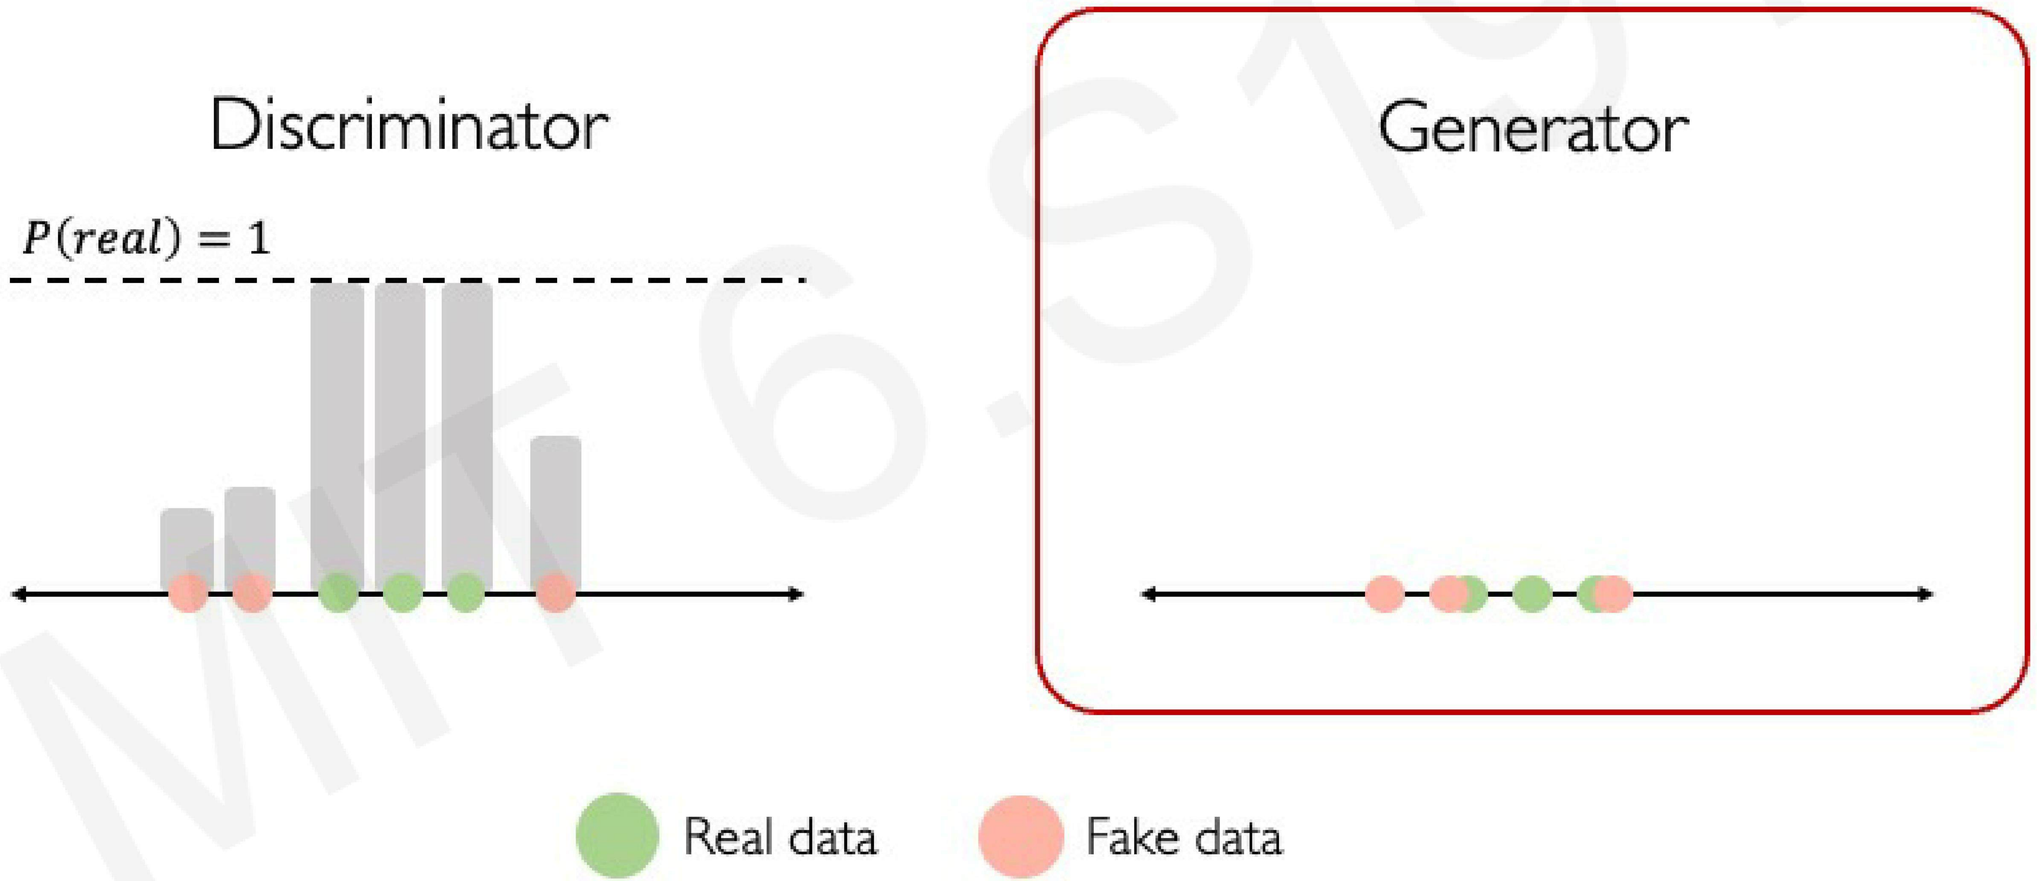

MIT 6.S191

Intuition behind GANs

Generator tries to improve its imitation of the data.

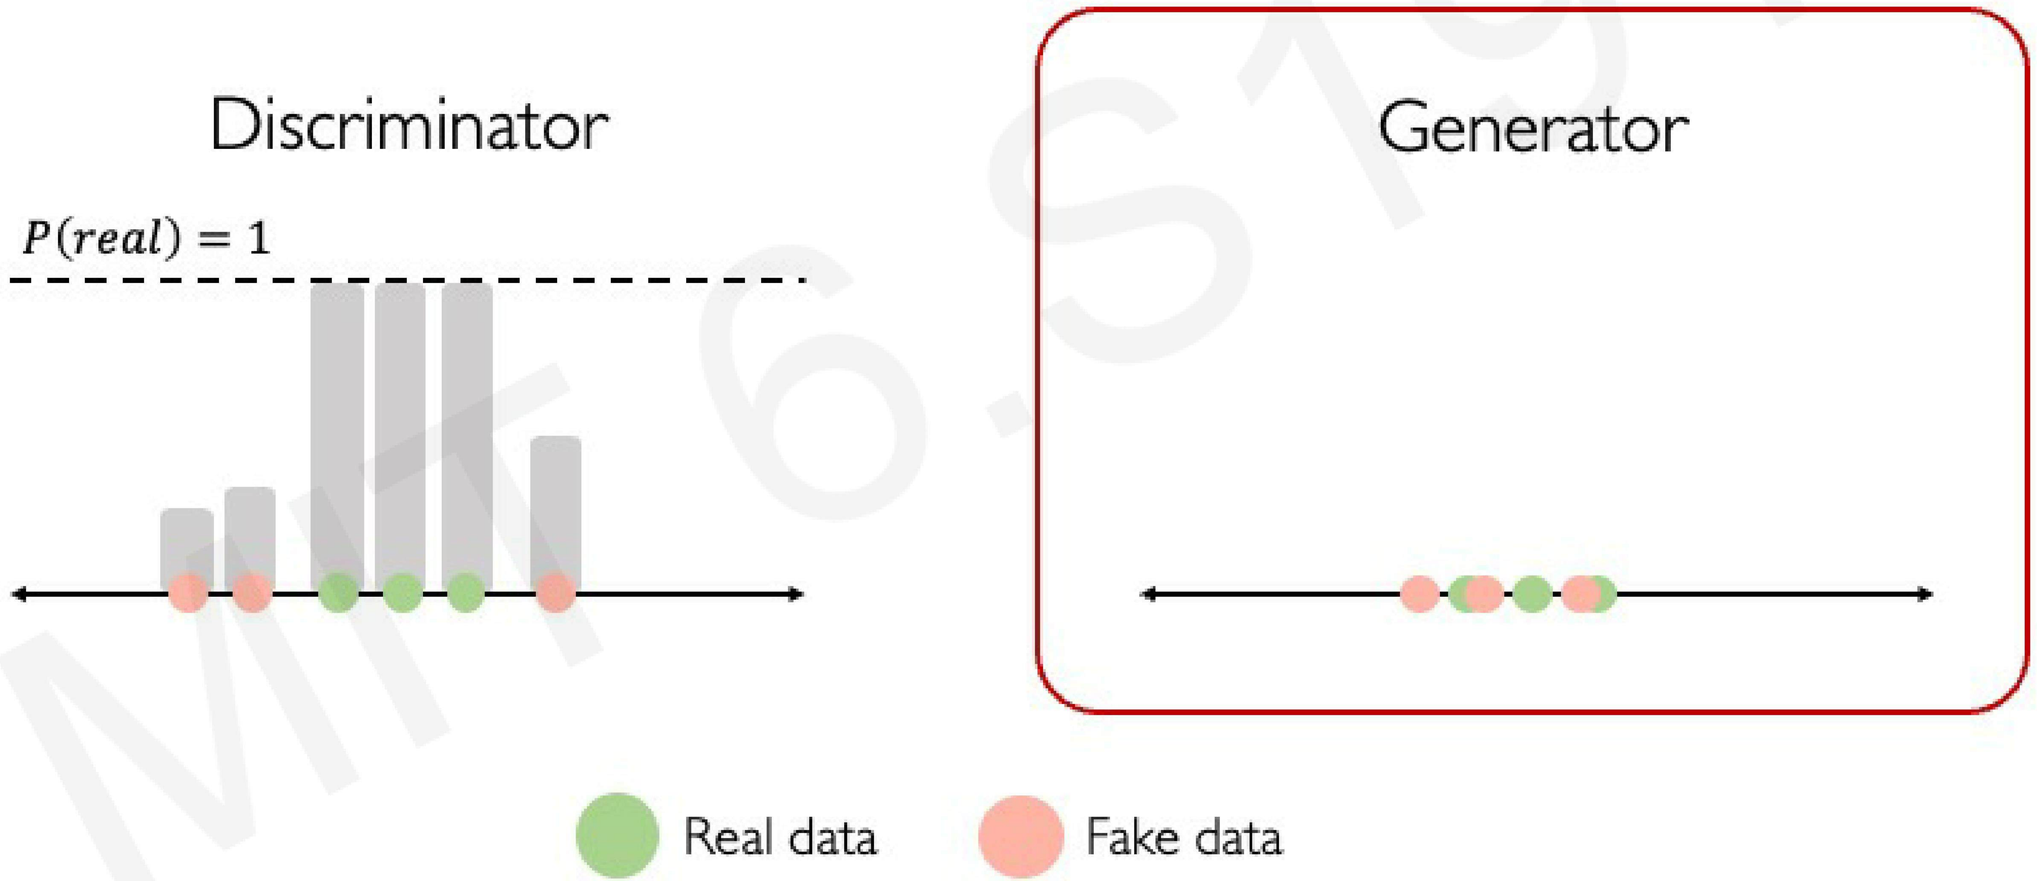

MIT 6.S191

Intuition behind GANs

Discriminator tries to identify real data from fakes created by the generator.
Generator tries to create imitations of data to trick the discriminator.

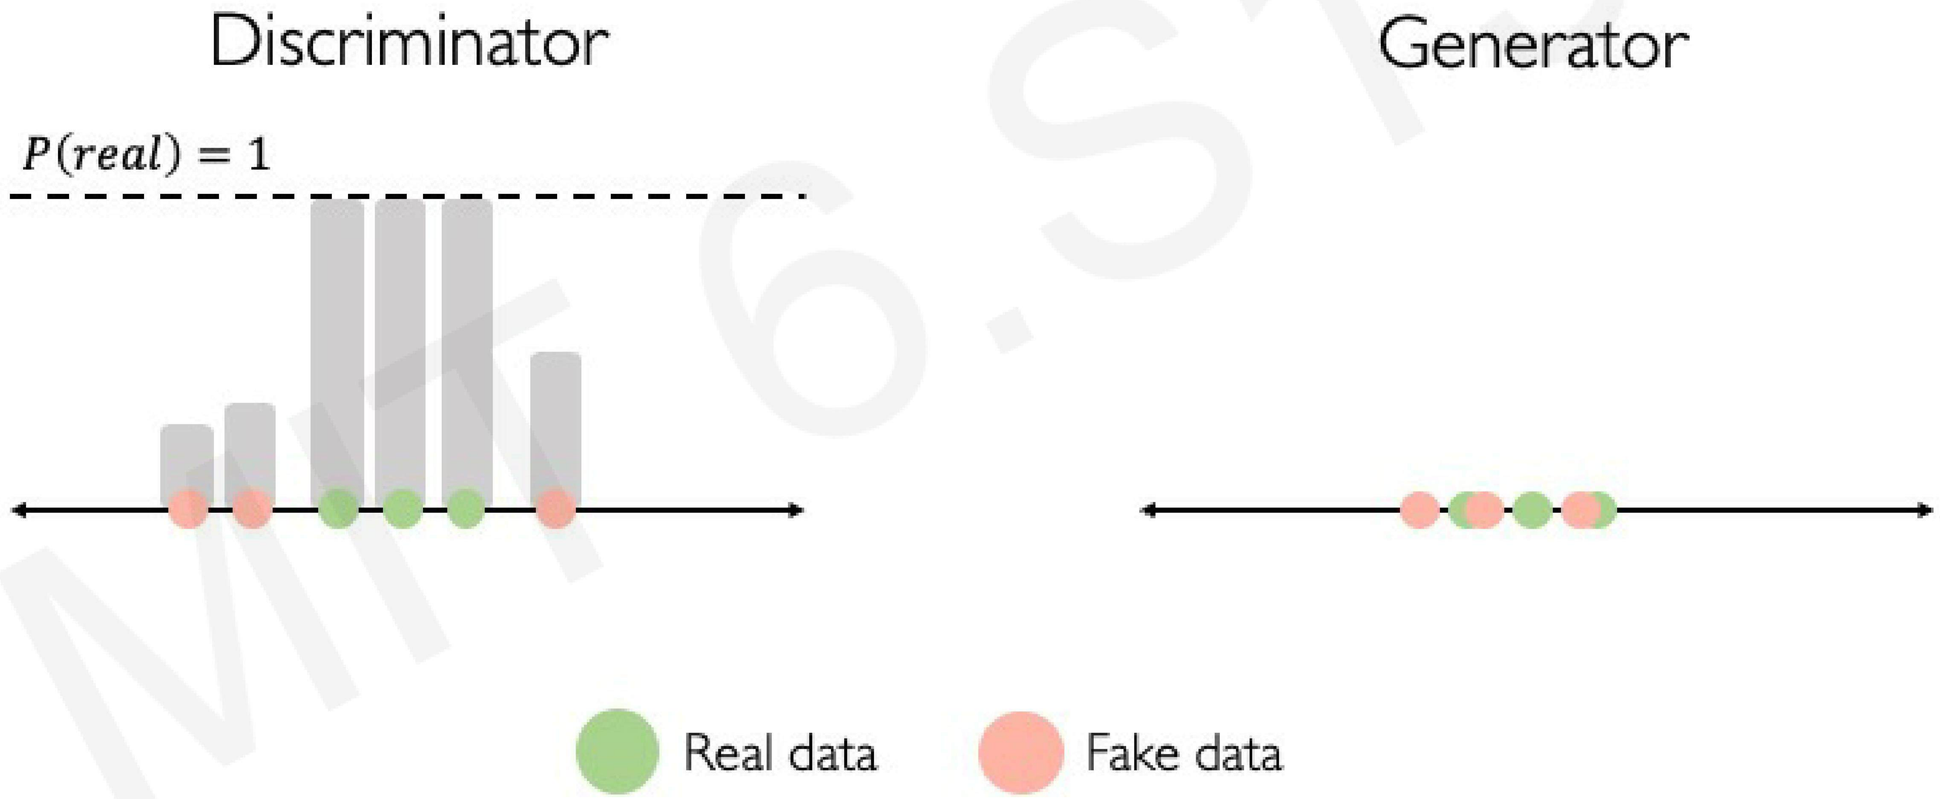

MIT 6.S191

Training GANs

G tries to synthesize fake instances that fool **D**

D tries to identify the synthesized instances

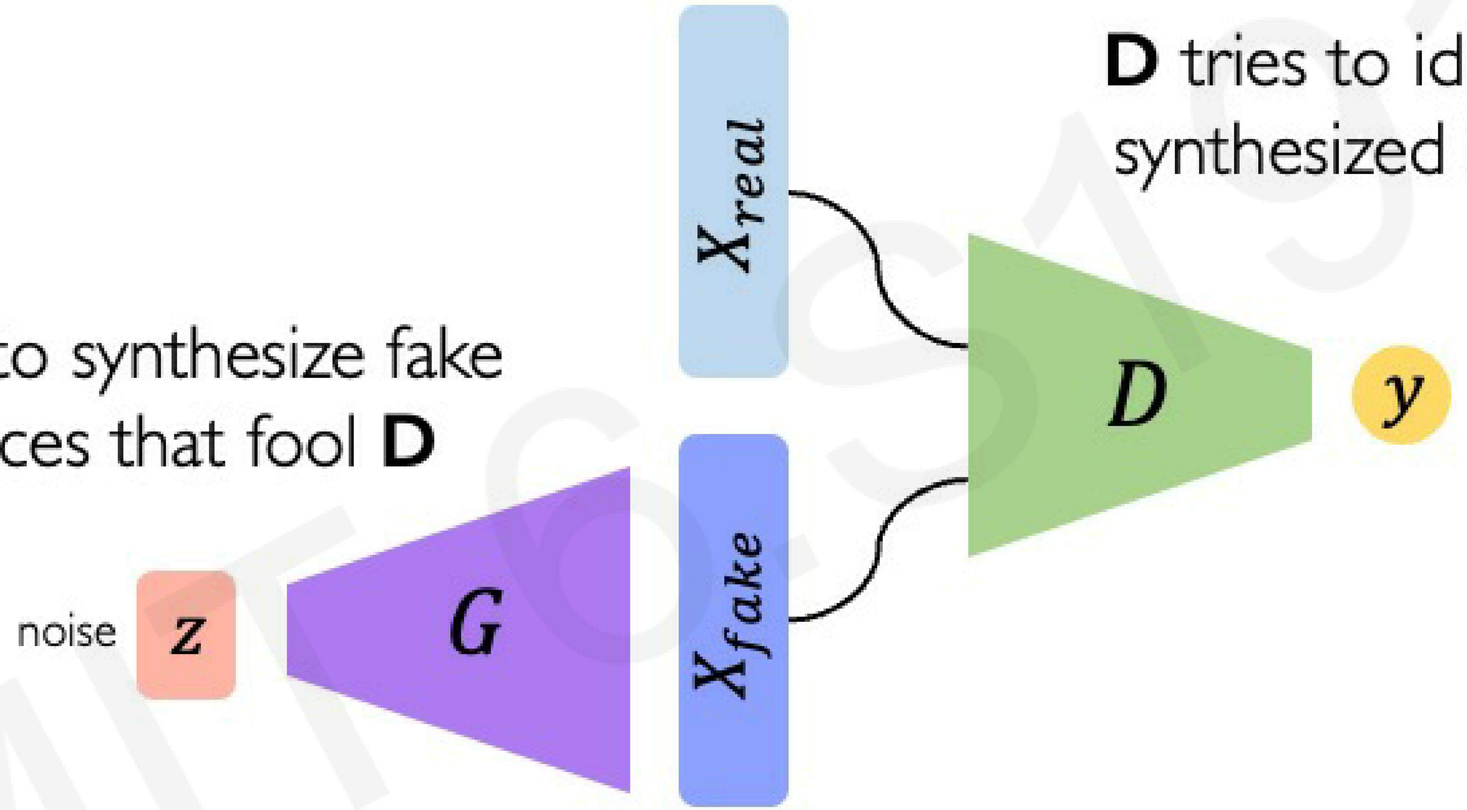

Training: adversarial objectives for **D** and **G**

Global optimum: G reproduces the true data distribution

Training GANs: loss function

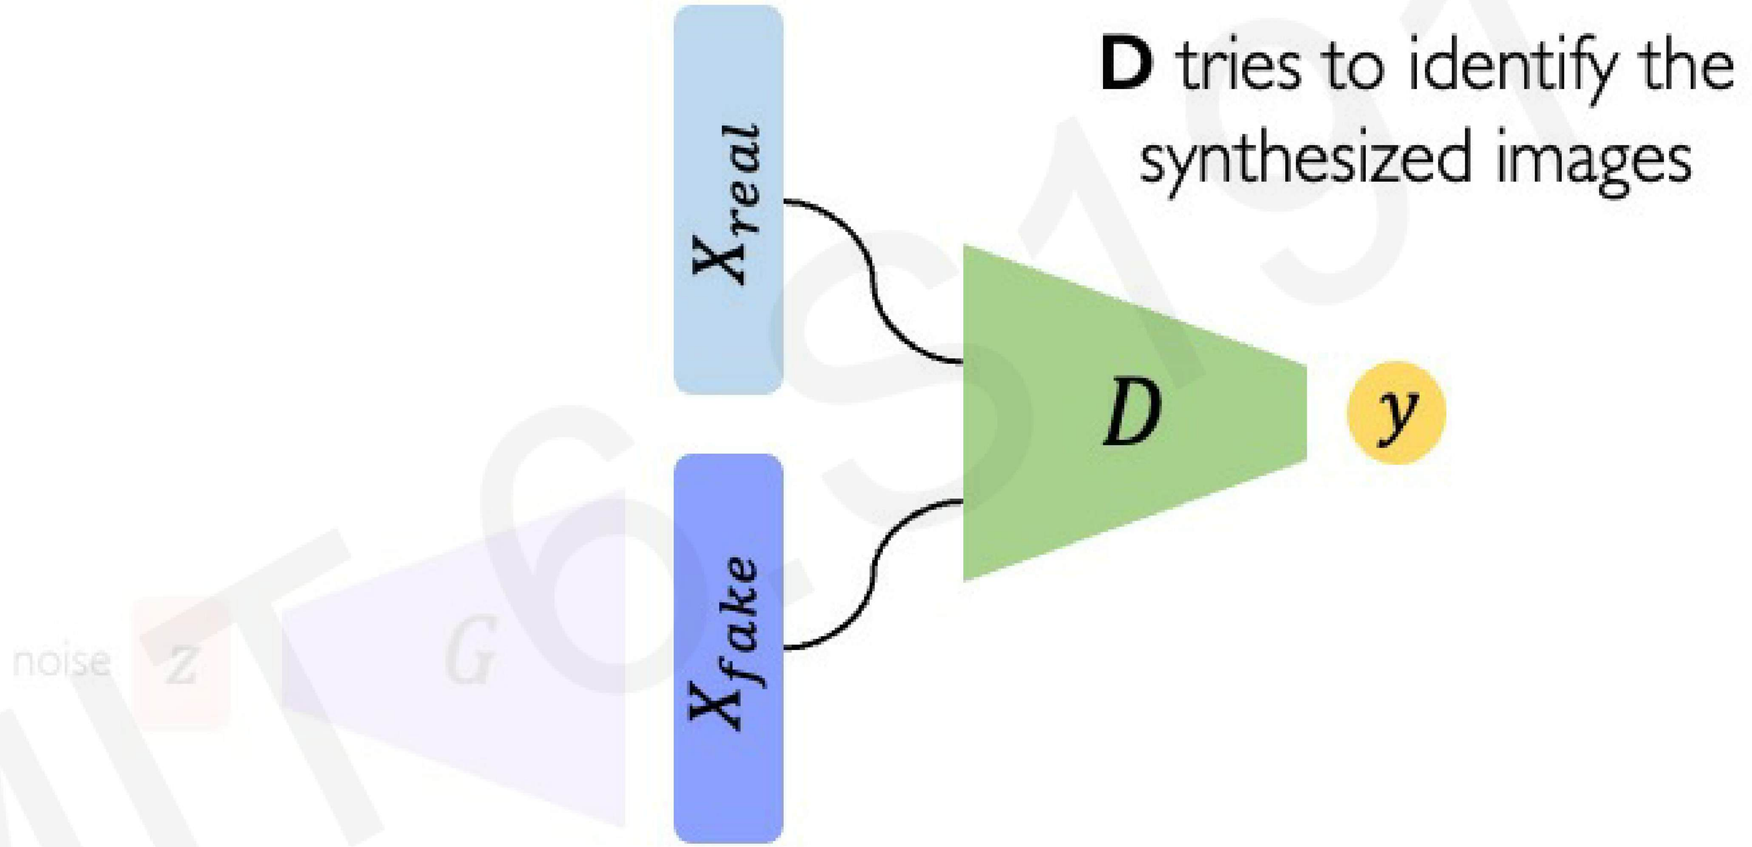

$$\arg\max_{D} \mathbb{E}_{\mathbf{z},\mathbf{x}}[\underbrace{\log D(G(\mathbf{z}))}_{\text{Fake}} + \underbrace{\log\left(1 - D(\mathbf{x})\right)}_{\text{Real}}]$$

Training GANs: loss function

G tries to synthesize fake images that fool **D**

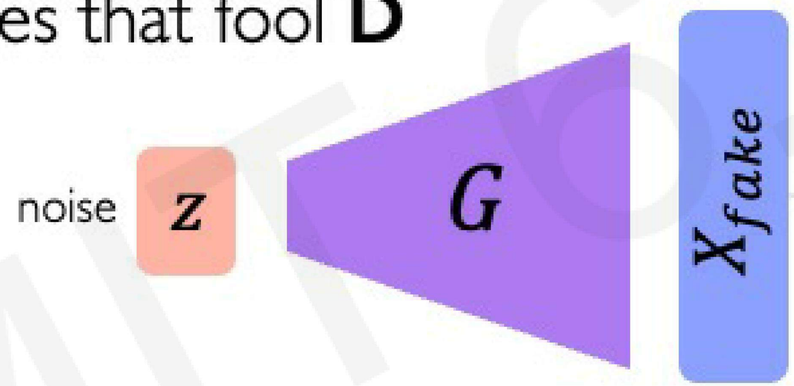

$$\arg\min_{G} \mathbb{E}_{\mathbf{z},\mathbf{x}}[\log D(G(\mathbf{z})) + \log(1 - D(\mathbf{x}))]$$

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Training GANs: loss function

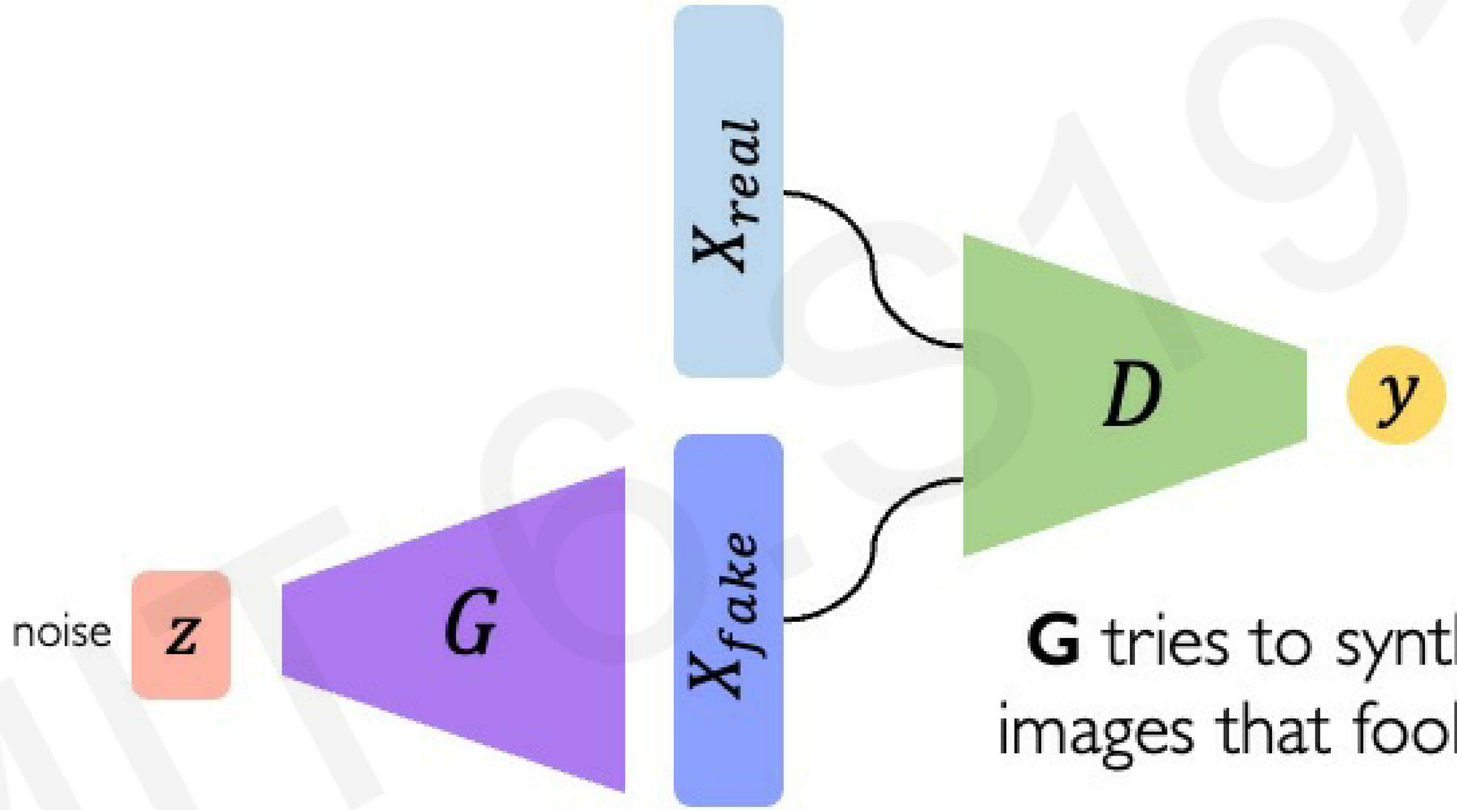

G tries to synthesize fake images that fool the *best* **D**

$$\arg\min_{G}\max_{D} \mathbb{E}_{\mathbf{z},\mathbf{x}}[\log D(G(\mathbf{z})) + \log(1 - D(\mathbf{x}))]$$

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Generating new data with GANs

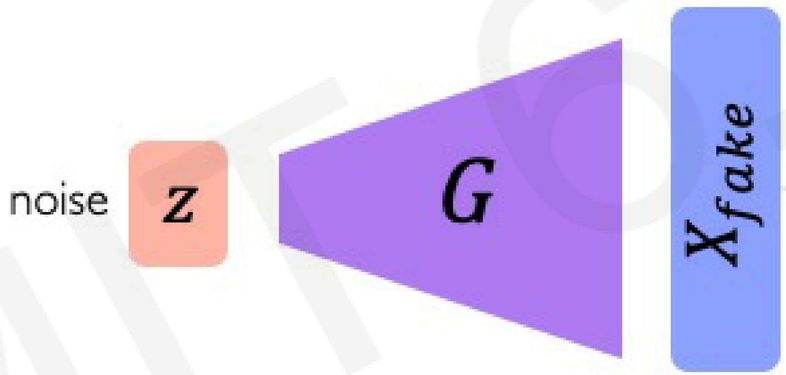

After training, use generator network to create **new data** that's never been seen before.

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

GANs are distribution transformers

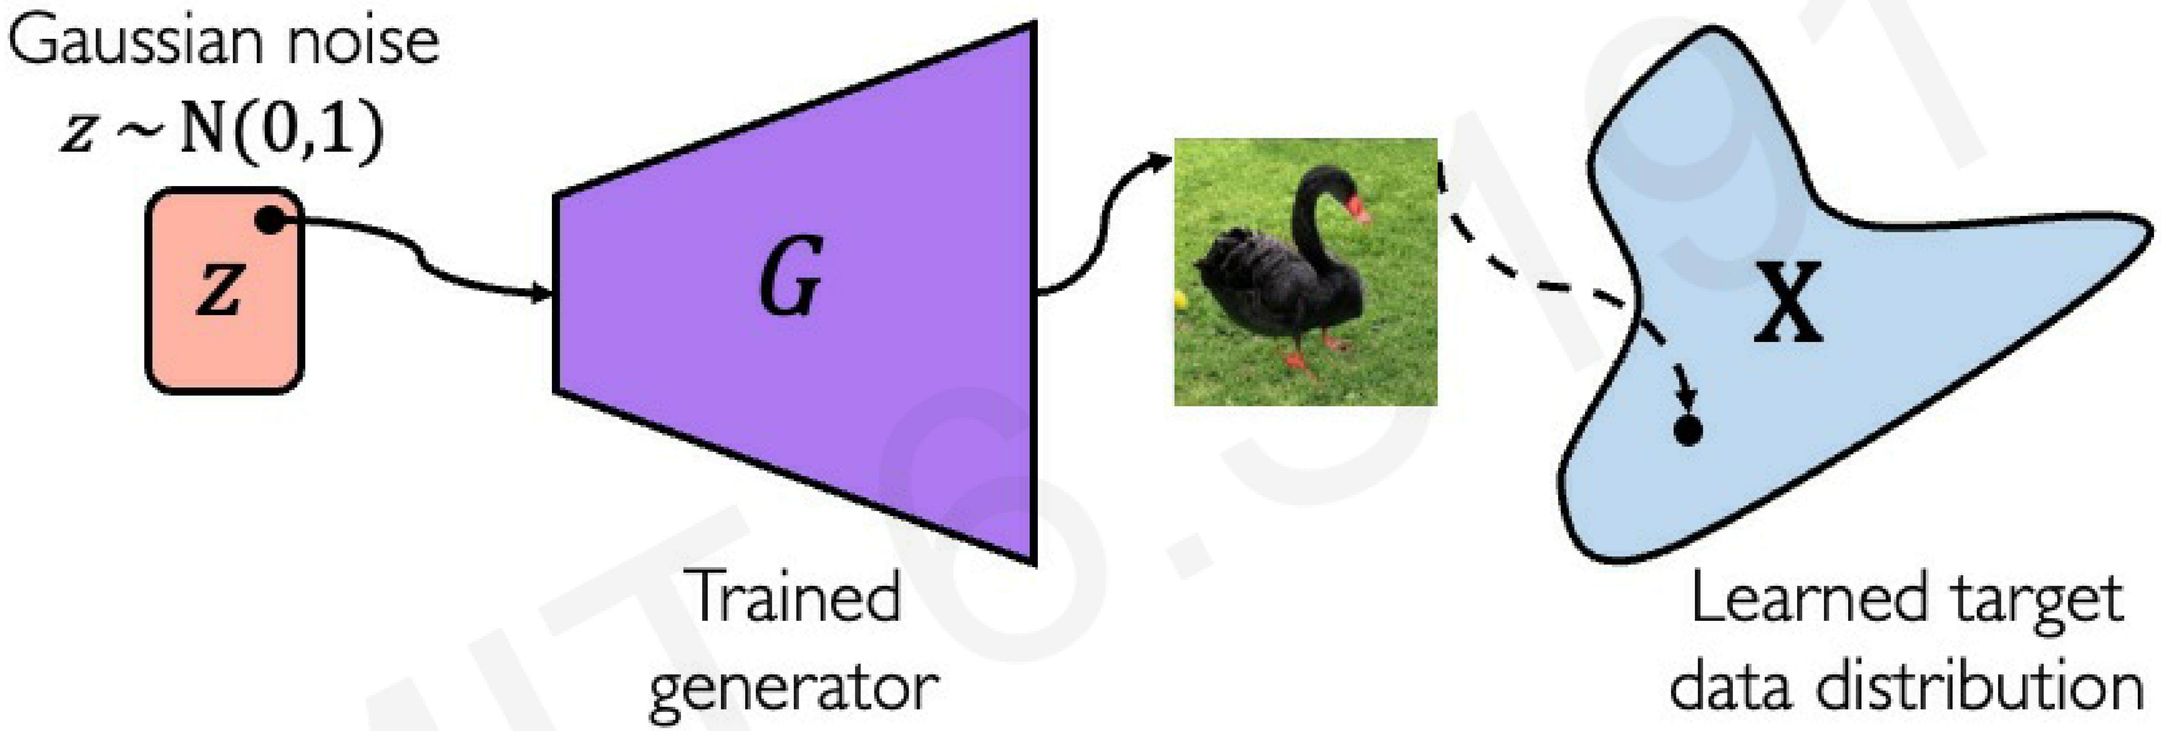

MIT 6.S191

GANs are distribution transformers

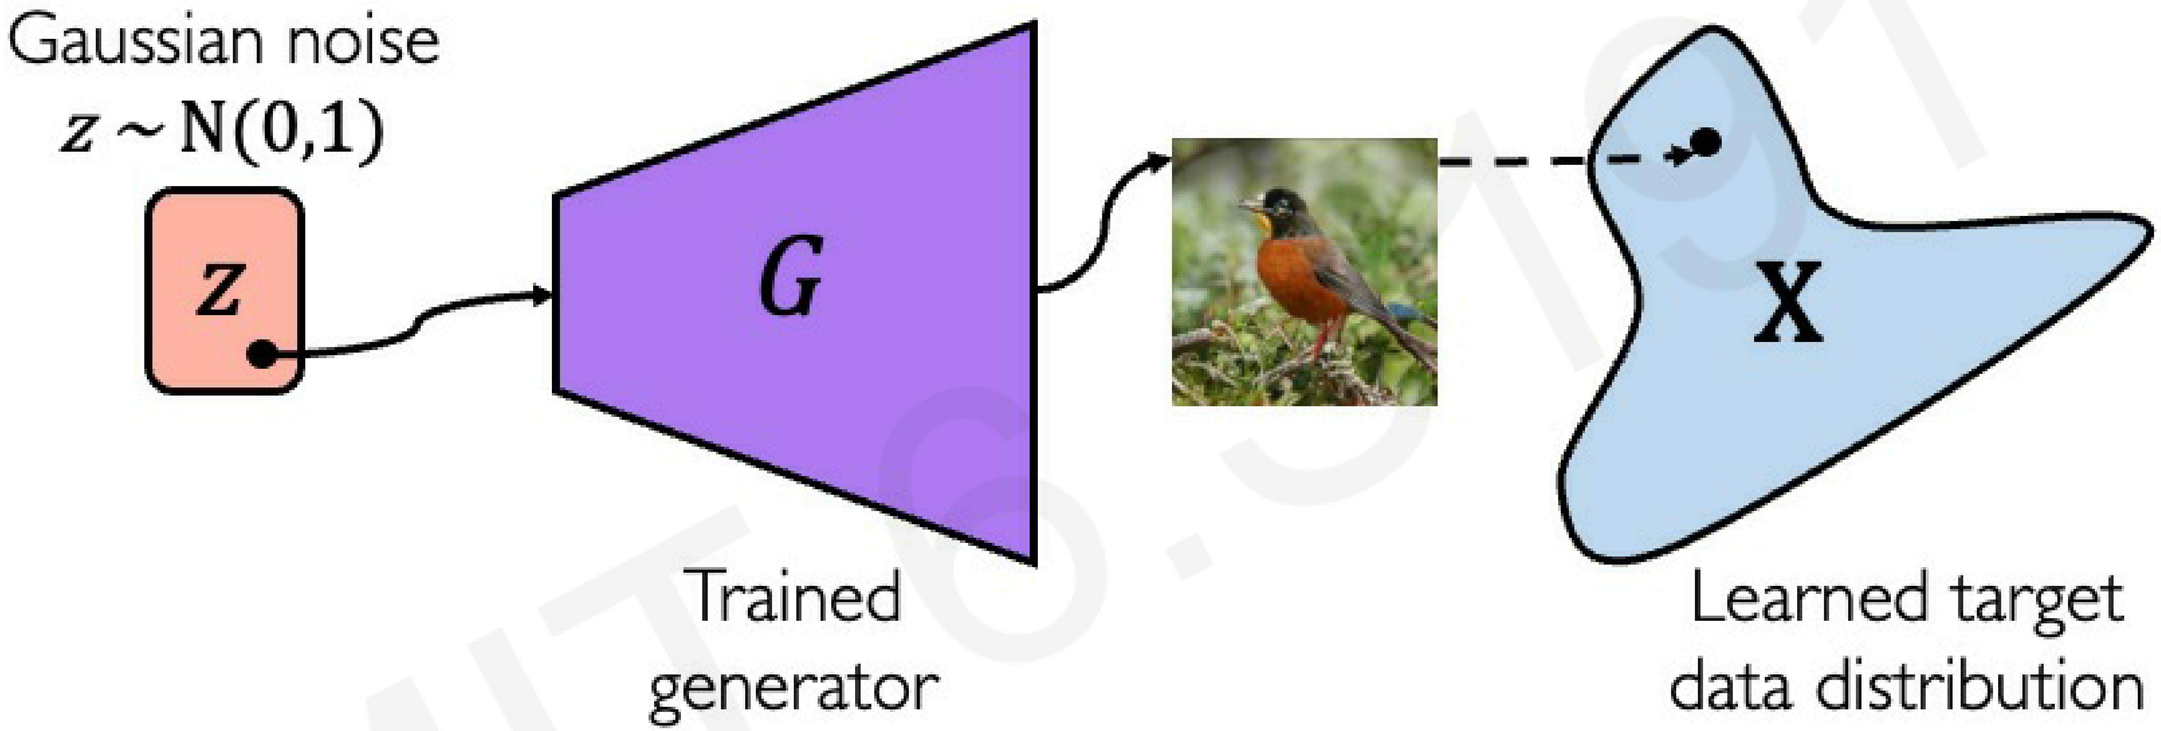

MIT.S191

GANs are distribution transformers

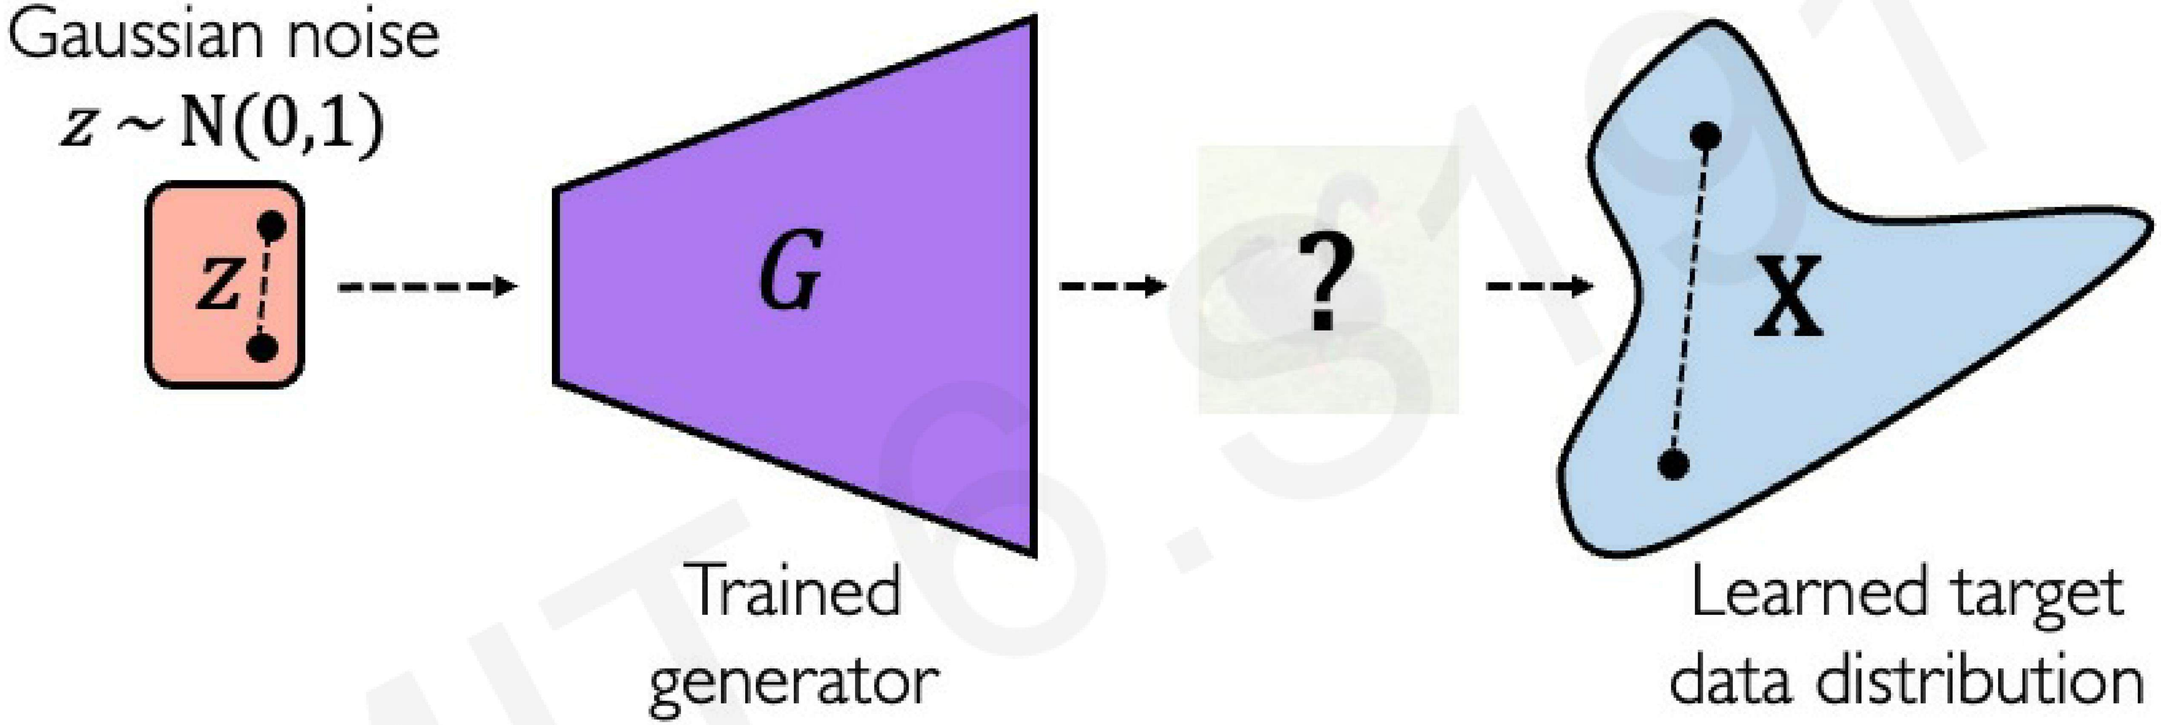

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

CycleGAN: domain transformation

CycleGAN learns transformations across domains with unpaired data.

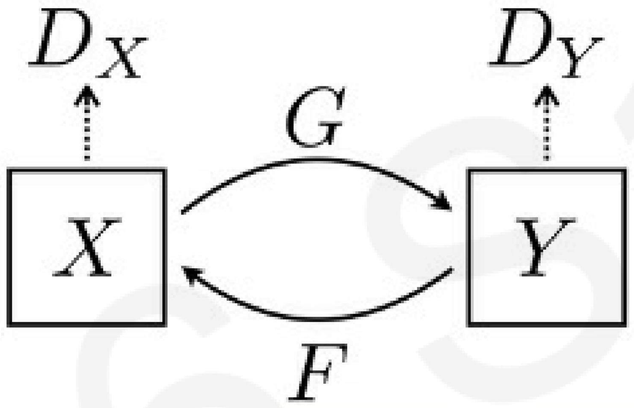

Distribution transformations

GANs:

Gaussian noise

$z \sim N(0,1)$

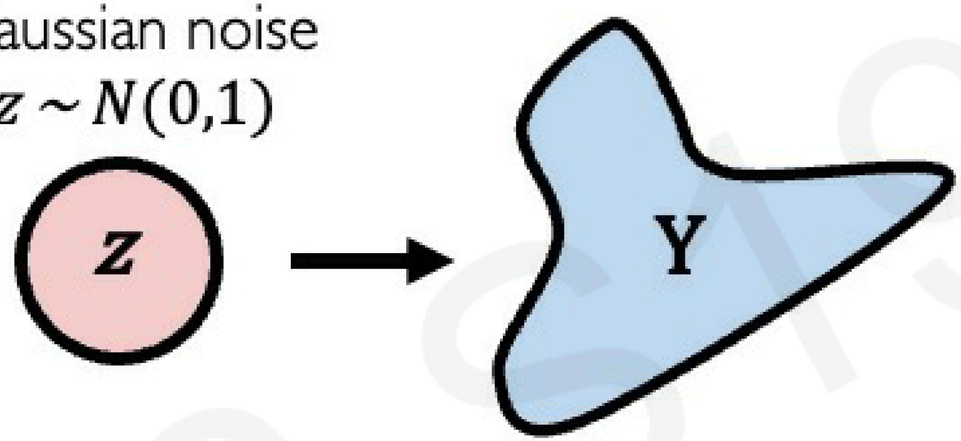

Gaussian noise → target data manifold

CycleGANs:

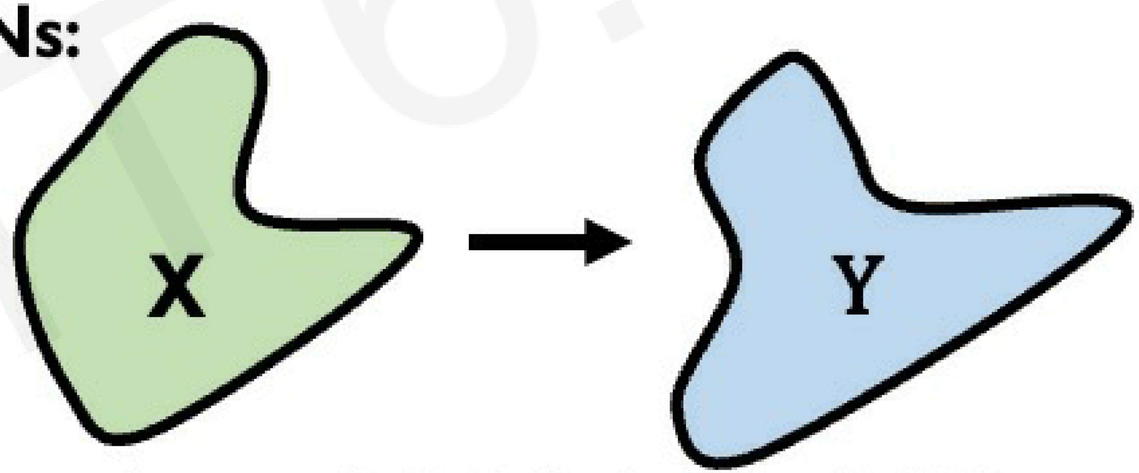

data manifold X → data manifold Y

MIT 6.S191

CycleGAN: transforming speech

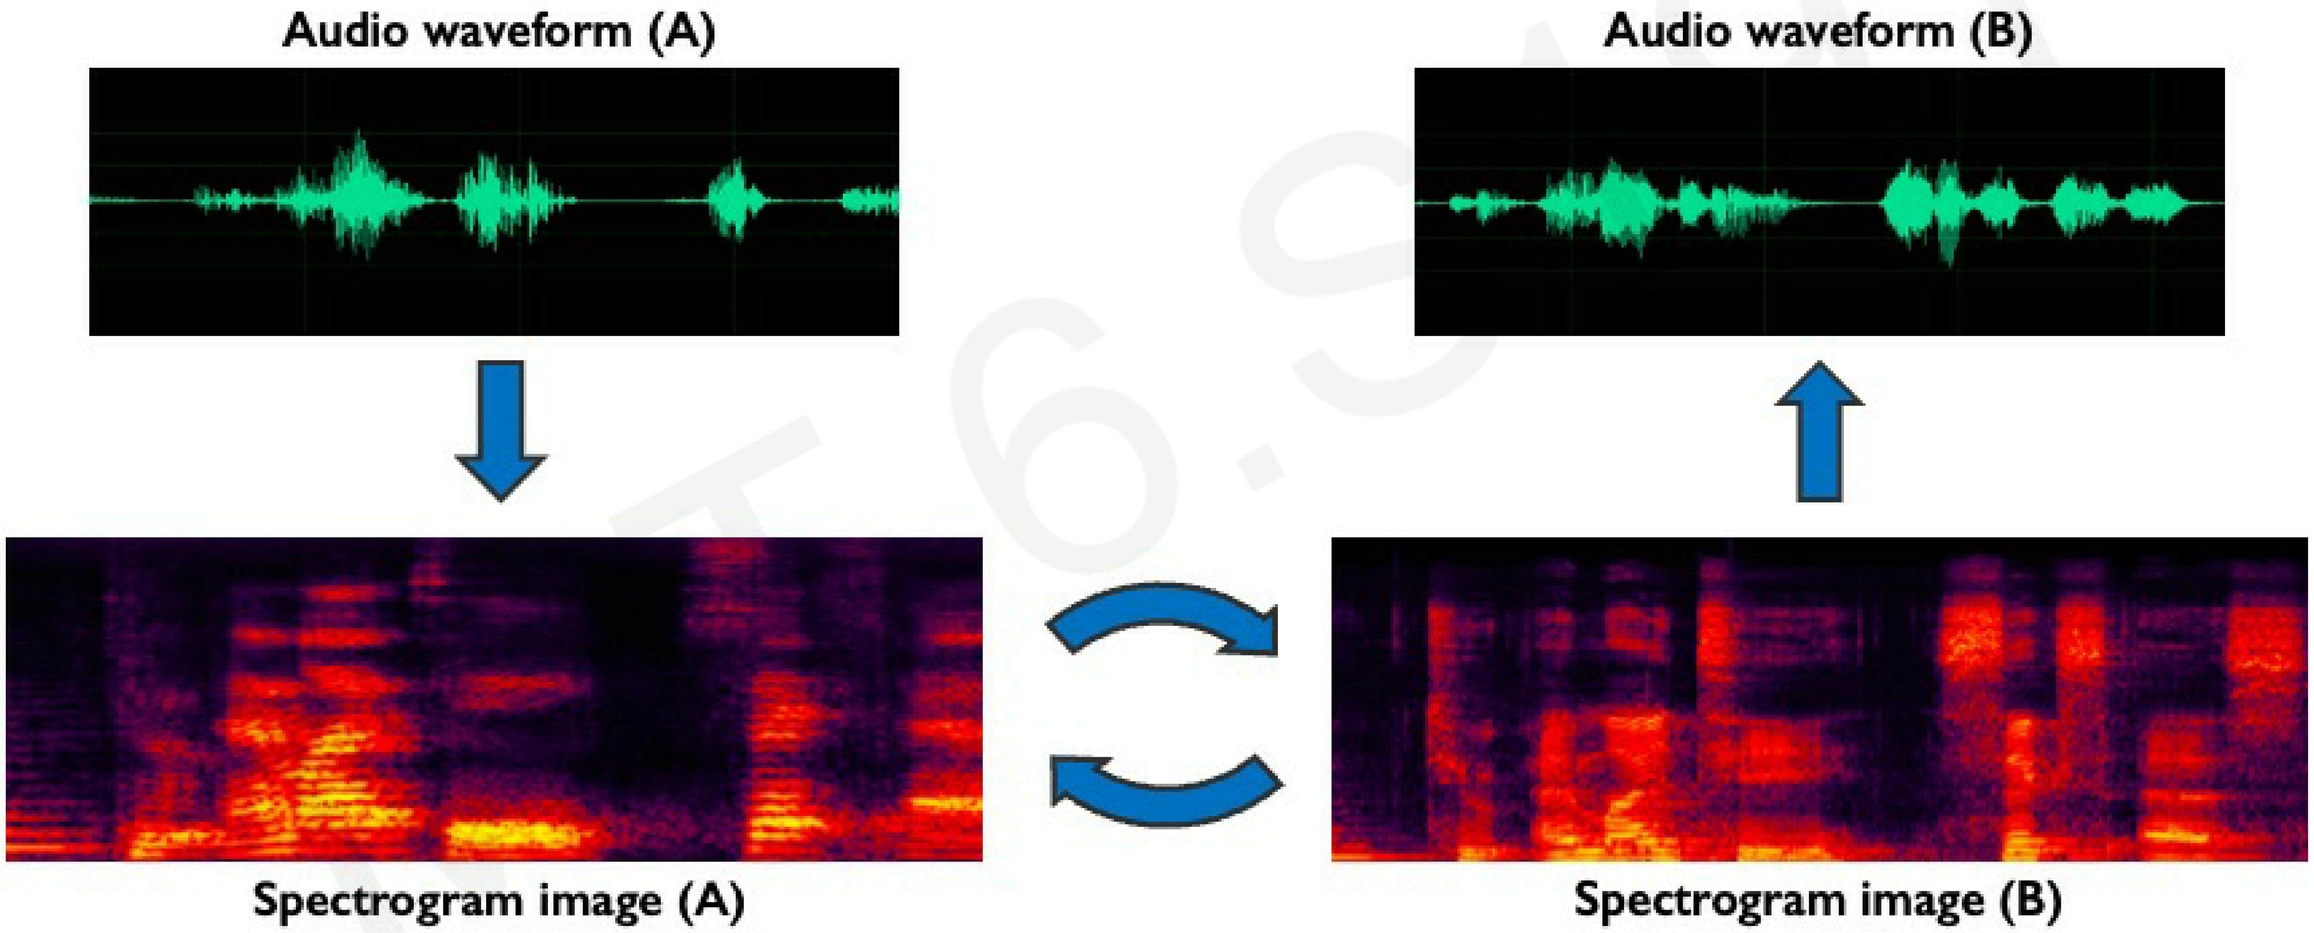

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

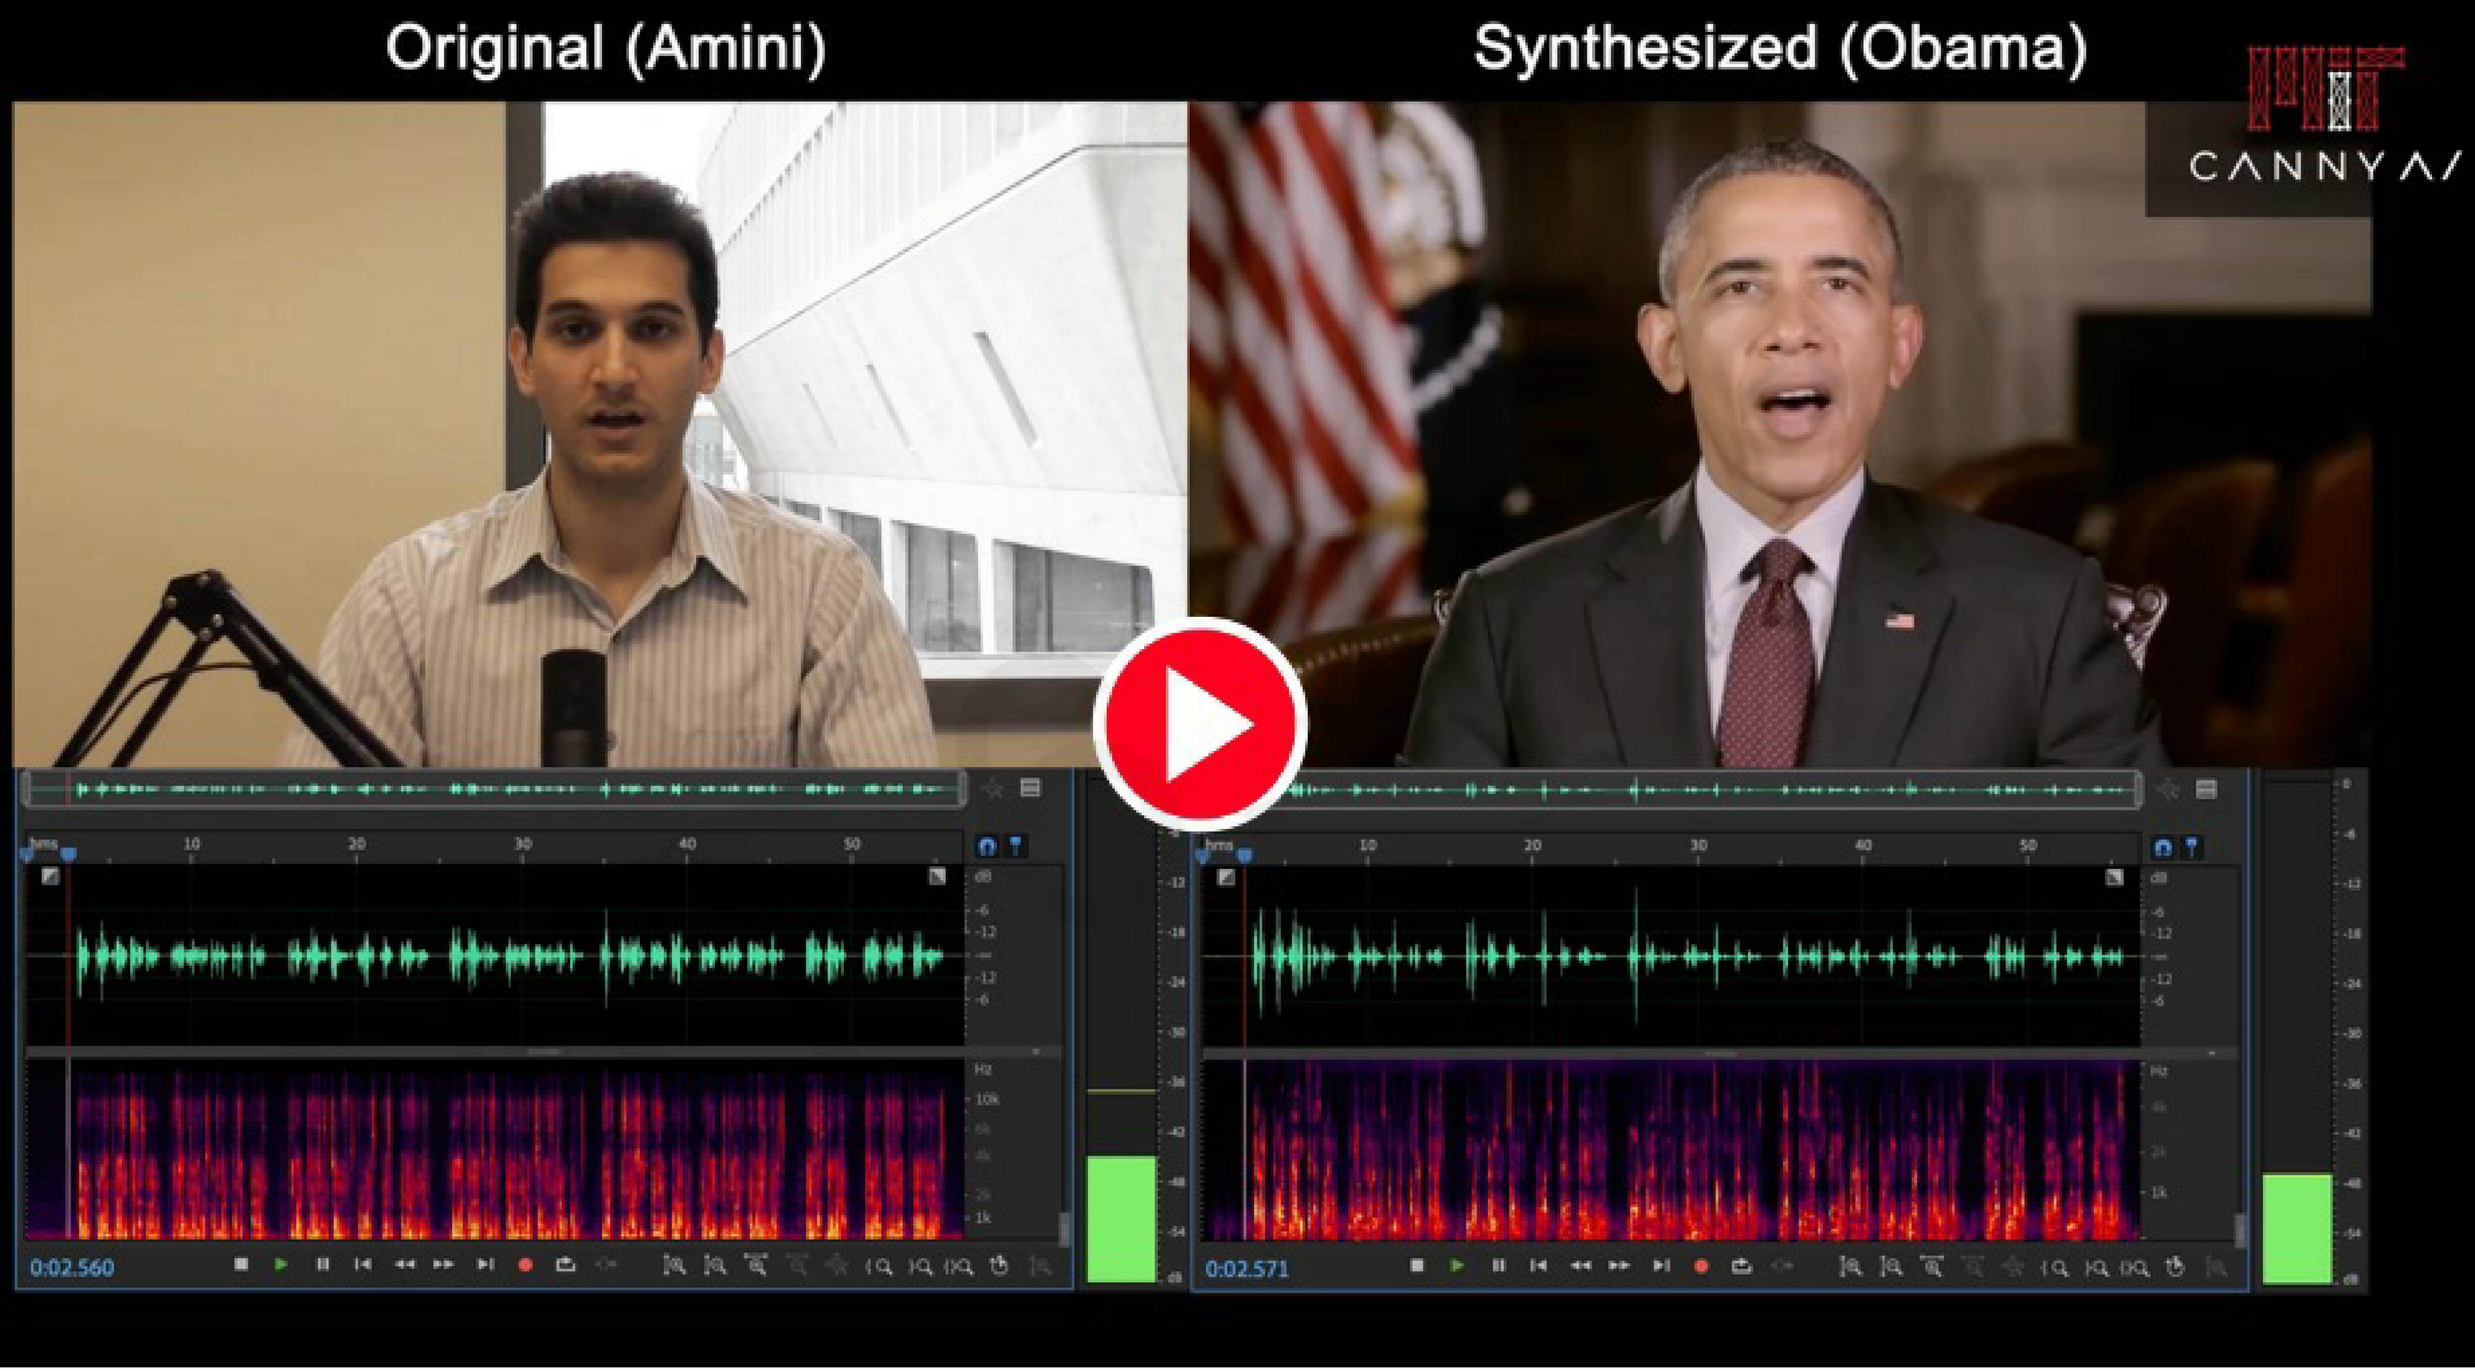

© Alexander Amini and Ava Amini MIT Introduction to Deep Learning IntroToDeepLearning.com

Deep Generative Modeling: Summary

Autoencoders and Variational Autoencoders (VAEs)

Learn lower-dimensional **latent space** and **sample** to generate input reconstructions

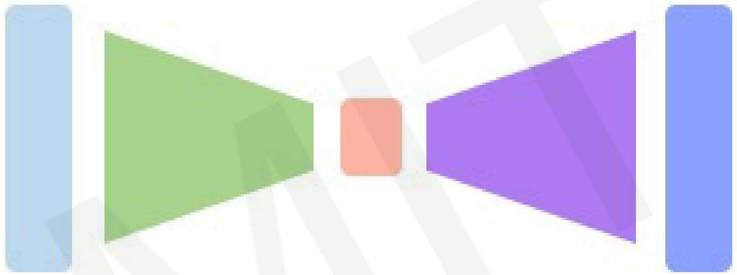

Generative Adversarial Networks (GANs)

Competing **generator** and **discriminator** networks

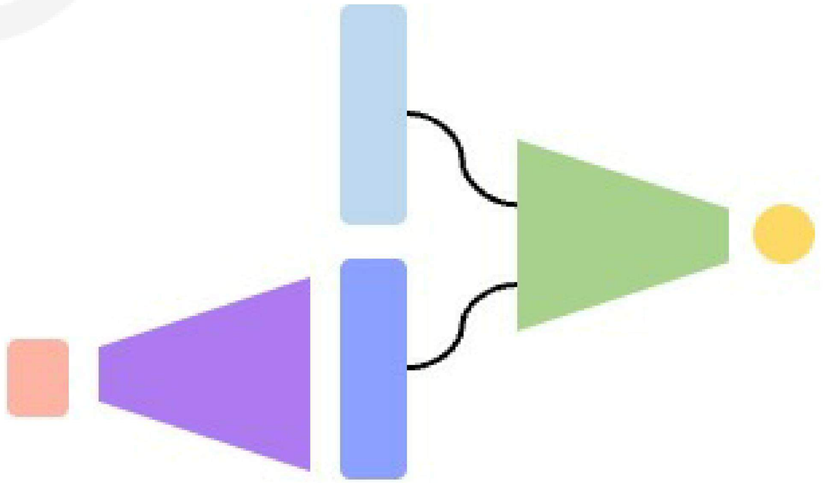

MIT 6.S191

Thank you !!!