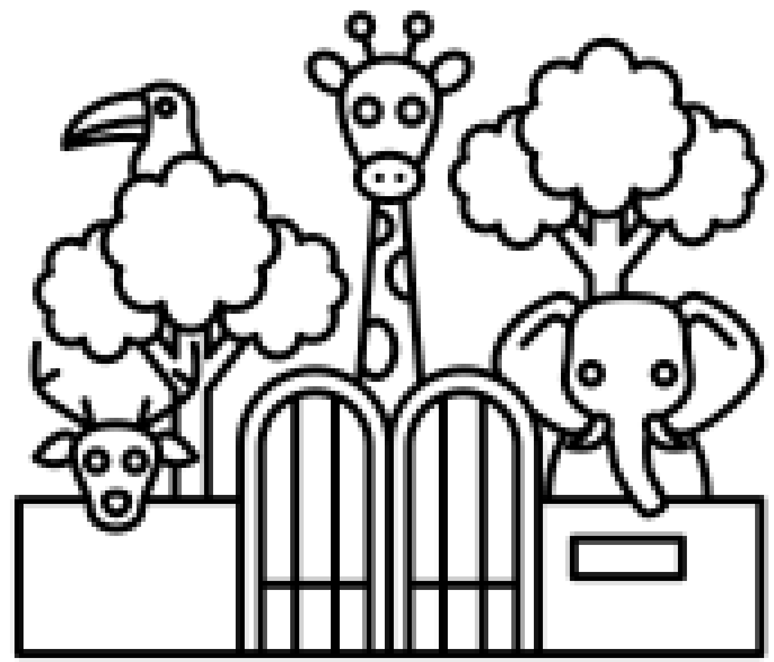

1

Generative Model Zoo (part I)

Jun-Yan Zhu

16-726 Learning-based Image Synthesis, Spring 2025

many slides from Phillip Isola, Kaiming He, Richard Zhang, Alyosha Efros

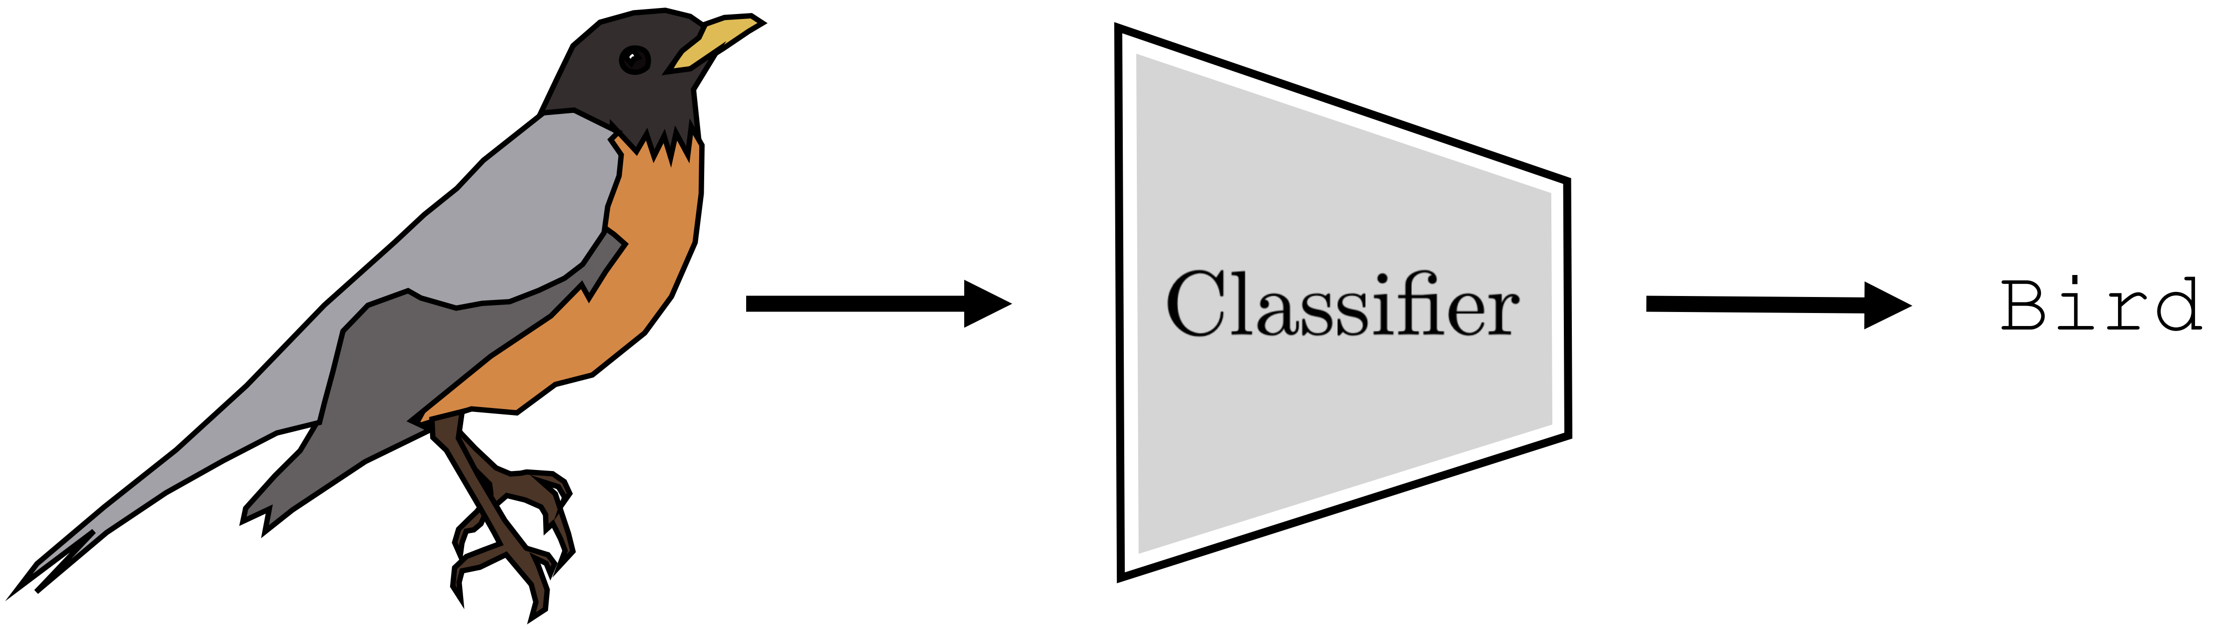

Slides credit: Phillip Isola

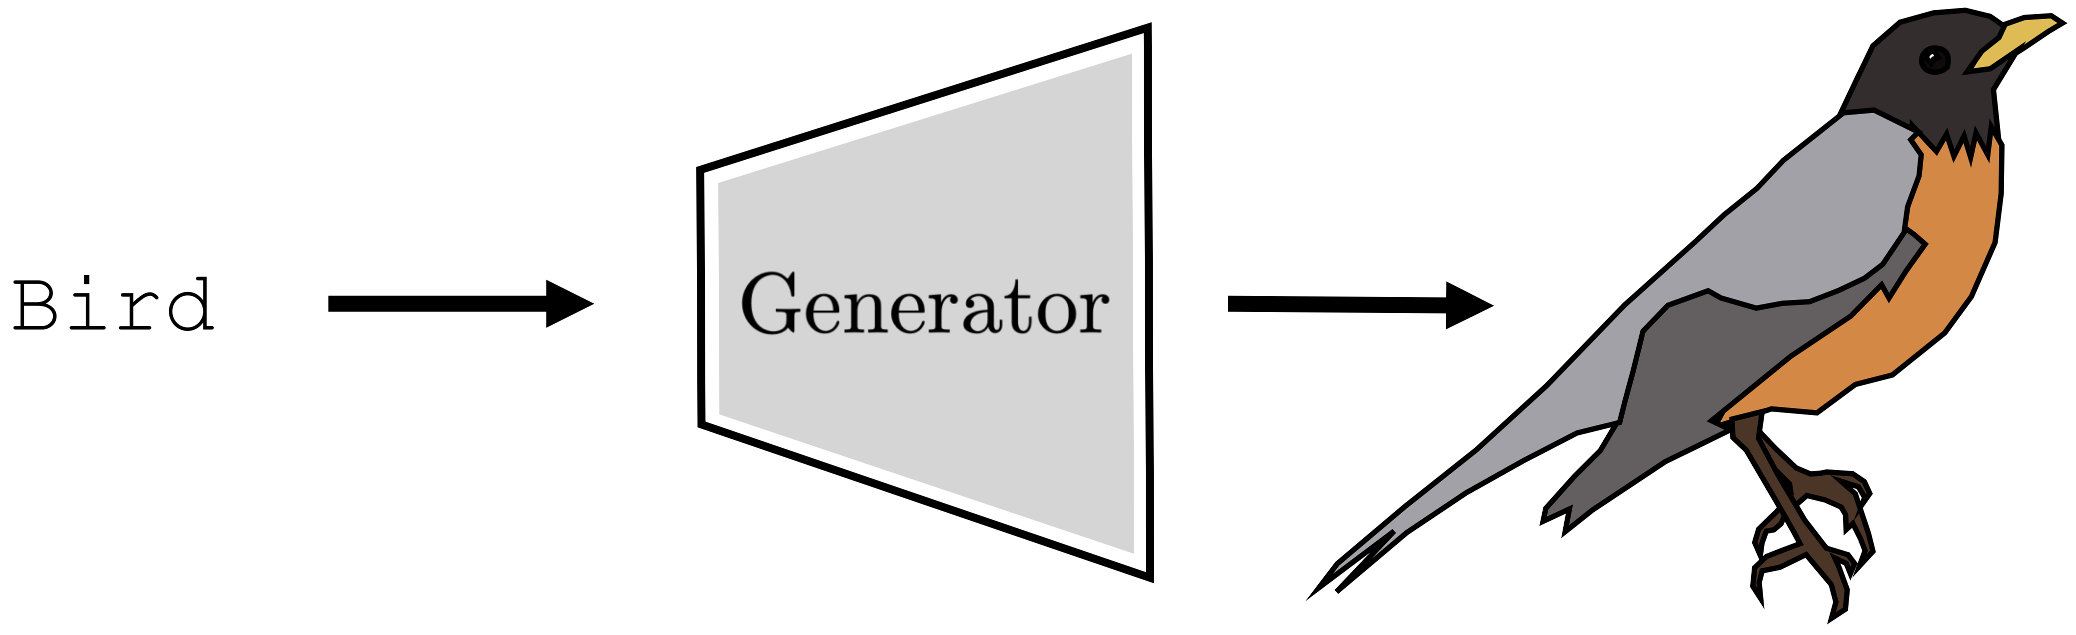

Slides credit: Phillip Isola

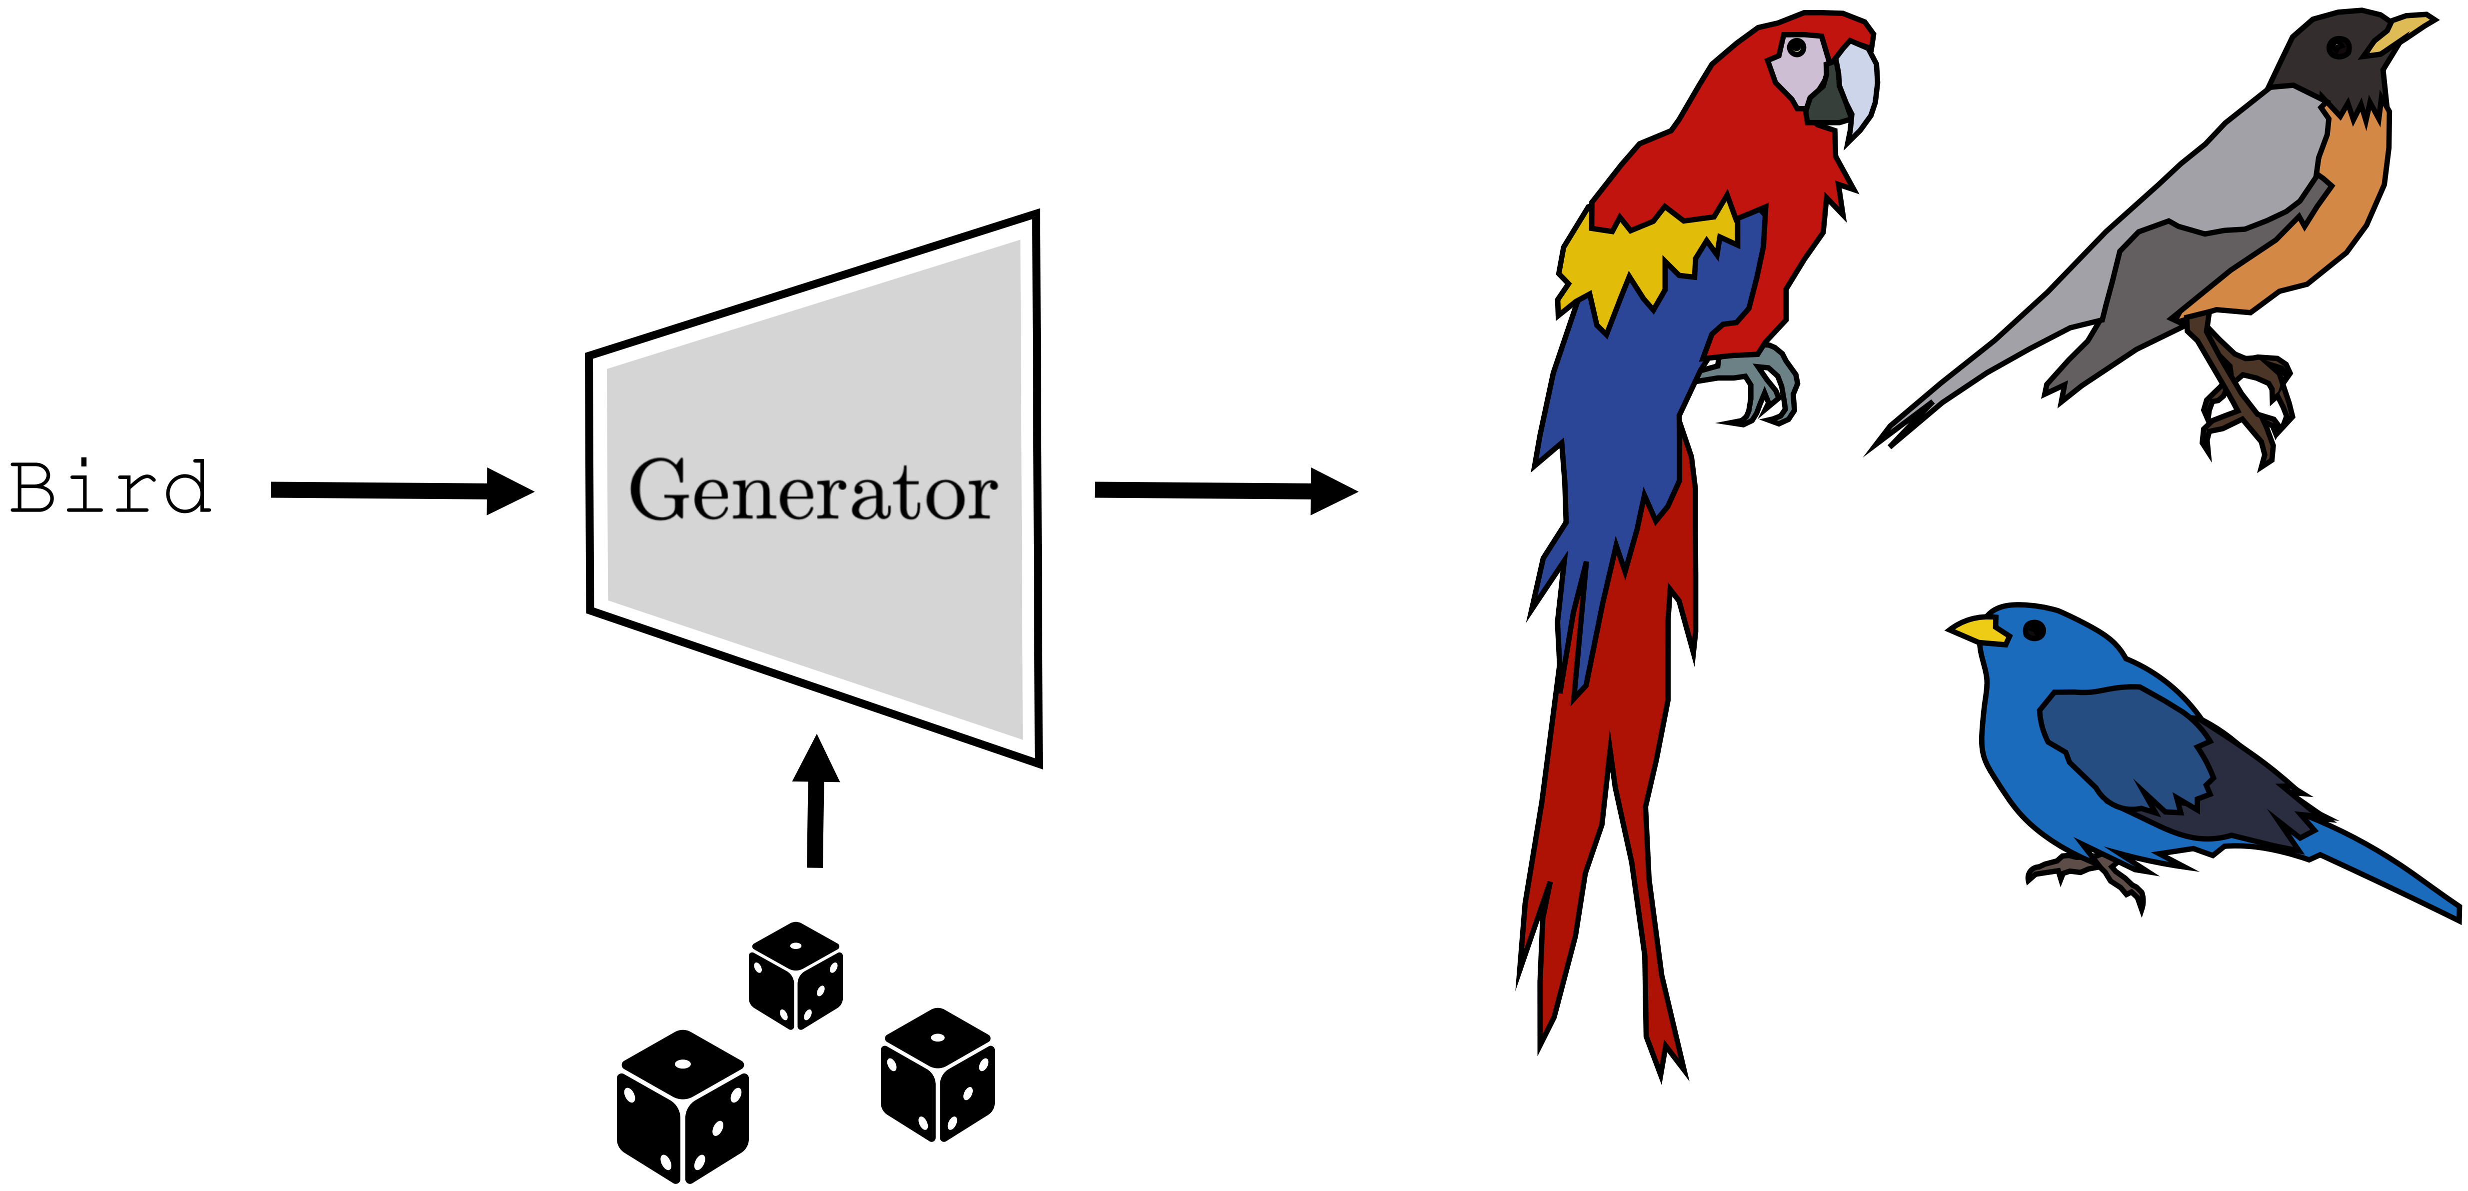

Slides credit: Phillip Isola

which color? 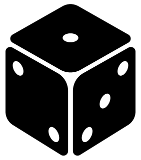

what angle? 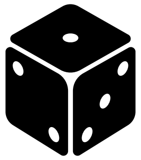

what size? 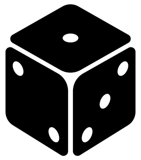

… …

Slides credit: Phillip Isola

which color?

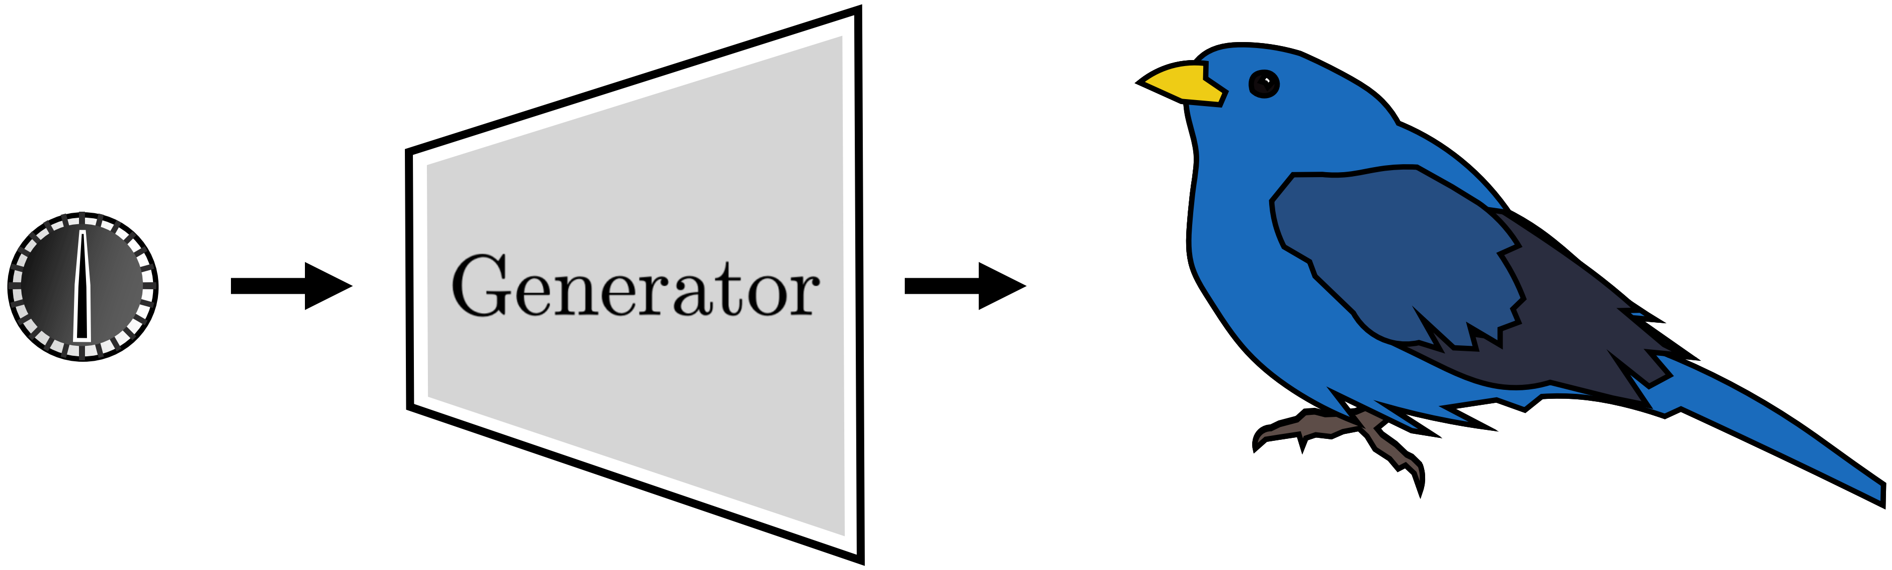

Slides credit: Phillip Isola

What’s the goal of generative modeling?

Make synthetic data that “looks like” real data.

How to measure “looks like”?

The main answer in deep generative models is: “has high probability under a density model fit to real data.”

Slides credit: Phillip Isola

What's the goal of generative modeling?

The goal is not to replicate the training data but to make ***new*** data that is ***realistic*** (captures the essential properties of real data)

(A model that memorizes the training data is overfit in exactly the same sense as a classifier can be overfit)

Slides credit: Phillip Isola

Learning data generators

Two approaches:

1. **Explicit Density Model** (indirect approach**):** learn a function that scores data; generate data by finding points that score highly under this function $E : \mathcal{X} \rightarrow \mathbb{R}$

2. **Implicit Generative Models** (Direct approach**):** learn a function that generates data directly $G : \mathcal{Z} \rightarrow \mathcal{X}$

Slides credit: Phillip Isola

Explicit Density Models

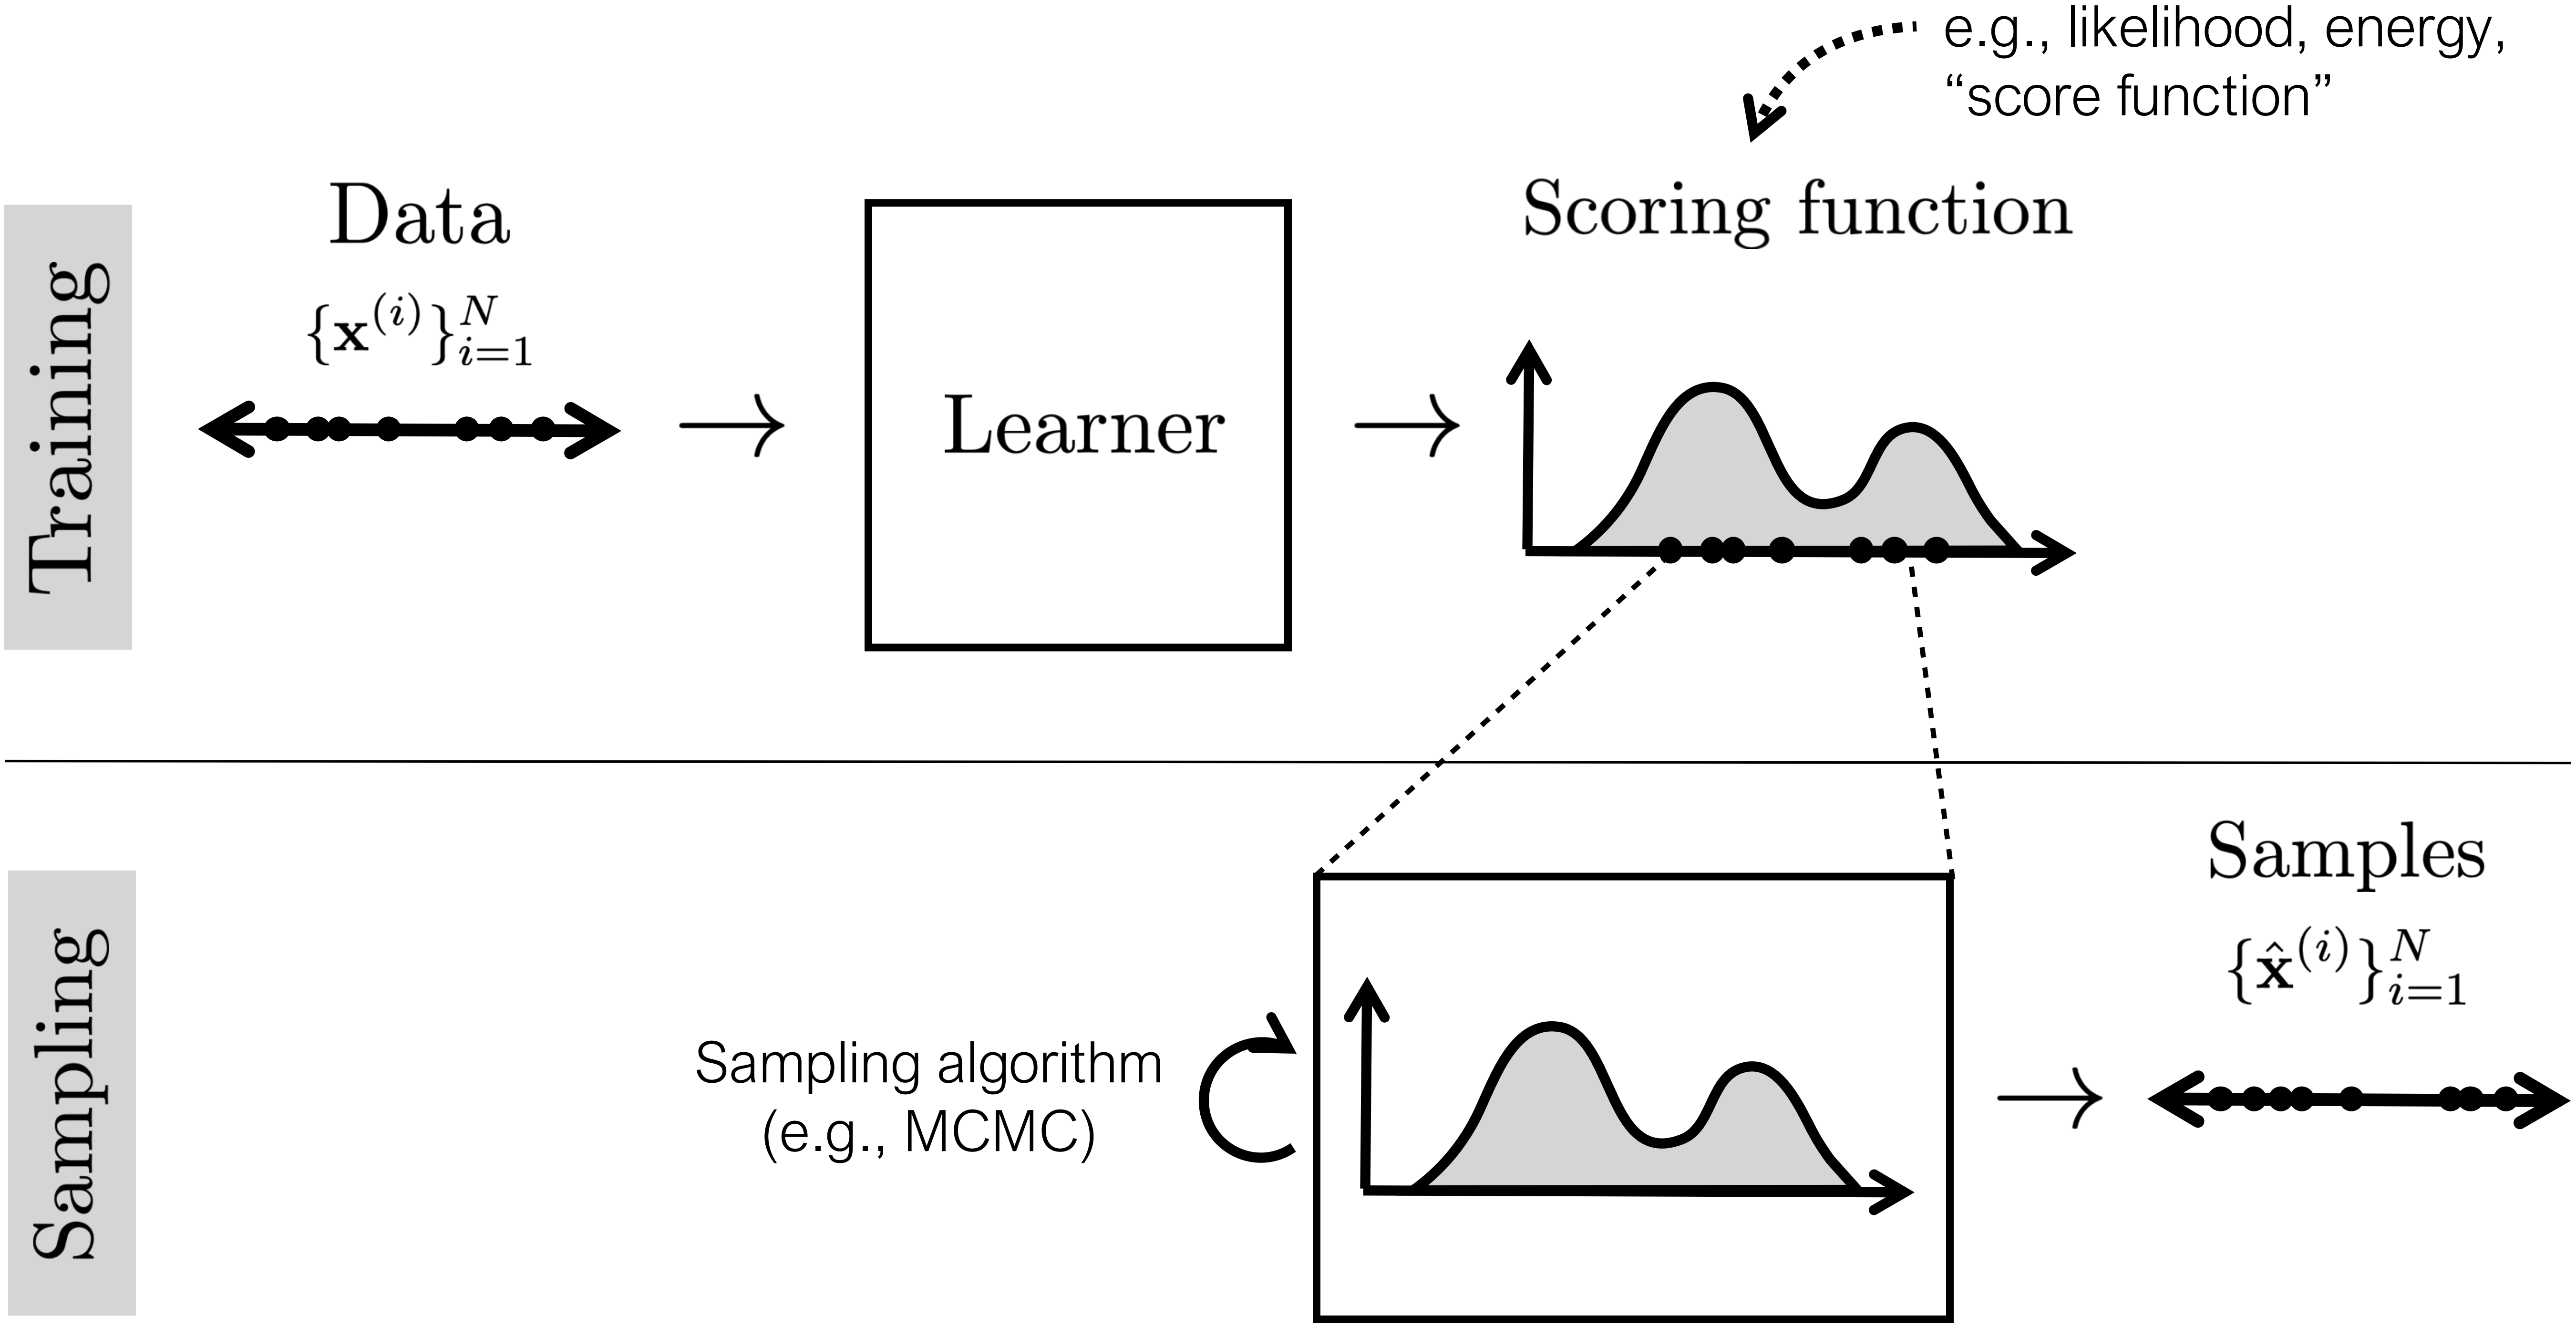

Slides credit: Phillip Isola

Explicit Density Models

$$p_\theta : \mathcal{X} \rightarrow [0, \infty)$$

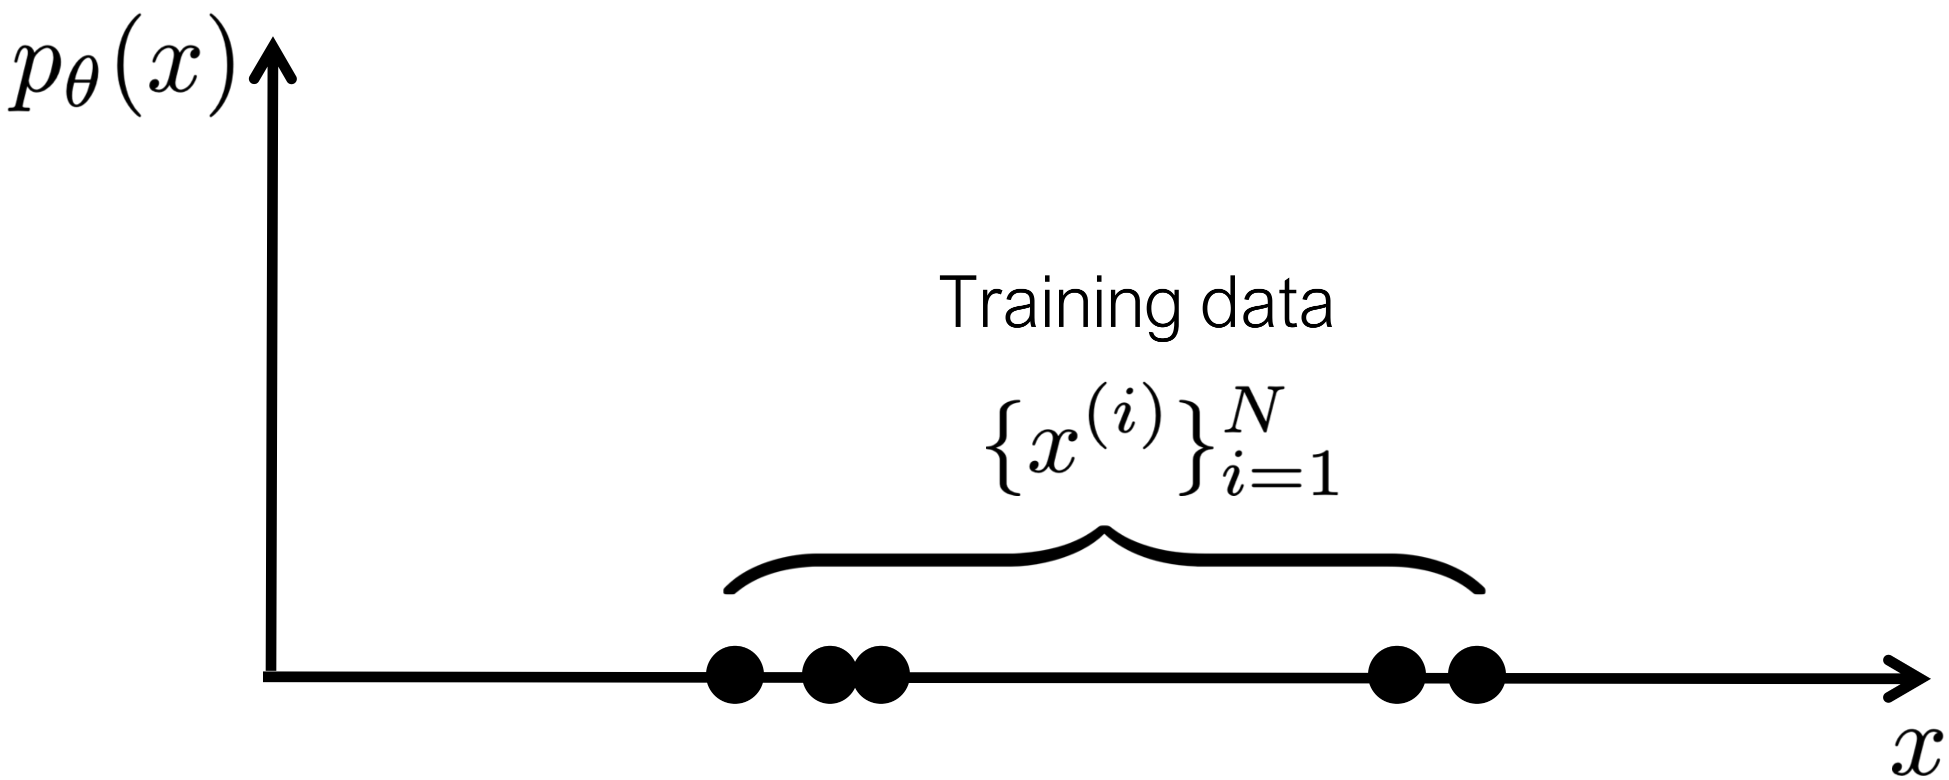

Slides credit: Phillip Isola

Explicit Density Models

$$p_\theta : \mathcal{X} \rightarrow [0, \infty)$$

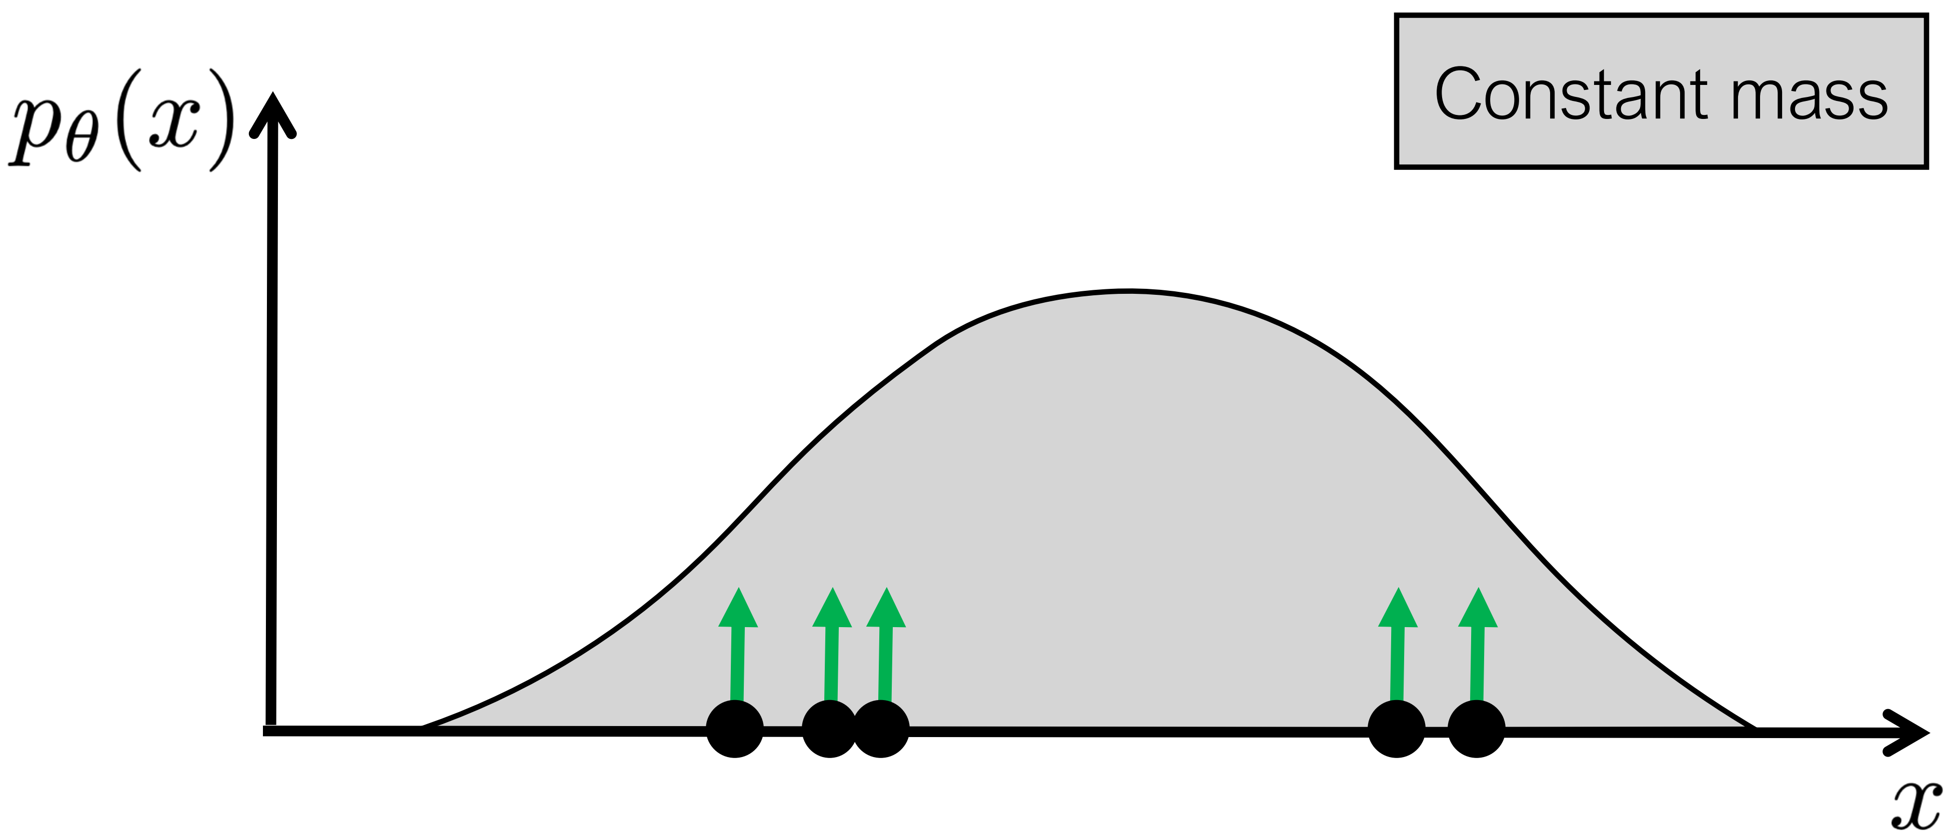

$$\int_x p_\theta(x)dx = 1$$

Slides credit: Phillip Isola

Explicit Density Models

$$p_\theta : \mathcal{X} \rightarrow [0, \infty)$$

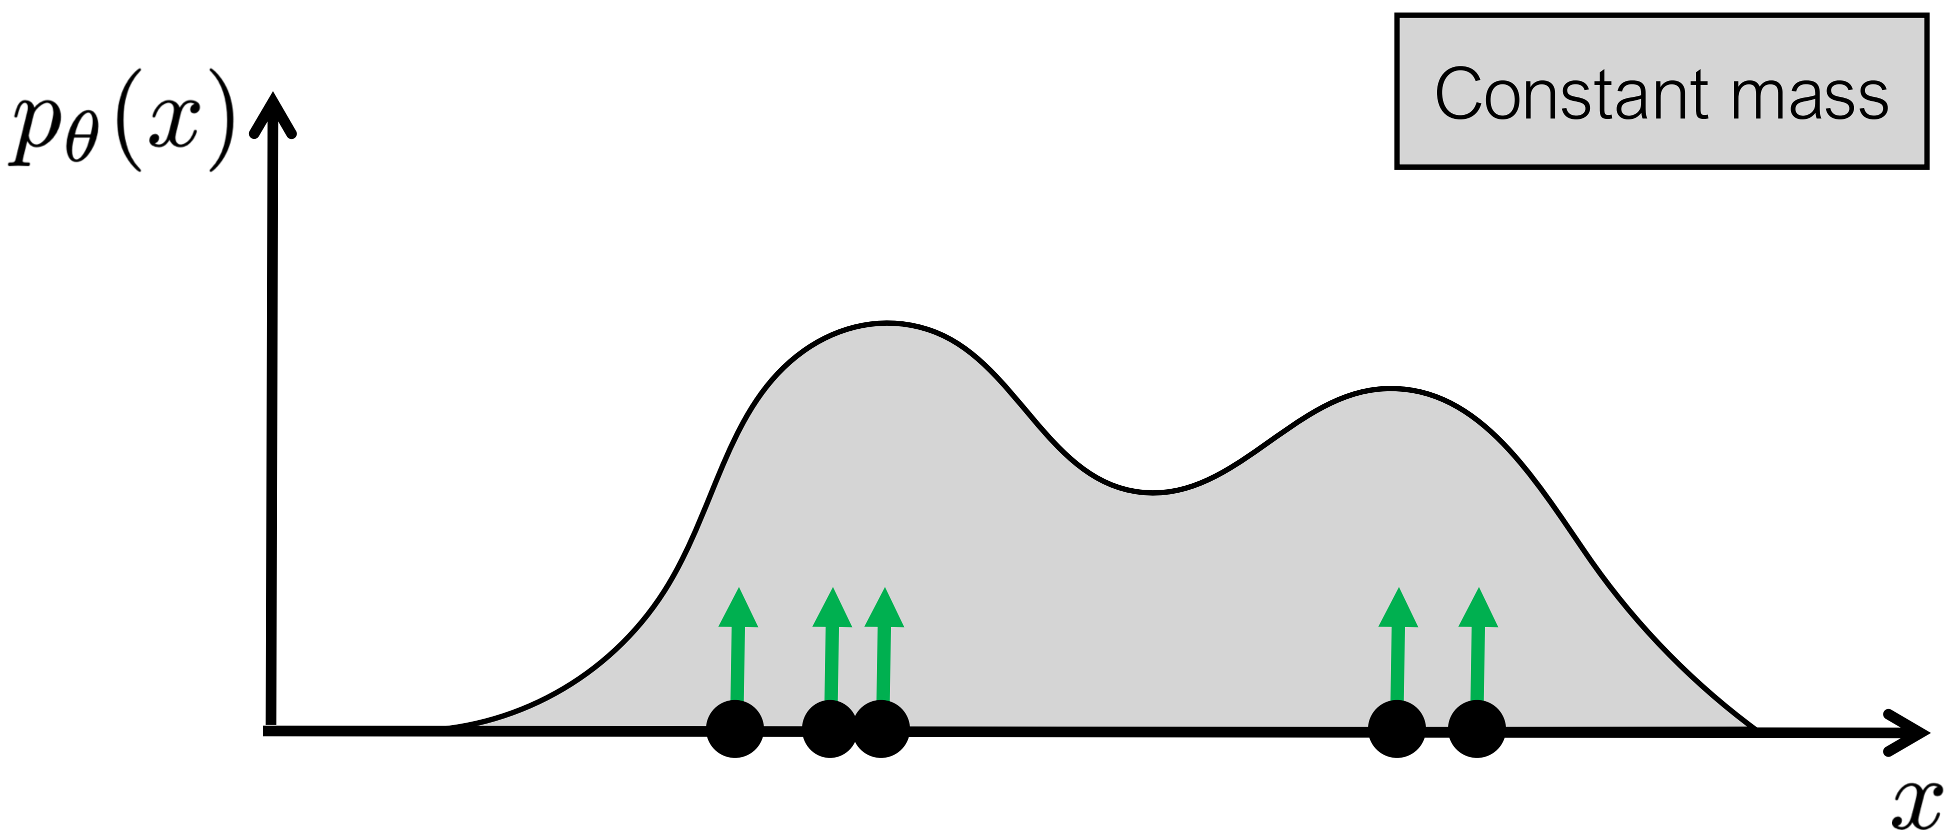

$$\int_x p_\theta(x)dx = 1$$

Slides credit: Phillip Isola

Explicit Density Models

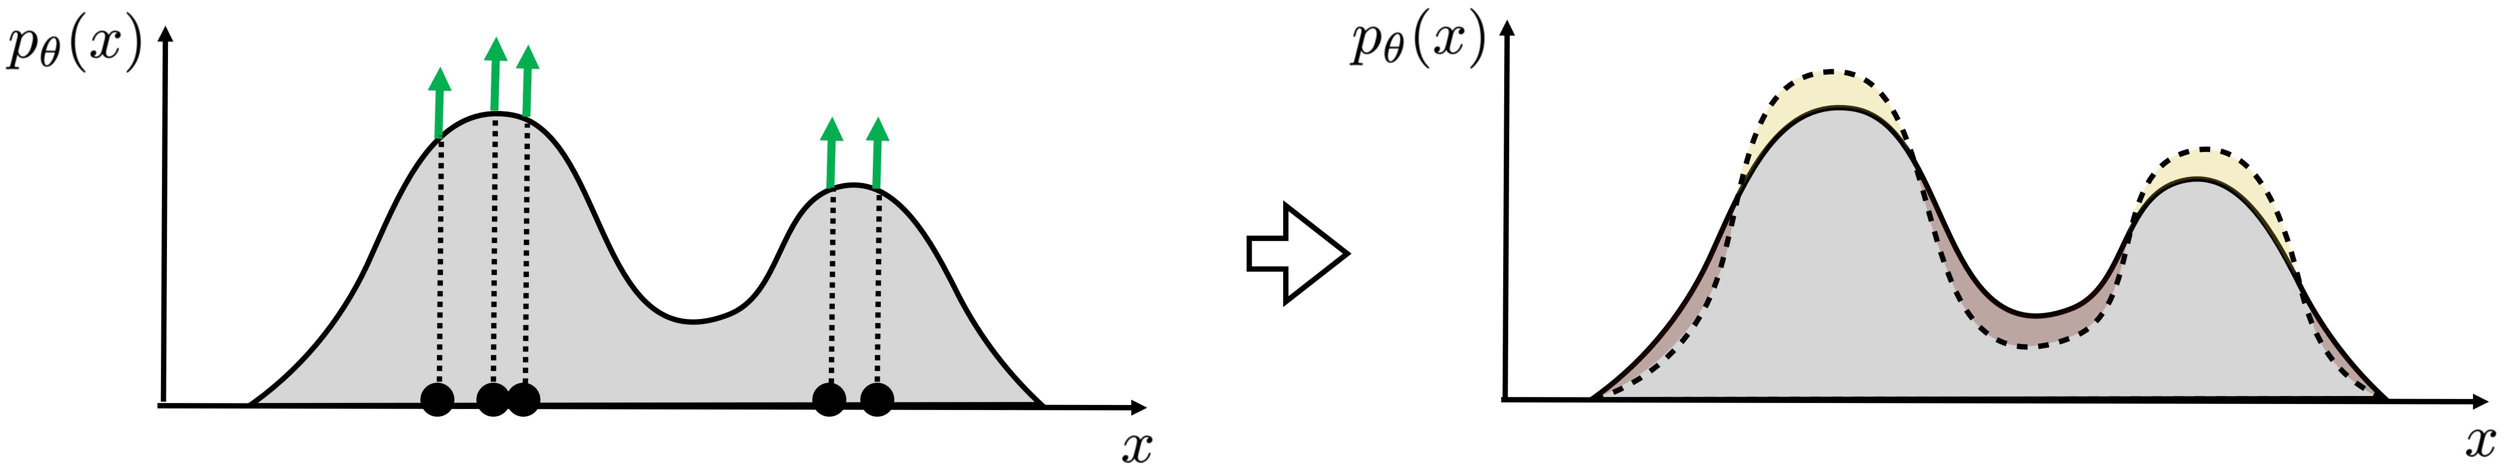

$$p^*_\theta = \arg\min_{p_\theta} \mathtt{KL}(p_{\mathtt{data}}, p_\theta)$$

$$= \arg\min_{p_\theta} \mathbb{E}_{\mathbf{x}\sim p_{\mathtt{data}}}\left[-\log \frac{p_\theta(\mathbf{x})}{p_{\mathtt{data}}(\mathbf{x})}\right]$$

$$= \arg\max_{p_\theta} \mathbb{E}_{\mathbf{x}\sim p_{\mathtt{data}}}\left[\log p_\theta(\mathbf{x})\right] - \mathbb{E}_{\mathbf{x}\sim p_{\mathtt{data}}}\left[\log p_{\mathtt{data}}(\mathbf{x})\right]$$

$$= \arg\max_{p_\theta} \mathbb{E}_{\mathbf{x}\sim p_{\mathtt{data}}}\left[\log p_\theta(\mathbf{x})\right] \quad \triangleleft \quad \text{dropped second term since no dependence on } p_\theta$$

$$\approx \arg\max_{p_\theta} \frac{1}{N}\sum_{i=1}^{N} \log p_\theta(\mathbf{x}^{(i)})$$

Slides credit: Phillip Isola

Energy-based models

i.e. unnormalized probability models

$\int_{\mathbf{x}} p_\theta(\mathbf{x}) d\mathbf{x} = 1$

$$p_\theta = \frac{e^{-E_\theta}}{Z(\theta)} \qquad Z(\theta) = \int_{\mathbf{x}} e^{-E_\theta(\mathbf{x})} d\mathbf{x}$$

$$\frac{p_\theta(\mathbf{x}_1)}{p_\theta(\mathbf{x}_2)} = \frac{e^{-E_\theta(\mathbf{x}_1)}/Z(\theta)}{e^{-E_\theta(\mathbf{x}_2)}/Z(\theta)} = \frac{e^{-E_\theta(\mathbf{x}_1)}}{e^{-E_\theta(\mathbf{x}_2)}}$$

<— Relative probabilities are often all you need (e.g., for sampling)

Nobel Prize in Physics, 2024

Hopfield Networks [Hopfield et al., 1982], Restricted Boltzmann Machines [Ackley et al., 1985]

Energy-based models

At convergence, green (data) and red (model) samples are identical and model update (green-red) cancels out

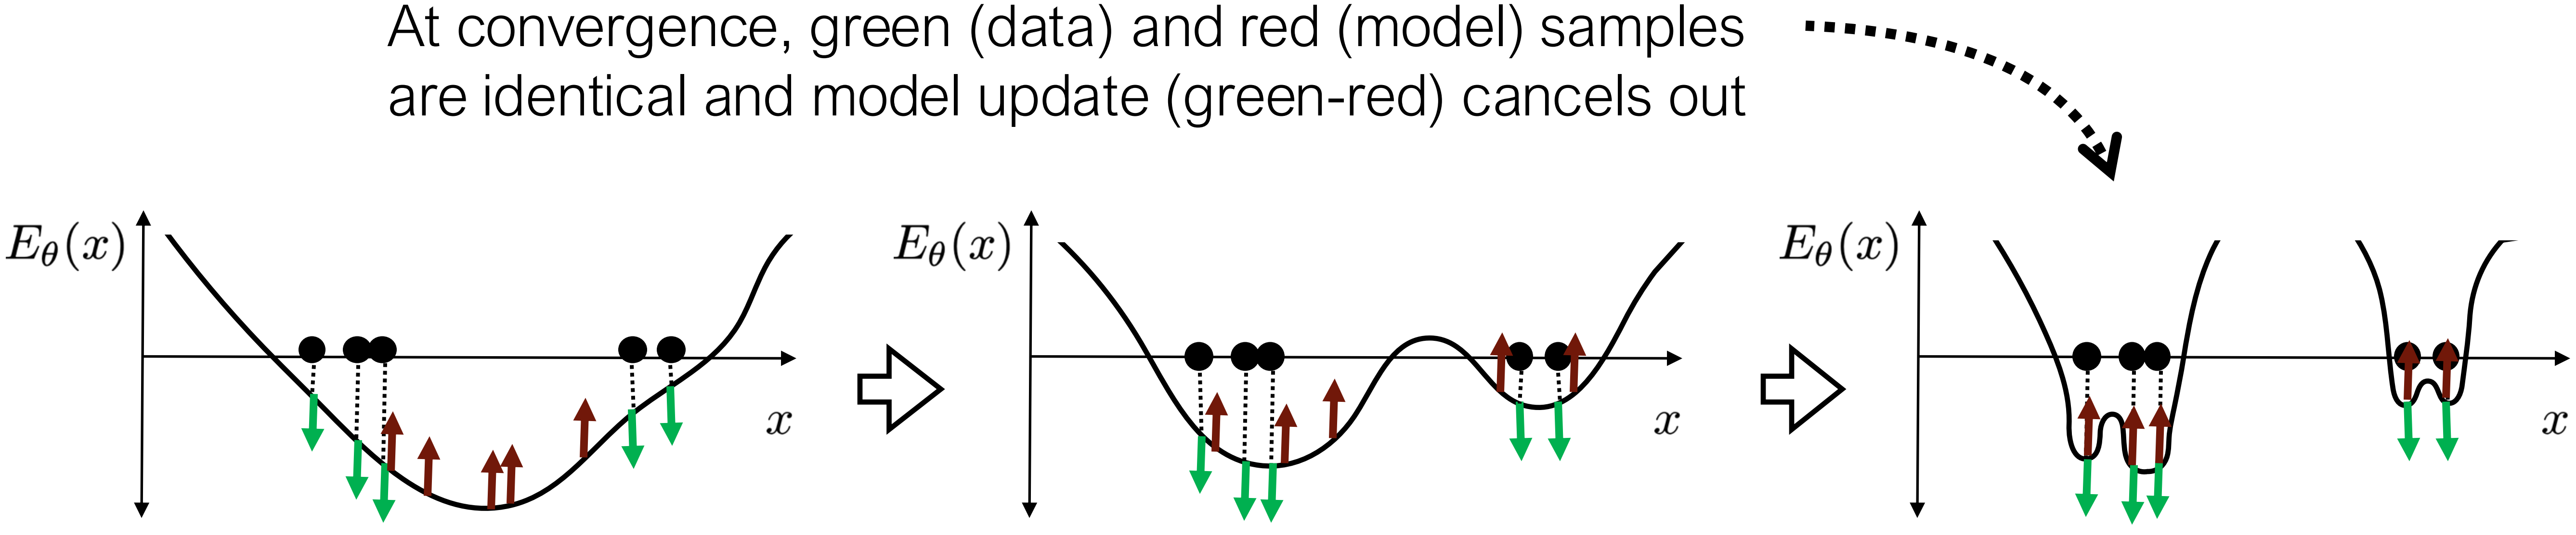

Slides credit: Phillip Isola

Energy-based models — learning algorithm

$$\nabla_\theta \mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\log p_\theta(\mathbf{x})] = \nabla_\theta \mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\log \frac{e^{-E_\theta(\mathbf{x})}}{Z(\theta)}]$$

$$= -\mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\nabla_\theta E_\theta(\mathbf{x})] - \underbrace{\nabla_\theta \log Z(\theta)}$$

How to measure this?

Slides credit: Phillip Isola

Energy-based models — learning algorithm

$$
\begin{aligned}
\boxed{-\nabla_\theta \log Z(\theta)} &= \frac{1}{Z(\theta)} \nabla_\theta Z(\theta) && \triangleleft \quad \nabla_x \log f(x) = \tfrac{1}{f(x)} \nabla_x f(x) \\
&= \frac{1}{Z(\theta)} \nabla_\theta \int_x e^{-E_\theta(\mathbf{x})} d\mathbf{x} && \triangleleft \quad \text{definition of } Z \\
&= \frac{1}{Z(\theta)} \int_x \nabla_\theta e^{-E_\theta(\mathbf{x})} d\mathbf{x} && \triangleleft \quad \text{exchange sum and grad} \\
&= \frac{1}{Z(\theta)} - \int_x e^{-E_\theta(\mathbf{x})} \nabla_\theta E_\theta(\mathbf{x}) d\mathbf{x} \\
&= -\int_x \frac{e^{-E_\theta(\mathbf{x})}}{Z(\theta)} \nabla_\theta E_\theta(\mathbf{x}) d\mathbf{x} \\
&= -\int_x p_\theta(\mathbf{x}) \nabla_\theta E_\theta(\mathbf{x}) d\mathbf{x} && \triangleleft \quad \text{definition of } p_\theta \\
&= -\mathbb{E}_{\mathbf{x} \sim p_\theta} [\nabla_\theta E_\theta(\mathbf{x})] && \triangleleft \quad \text{definition of expectation}
\end{aligned}
$$

Slides credit: Phillip Isola

Energy-based models — learning algorithm

$$\nabla_\theta \mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\log p_\theta(\mathbf{x})] = \nabla_\theta \mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\log \frac{e^{-E_\theta(\mathbf{x})}}{Z(\theta)}]$$

$$= -\mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\nabla_\theta E_\theta(\mathbf{x})] - \nabla_\theta \log Z(\theta)$$

$$= -\mathbb{E}_{\mathbf{x}\sim p_{\texttt{data}}}[\nabla_\theta E_{\theta(\mathbf{x})}] + \mathbb{E}_{\mathbf{x}\sim p_\theta}[\nabla_\theta E_\theta(\mathbf{x})]$$

$$\approx -\frac{1}{N}\sum_{i=1}^{N} \nabla_\theta E_\theta(\mathbf{x}^{(i)}) + \frac{1}{N}\sum_{i=1}^{N} \nabla_\theta E_\theta(\hat{\mathbf{x}}^{(i)})$$

$\mathbf{x}^{(i)} \sim p_{\texttt{data}}$ $\hat{\mathbf{x}}^{(i)} \sim p_\theta$

Slides credit: Phillip Isola

Energy-based models

At convergence, green (data) and red (model) samples are identical and model update (green-red) cancels out

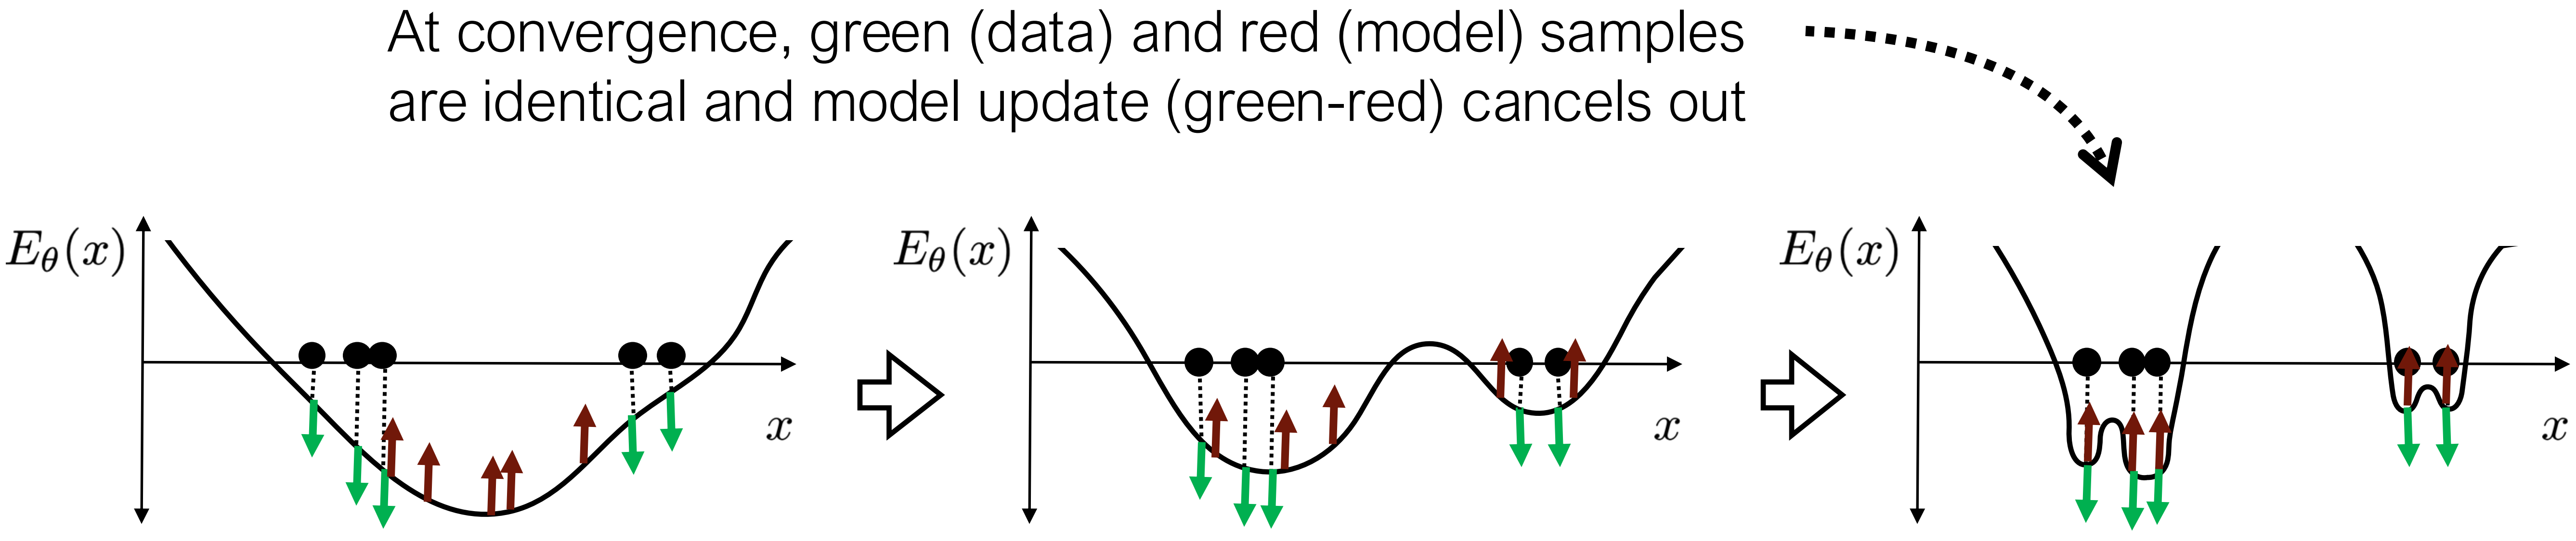

Slides credit: Phillip Isola

Learning an Explicit Density Model

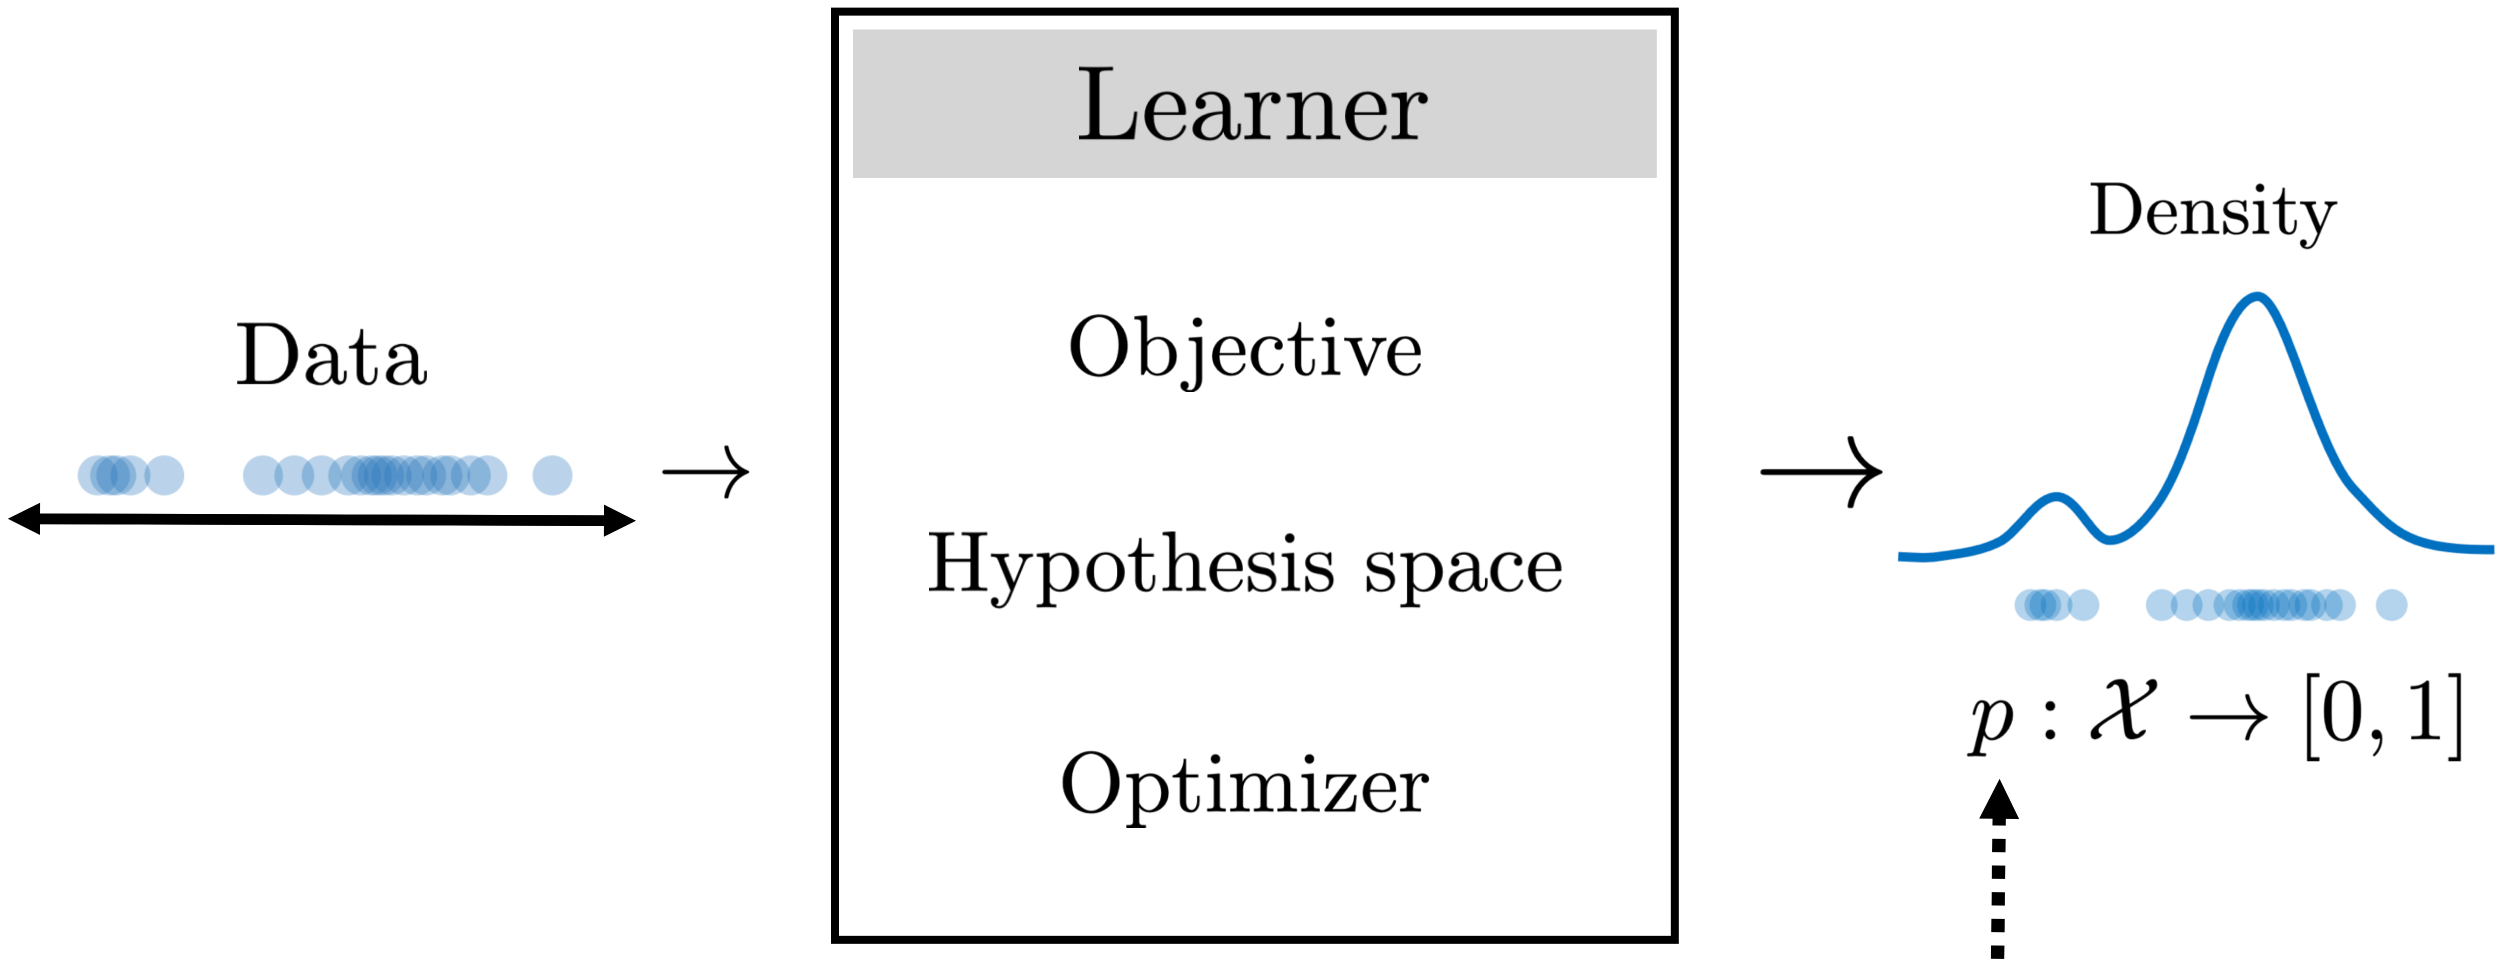

Integral of probability density function needs to be 1 → Normalized distribution

(some models output unnormalized ***energy functions***)

[figs modified from: http://introtodeeplearning.com/materials/2019_6S191_L4.pdf]

Useful for abnormality/outlier detection (detect unlikely events)

Slides credit: Phillip Isola

Implicit Generative Models

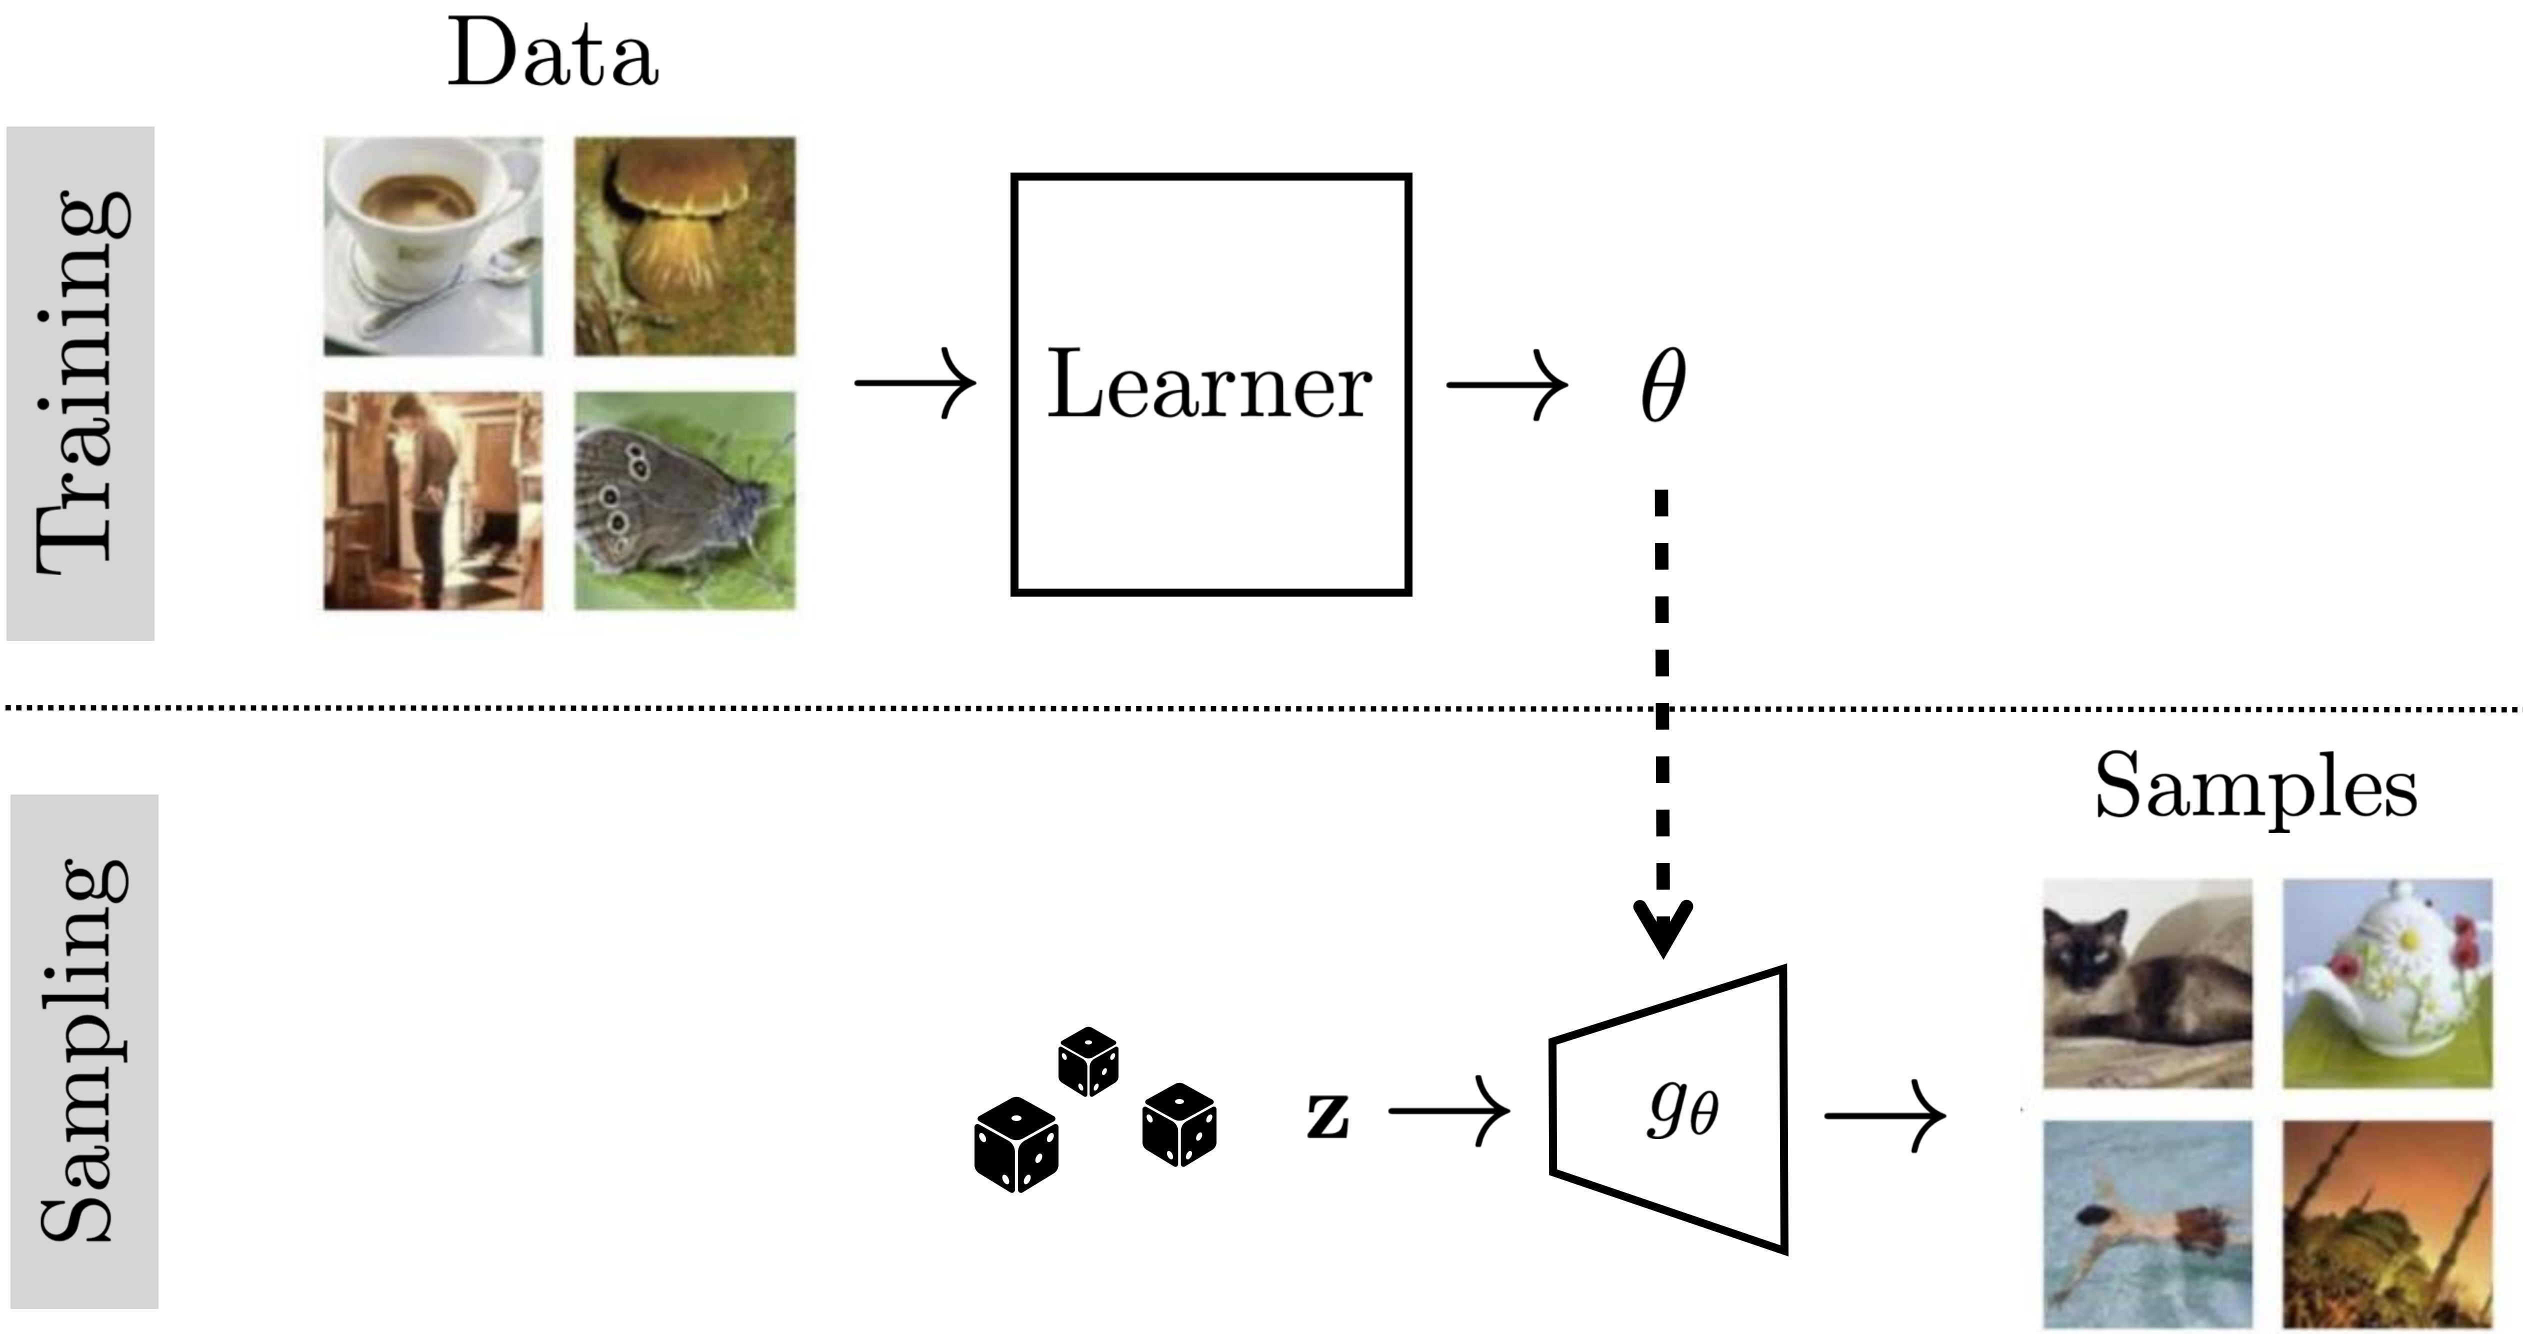

Slides credit: Phillip Isola

Data

$\{x^{(i)}\}_{i=1}^{N}$ $\rightarrow$ **Generative modeling** $\rightarrow$

Density function

$p_\theta : \mathcal{X} \rightarrow [0, \infty)$

Energy function

$E_\theta : \mathcal{X} \rightarrow \mathbb{R}$

Generator

$G_\theta : \mathcal{Z} \rightarrow \mathcal{X}$

You can represent the data generating process directly or indirectly

Slides credit: Phillip Isola