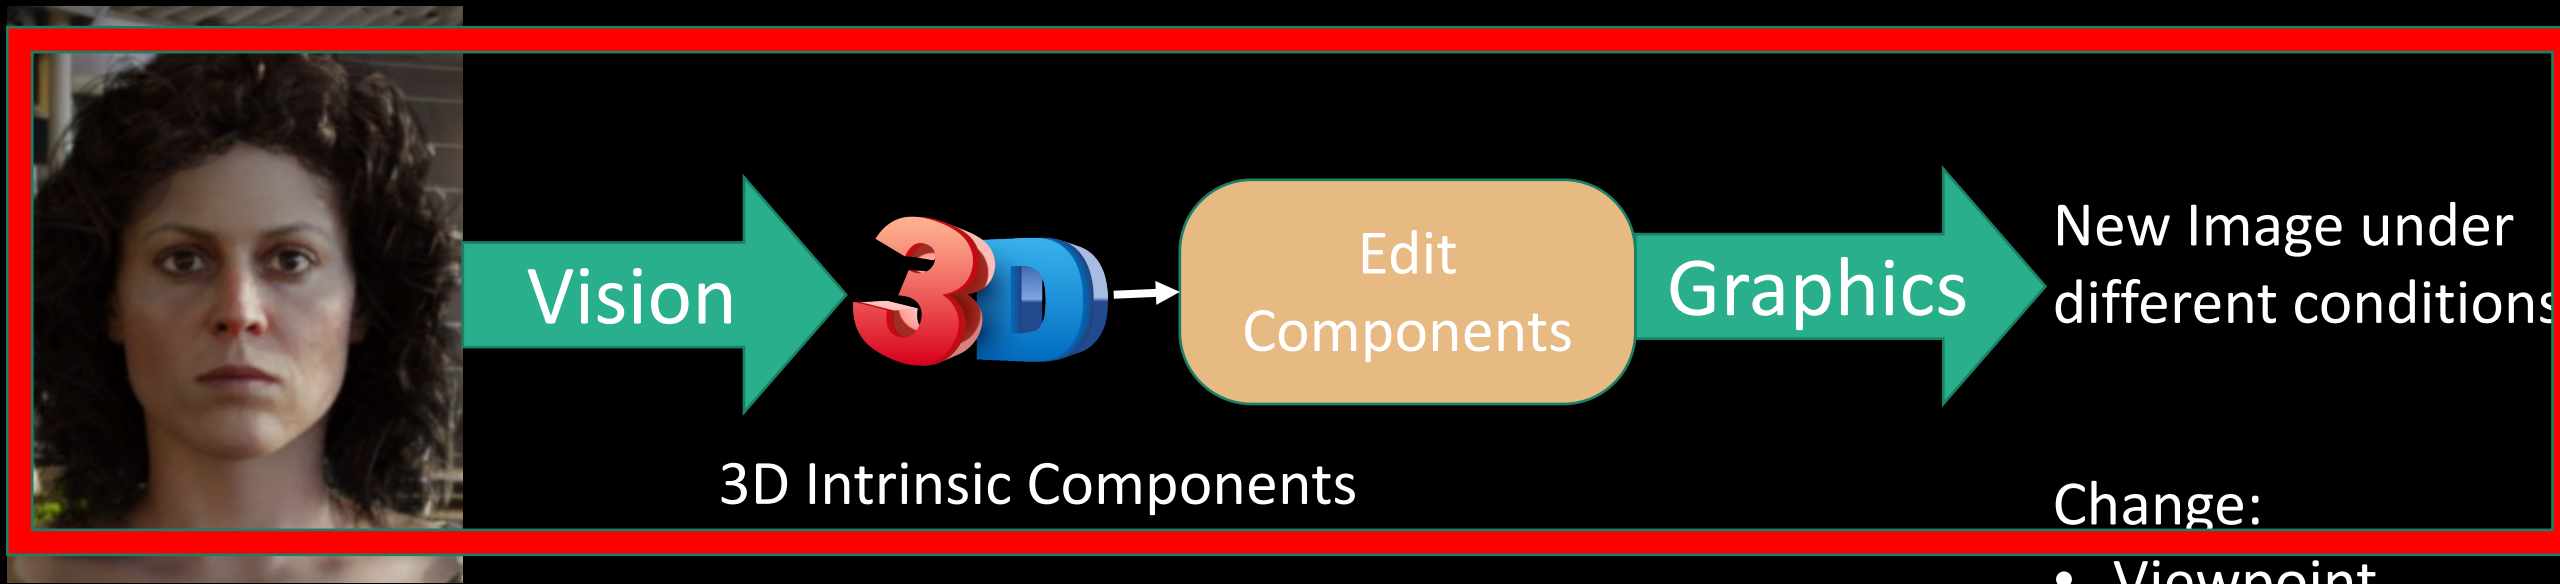


Lecture 4: Generative Models

Next few lectures: Generative models for direct image based rendering.



Current Image

Implicit: Use a Neural Network
(Conditional Generative networks)
Often, end-to-end.

- Change:
- Viewpoint
 - Lighting
 - Reflectance
 - Background
 - Attributes
 - Many others...

Supervised vs Unsupervised Learning

Supervised Learning

Data: (x, y)

x is data, y is label

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

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

Classification



Cat

Supervised vs Unsupervised Learning

Supervised Learning

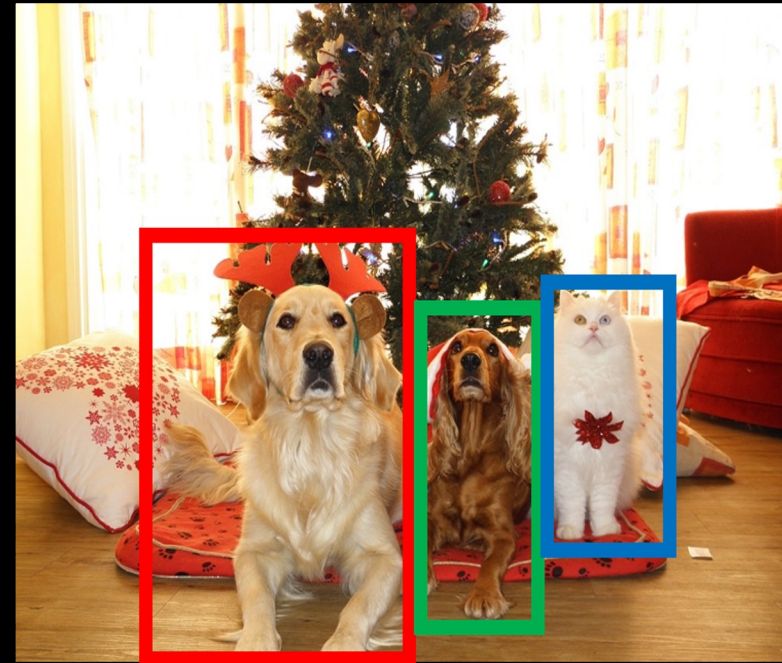
Data: (x, y)

x is data, y is label

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

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

Object Detection



DOG, DOG, CAT

Supervised vs Unsupervised Learning

Supervised Learning

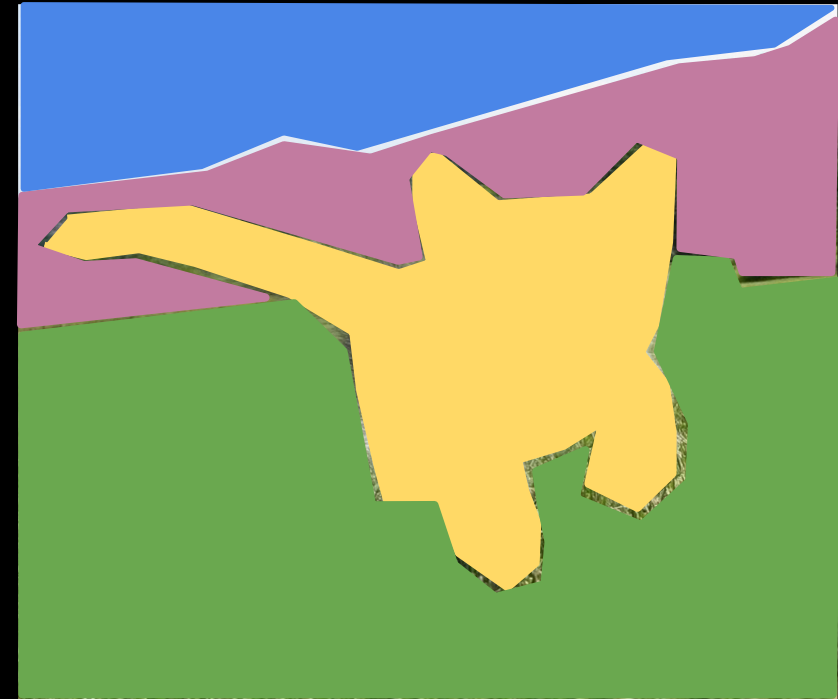
Data: (x, y)

x is data, y is label

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

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

Semantic Segmentation



GRASS, CAT, TREE, SKY

Supervised vs Unsupervised Learning

Supervised Learning

Data: (x, y)

x is data, y is label

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

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

Image captioning



A cat sitting on a suitcase on the floor

Supervised vs Unsupervised Learning

Supervised Learning

Data: (x, y)

x is data, y is label

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

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

Unsupervised Learning

Data: x

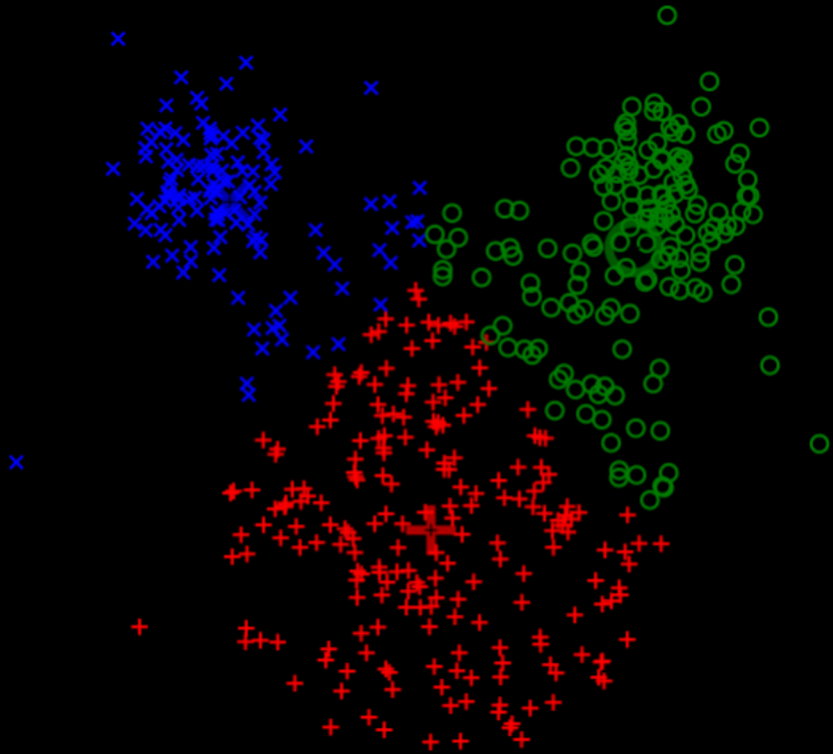
Just data, no labels!

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

Examples: Clustering, dimensionality reduction, feature learning, density estimation, etc.

Supervised vs Unsupervised Learning

Clustering
(e.g. K-Means)



Unsupervised Learning

Data: x

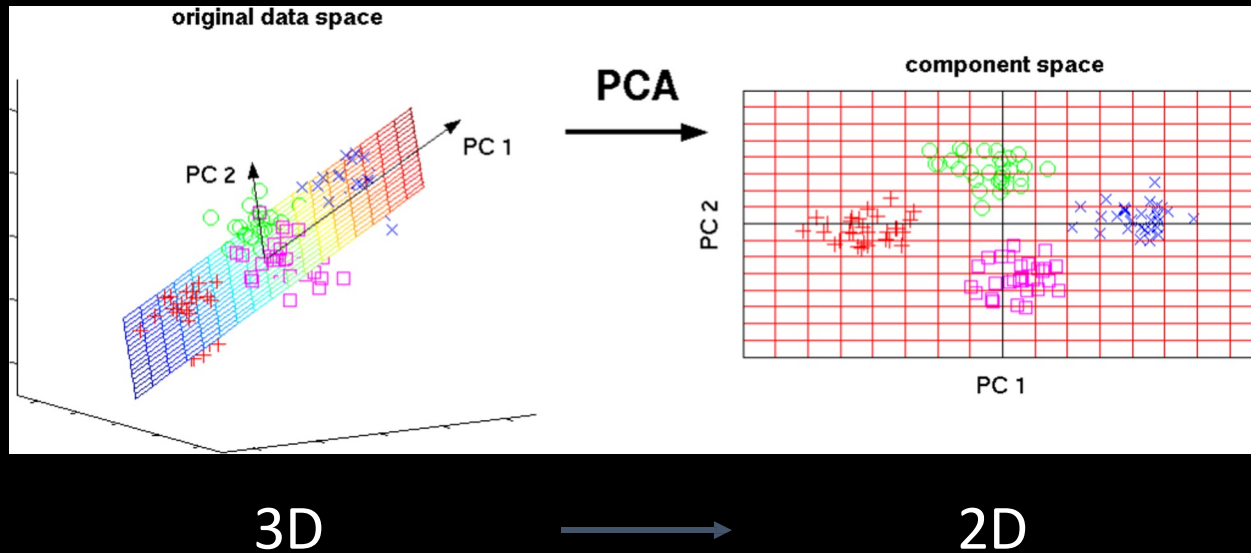
Just data, no labels!

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

Examples: Clustering, dimensionality reduction, feature learning, density estimation, etc.

Supervised vs Unsupervised Learning

Dimensionality Reduction (e.g. Principal Components Analysis)



Unsupervised Learning

Data: x

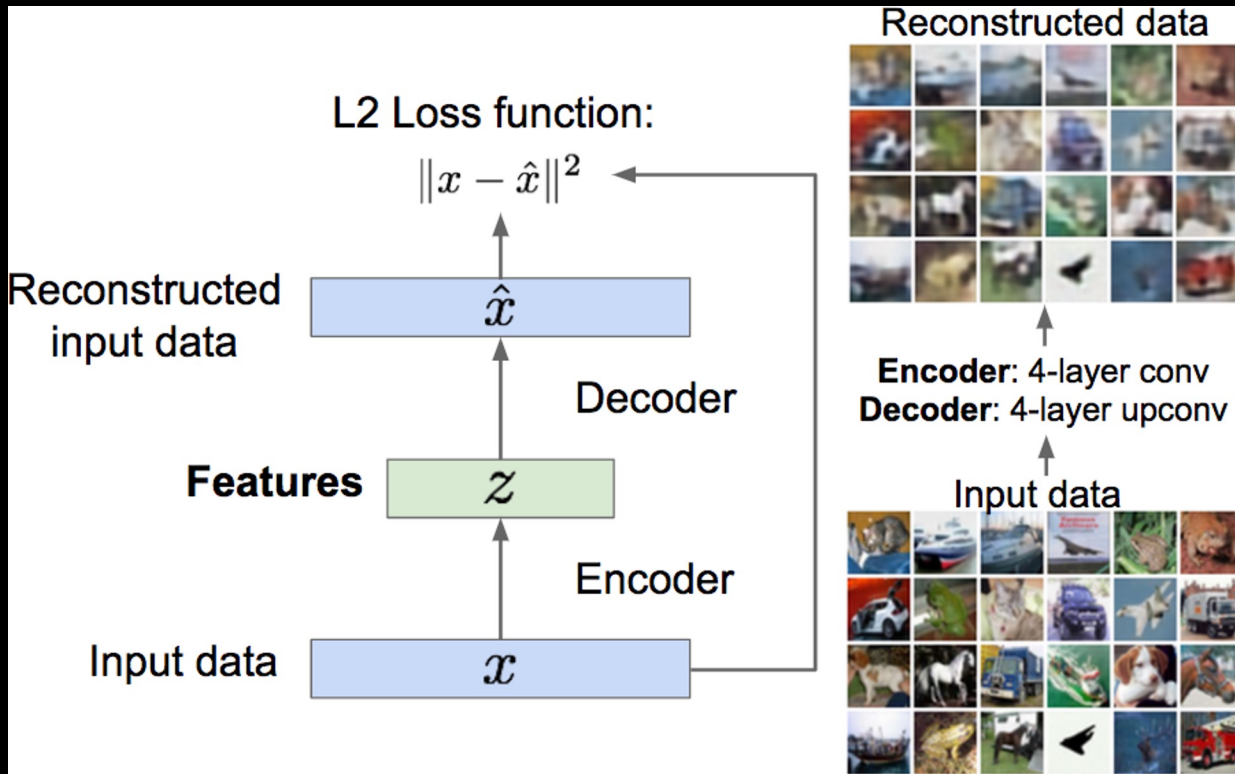
Just data, no labels!

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

Examples: Clustering, dimensionality reduction, feature learning, density estimation, etc.

Supervised vs Unsupervised Learning

Feature Learning (e.g. autoencoders)



Unsupervised Learning

Data: x

Just data, no labels!

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

Examples: Clustering, dimensionality reduction, feature learning, density estimation, etc.

Supervised vs Unsupervised Learning

Supervised Learning

Data: (x, y)

x is data, y is label

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

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

Unsupervised Learning

Data: x

Just data, no labels!

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

Examples: Clustering, dimensionality reduction, feature learning, density estimation, etc.

Discriminative vs Generative Models

Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

Conditional Generative Model: Learn $p(x|y)$

Data: x



Label: y

Cat

Discriminative vs Generative Models

Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

Conditional Generative Model: Learn $p(x|y)$

Data: x



Label: y

Cat

Probability Recap:

Density Function

$p(x)$ assigns a positive number to each possible x ; higher numbers mean x is more likely

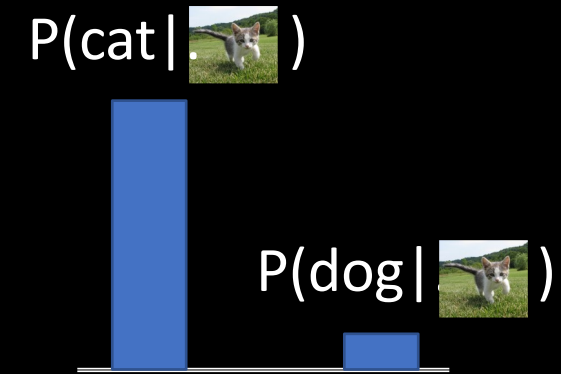
Density functions are **normalized**:

$$\int_{\mathcal{X}} p(x) dx = 1$$

Different values of x **compete** for density

Discriminative vs Generative Models

Discriminative Model:
Learn a probability distribution $p(y|x)$

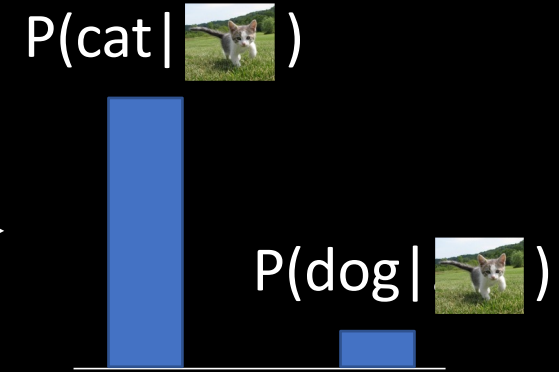


Generative Model:
Learn a probability distribution $p(x)$

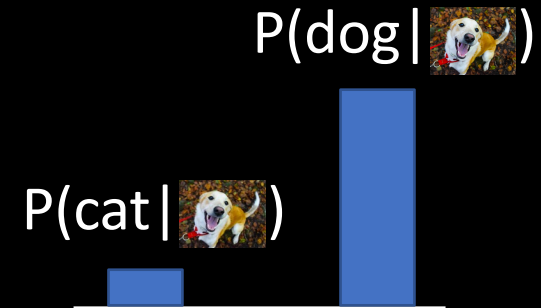
Conditional Generative Model: Learn $p(x|y)$

Discriminative vs Generative Models

Discriminative Model:
Learn a probability distribution $p(y|x)$



Generative Model:
Learn a probability distribution $p(x)$

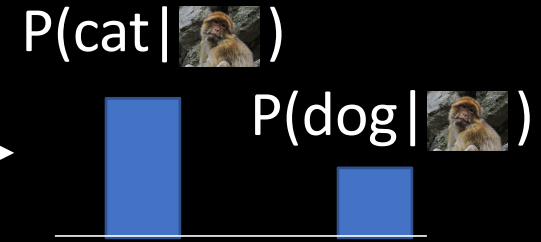


Conditional Generative Model: Learn $p(x|y)$

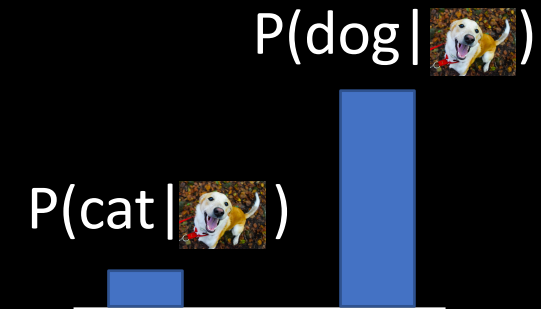
Discriminative model: the possible labels for each input "compete" for probability mass.
But no competition between **images**

Discriminative vs Generative Models

Discriminative Model:
Learn a probability distribution $p(y|x)$



Generative Model:
Learn a probability distribution $p(x)$

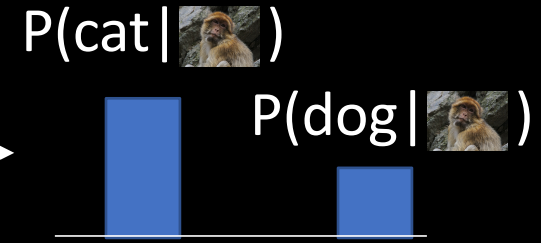


Conditional Generative Model: Learn $p(x|y)$

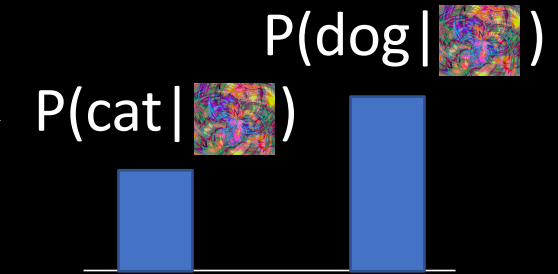
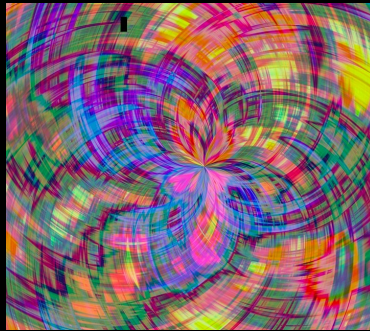
Discriminative model: No way for the model to handle unreasonable inputs; it must give label distributions for all images

Discriminative vs Generative Models

Discriminative Model:
Learn a probability distribution $p(y|x)$



Generative Model:
Learn a probability distribution $p(x)$



Conditional Generative Model: Learn $p(x|y)$

Discriminative model: No way for the model to handle unreasonable inputs; it must give label distributions for all images

Discriminative vs Generative Models

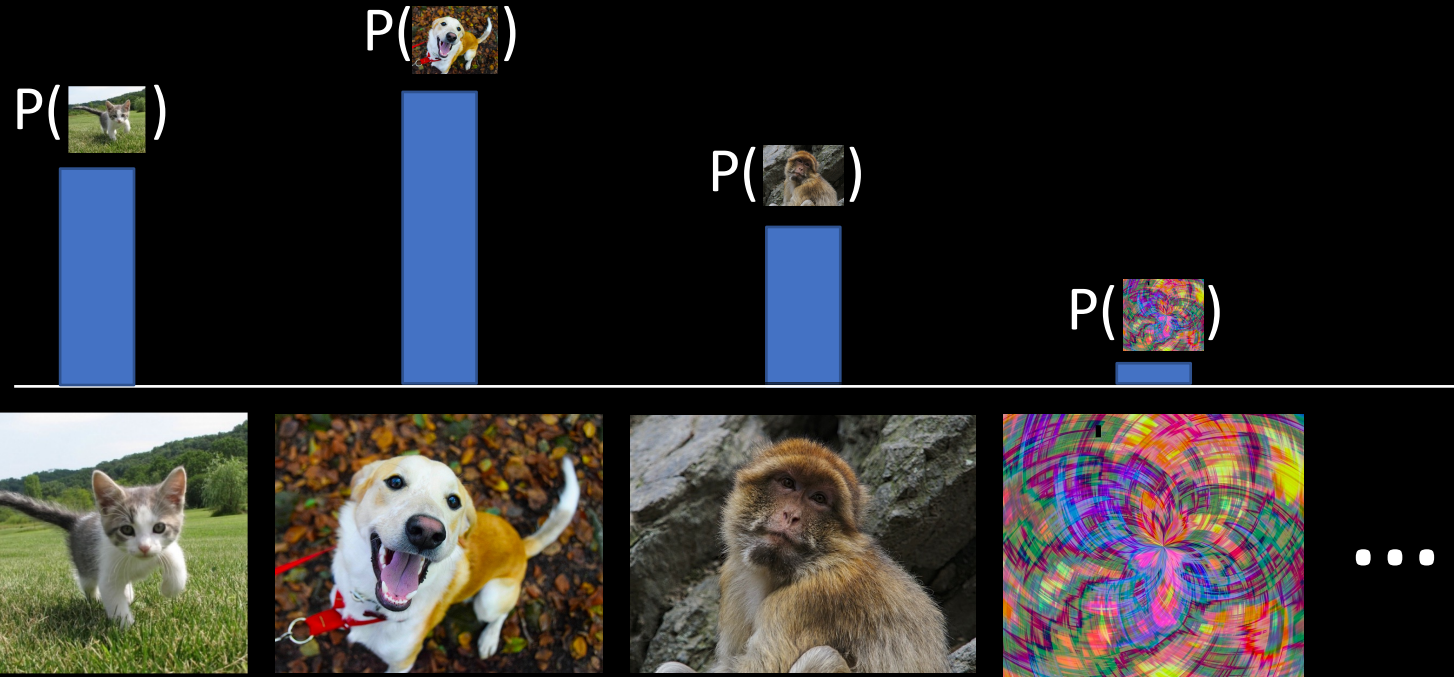
Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

Conditional Generative Model: Learn $p(x|y)$



Generative model: All possible images compete with each other for probability mass

Requires deep image understanding! Is a dog more likely to sit or stand? How about 3-legged dog vs 3-armed monkey?

Model can “reject” unreasonable inputs by assigning them small values

Discriminative vs Generative Models

Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

$$P(\text{cat} | \text{cat})$$

$$P(\text{dog} | \text{cat})$$

$$P(\text{monkey} | \text{cat})$$

$$P(\text{noise} | \text{cat})$$

$$P(\text{cat} | \text{dog})$$

$$P(\text{dog} | \text{dog})$$

$$P(\text{monkey} | \text{dog})$$

$$P(\text{noise} | \text{dog})$$



...

Conditional Generative Model: Learn $p(x|y)$

Conditional Generative Model: Each possible label induces a competition among all images

Discriminative vs Generative Models

Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

Conditional Generative Model: Learn $p(x|y)$

Recall Bayes' Rule:

$$P(x|y) = \frac{P(y|x)P(x)}{P(y)}$$

Discriminative vs Generative Models

Discriminative Model:

Learn a probability distribution $p(y|x)$

Generative Model:

Learn a probability distribution $p(x)$

Conditional Generative Model: Learn $p(x|y)$

Recall Bayes' Rule:

$$\boxed{P(x|y)} = \frac{\boxed{P(y|x)}}{\boxed{P(y)}} \boxed{P(x)}$$

Conditional Generative Model

Discriminative Model

Prior over labels

(Unconditional) Generative Model

We can build a conditional generative model from other components!

What can we do with a discriminative model?

- **Discriminative Model:**

Learn a probability distribution $p(y|x)$



Assign labels to data

Feature learning (with labels)

- **Generative Model:**

Learn a probability distribution $p(x)$

- **Conditional Generative**

Model: Learn $p(x|y)$

What can we do with a generative model?

- **Discriminative Model:**

Learn a probability distribution $p(y|x)$



Assign labels to data

Feature learning (with labels)

- **Generative Model:**

Learn a probability distribution $p(x)$



Detect outliers

Feature learning (without labels)

Sample to **generate** new data

- **Conditional Generative**

Model: Learn $p(x|y)$

What can we do with a generative model?

- **Discriminative Model:**

Learn a probability distribution $p(y|x)$



Assign labels to data

Feature learning (with labels)

- **Generative Model:**

Learn a probability distribution $p(x)$



Detect outliers

Feature learning (without labels)

Sample to **generate** new data

- **Conditional Generative**

Model: Learn $p(x|y)$

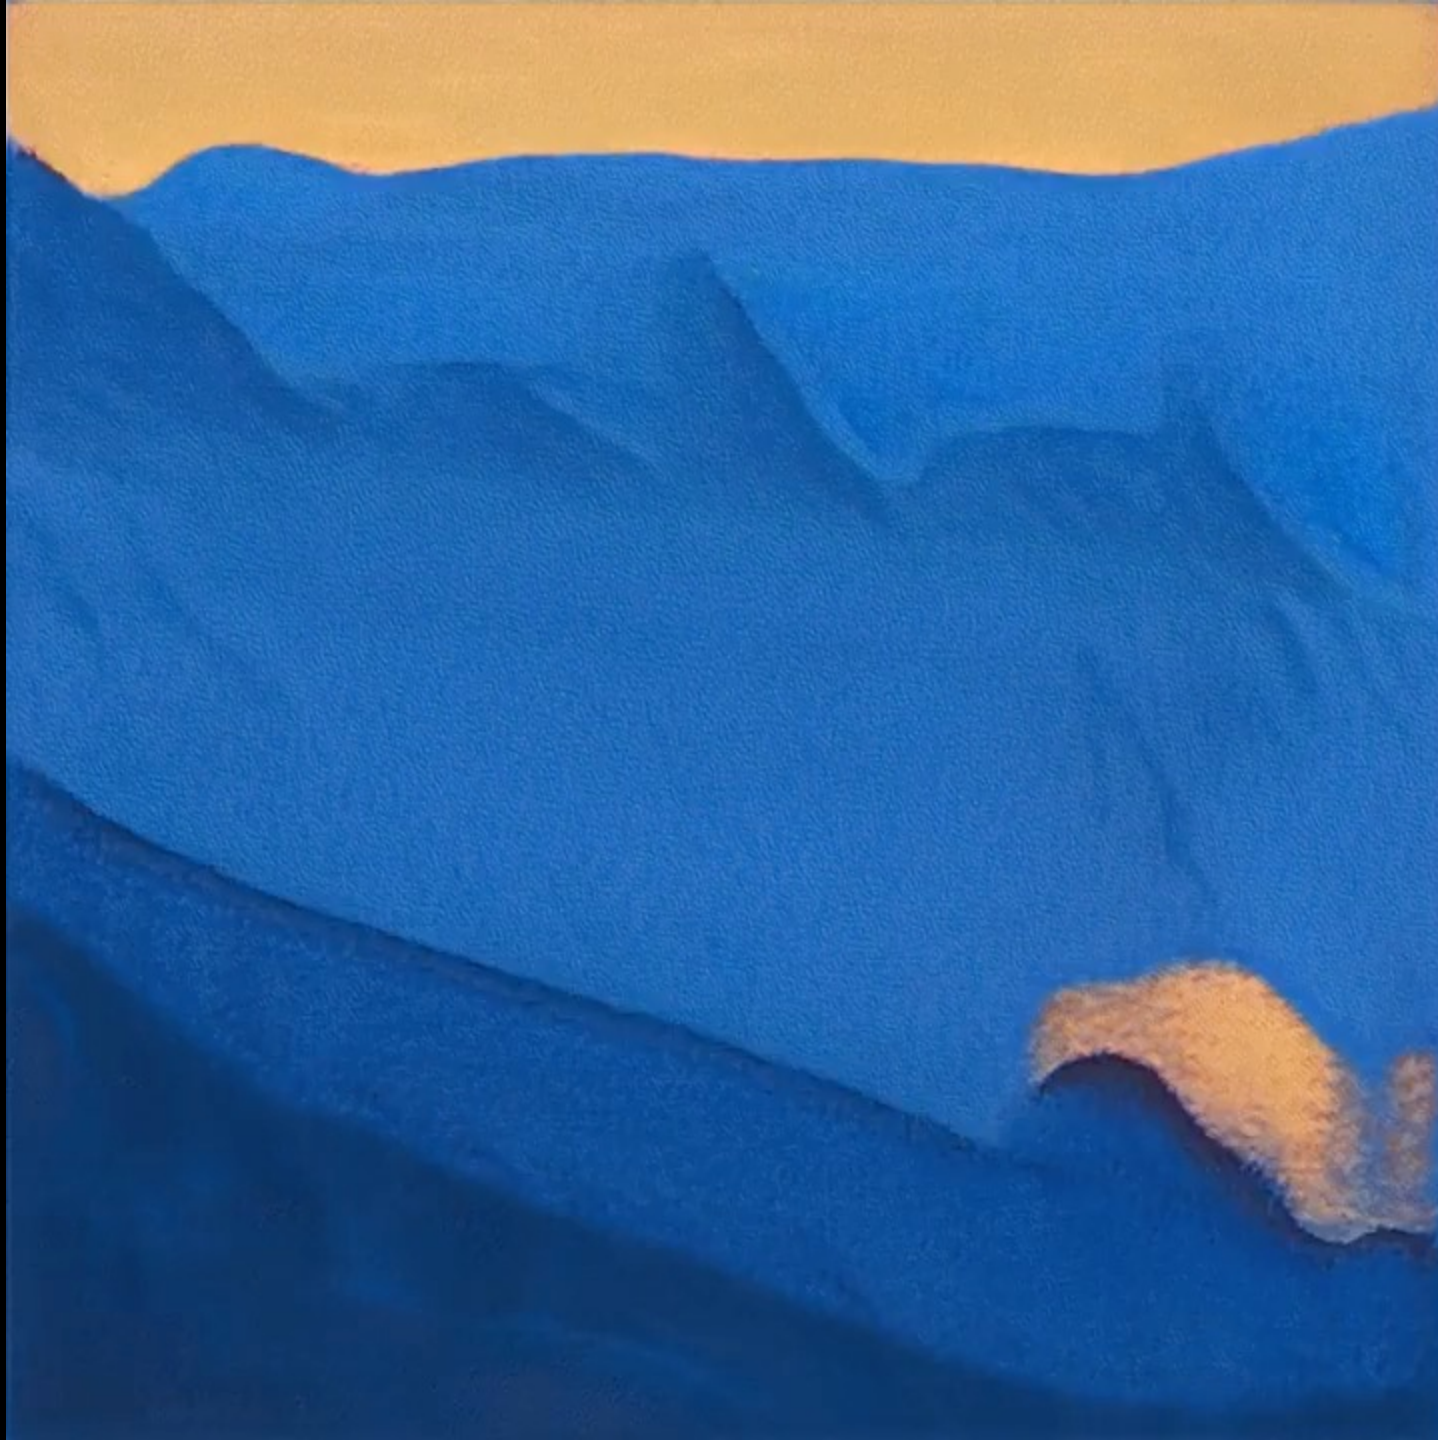


Assign labels, while rejecting outliers!

Generate new data conditioned on input labels

Introduction to Generative Models (Conditional and Unconditional)

What cool things can we do with it?



Credits: Neural Synesthesia

Click on the person who is real.

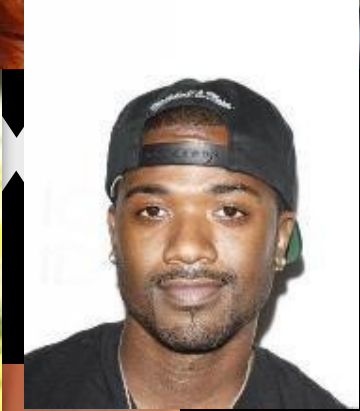
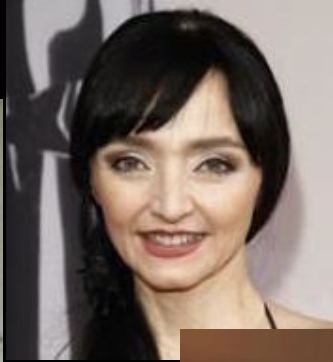
<https://www.whichfaceisreal.com/index.php>

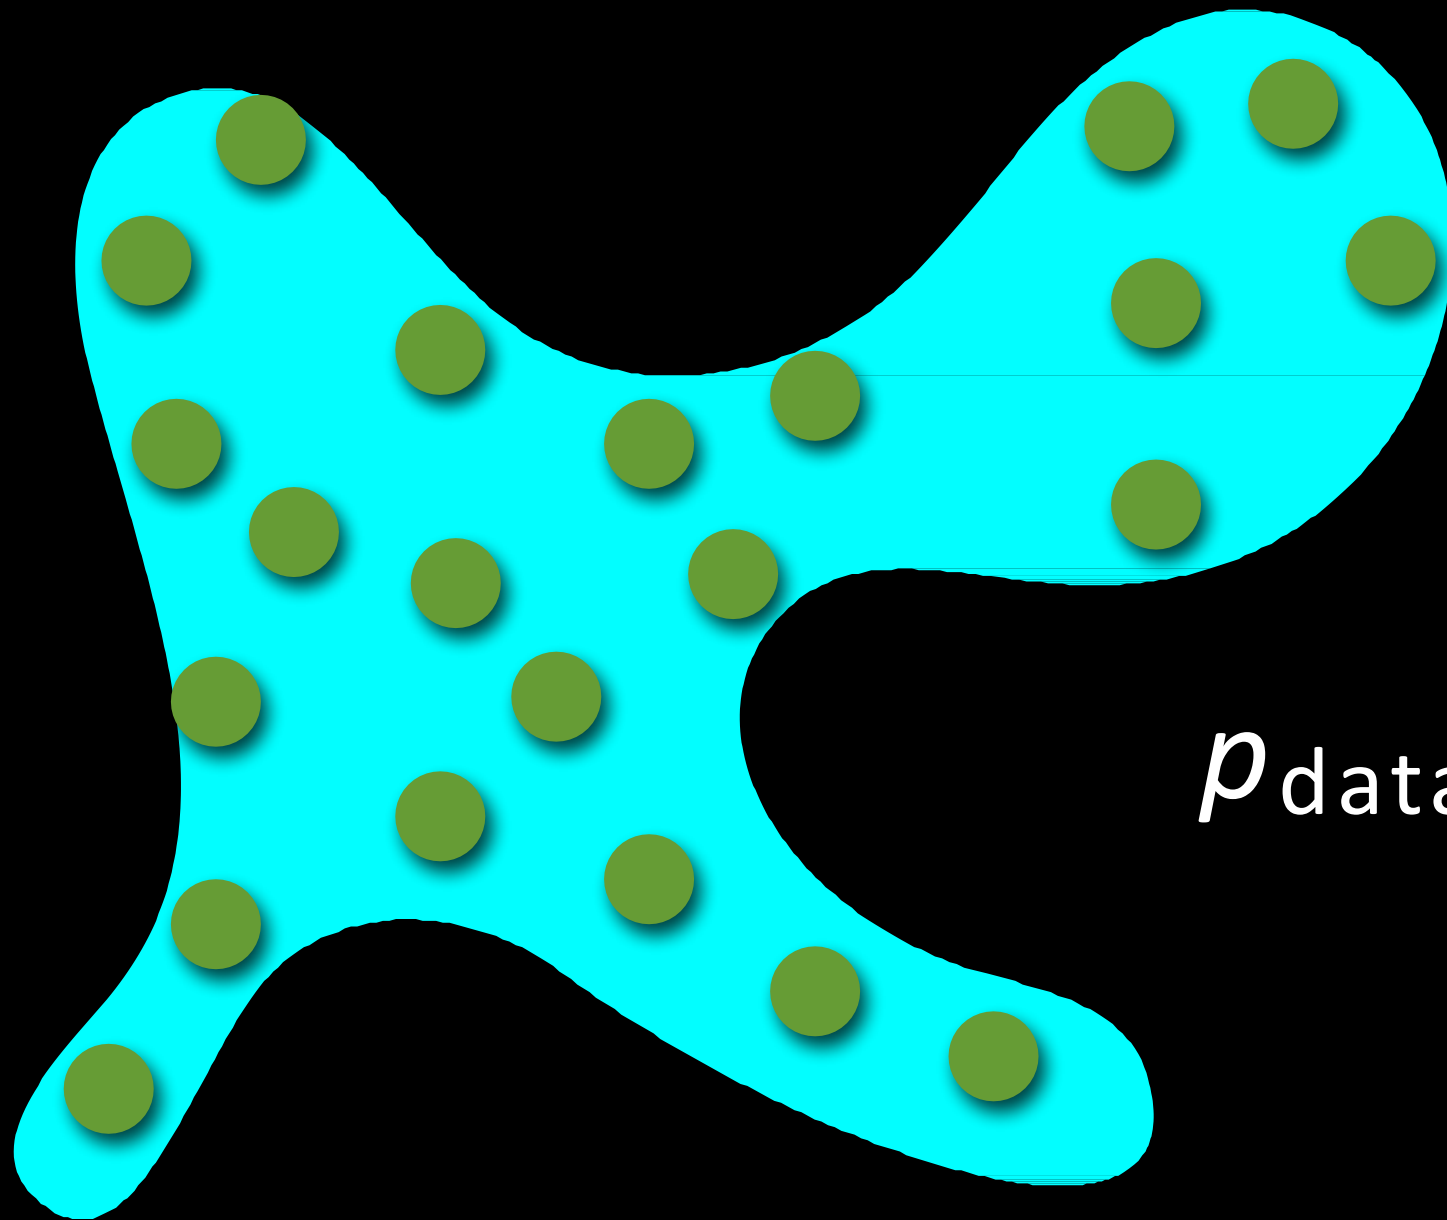




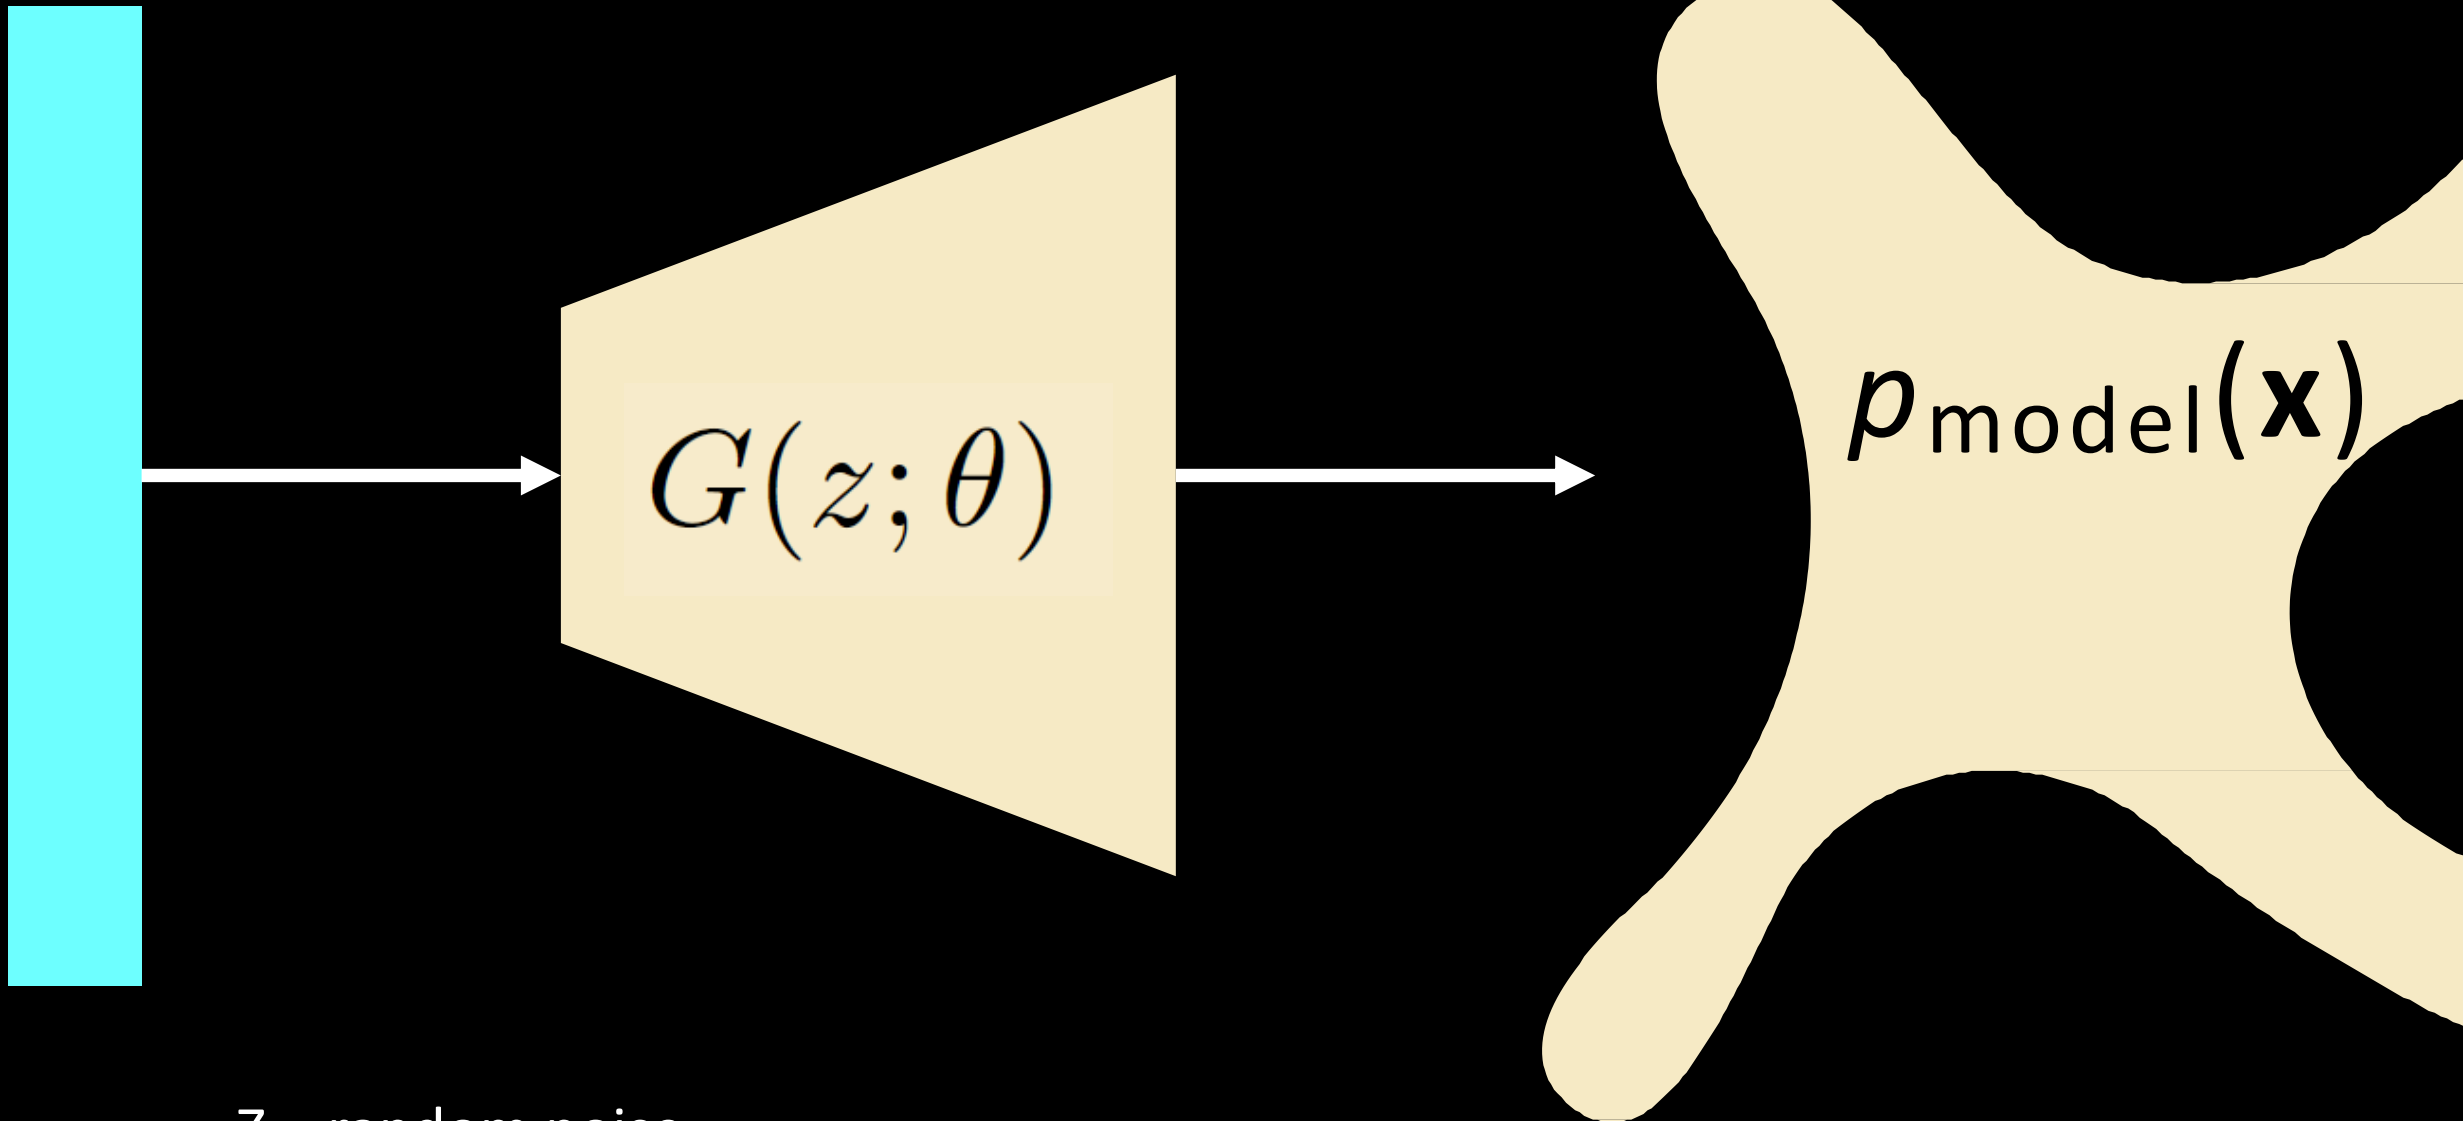
$p \text{ data}(x)$

$$p_{\text{data}}(\mathbf{x}) \approx p_{\text{model}}(\mathbf{x})$$

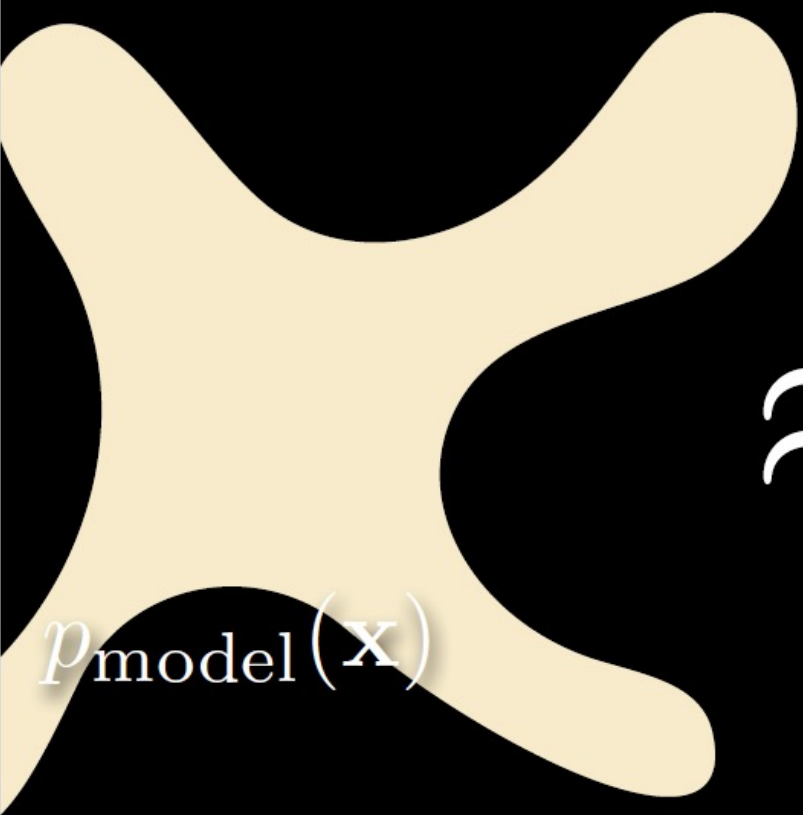




$p_{\text{data}}(\mathbf{x})$

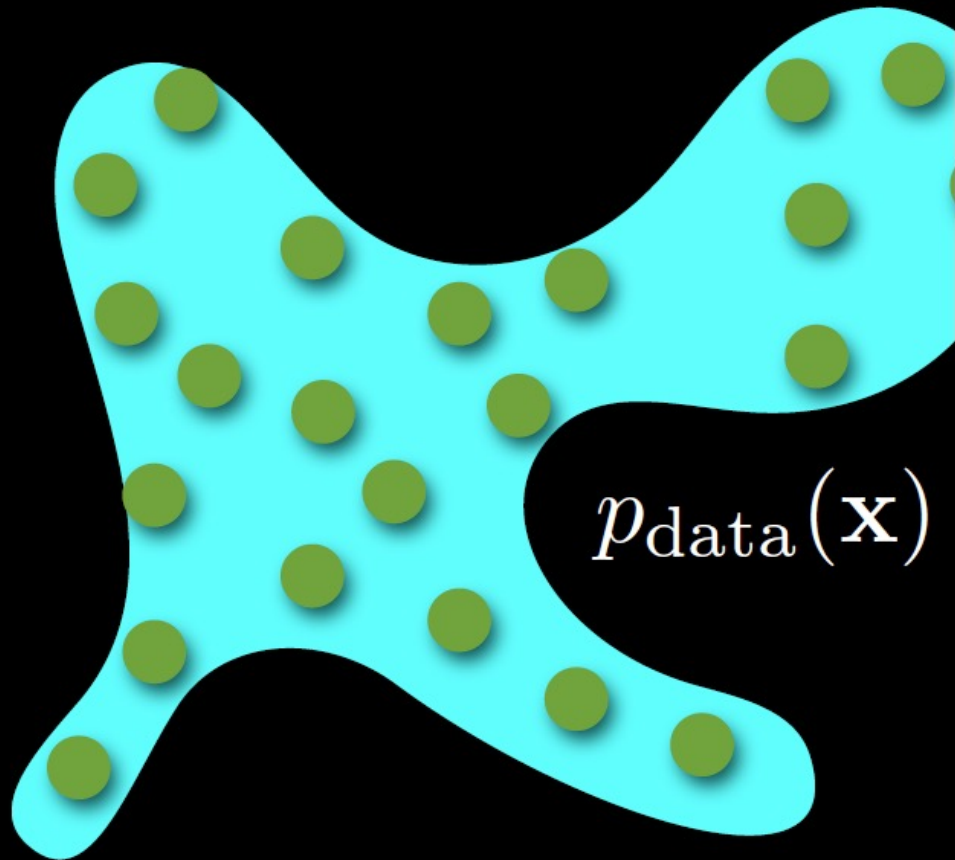


$Z = \text{random noise}$
(samples from latent space)



$p_{\text{model}}(\mathbf{x})$

\approx



$p_{\text{data}}(\mathbf{x})$

Taxonomy of Generative Models

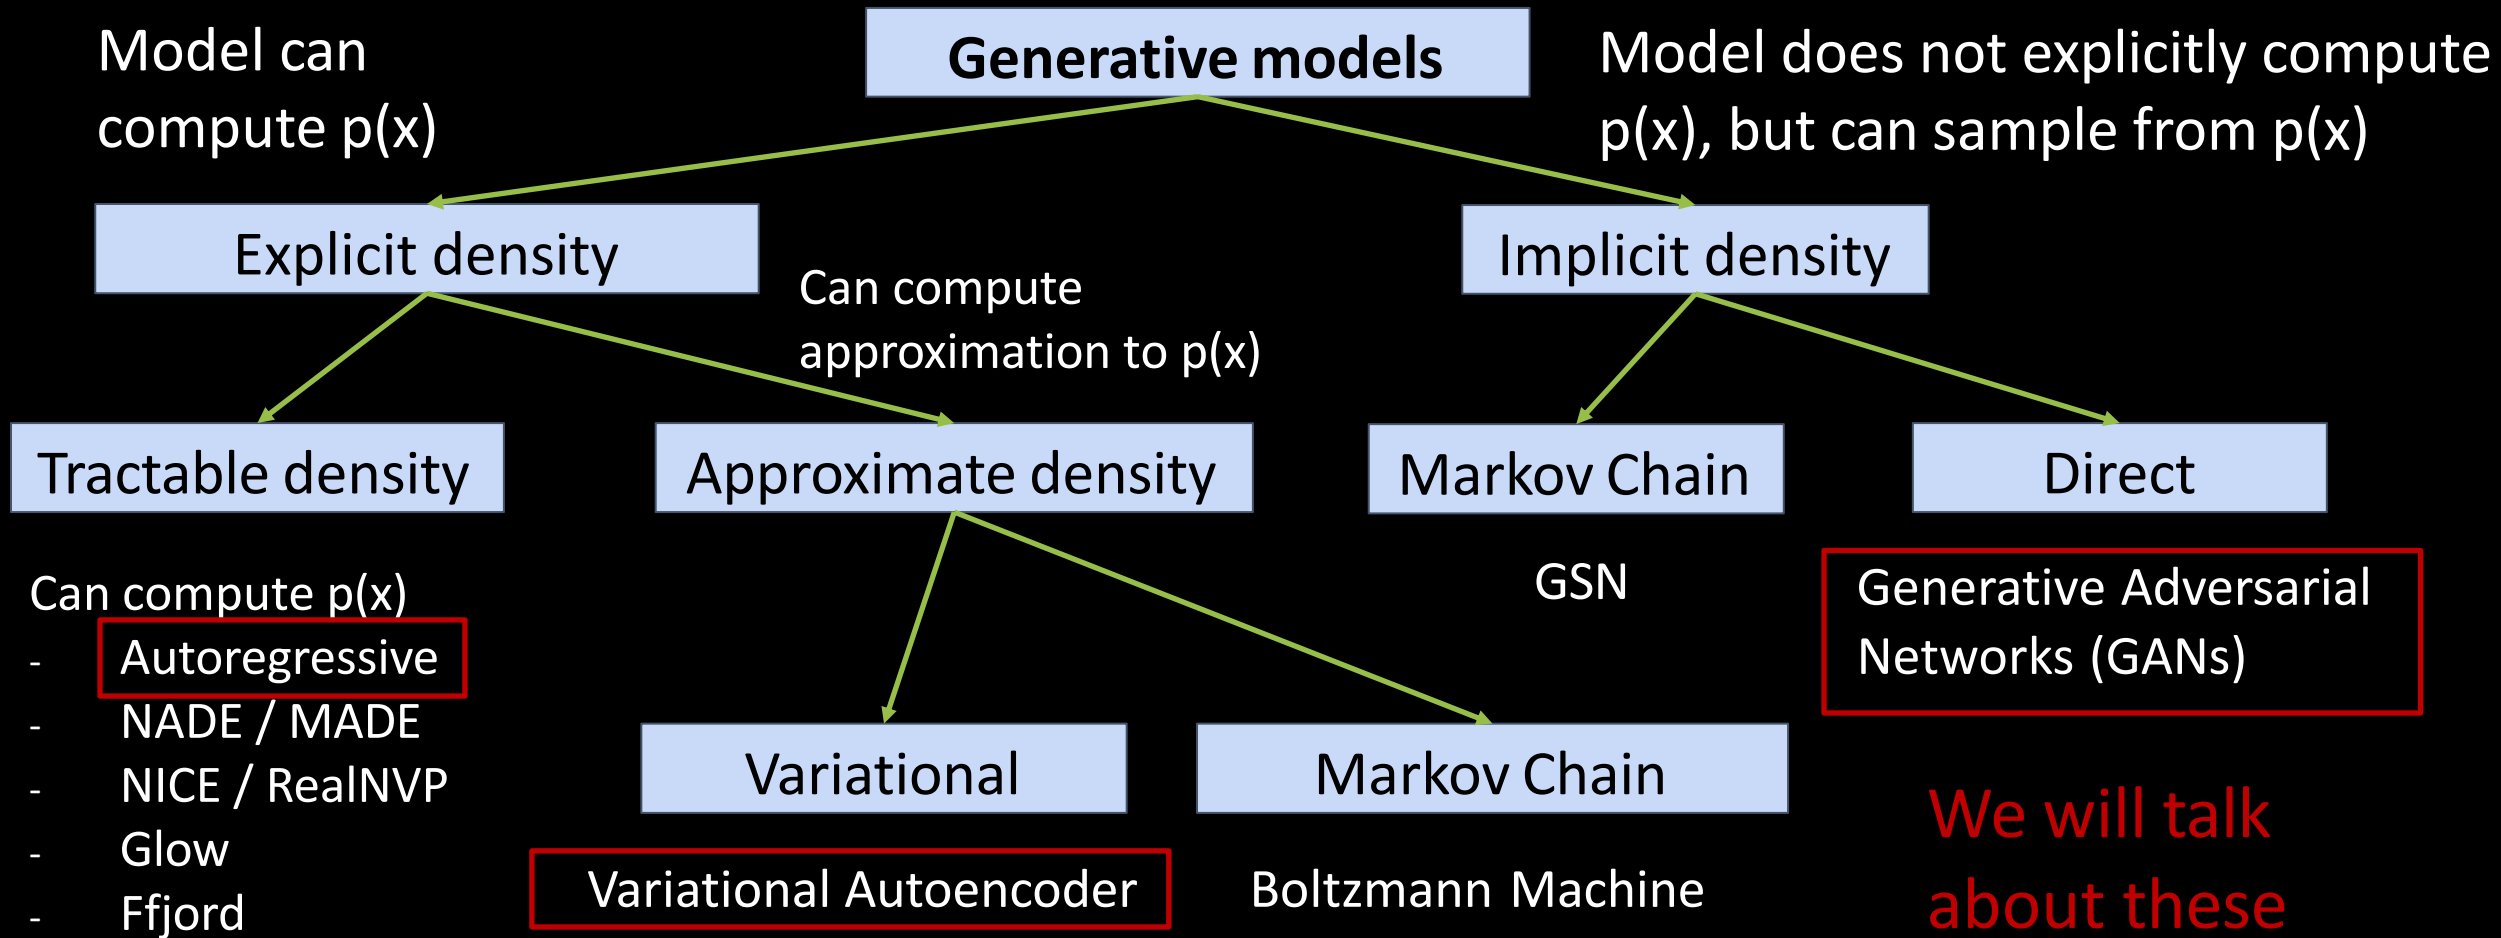


Figure adapted from Ian Goodfellow, Tutorial on Generative Adversarial Networks, 2017.

Autoregressive models: PixelRNN & PixelCNN

Explicit Density Estimation

Learn weights W

Goal: Write down an explicit function for $p(x) = f(x, W)$

Given dataset $x^{(1)}, x^{(2)}, \dots, x^{(N)}$, train the model by solving:

$$W^* = \arg \max_W \prod_i p(x^{(i)})$$

Maximize probability of training data
(Maximum likelihood estimation)

$$= \arg \max_W \sum_i \log p(x^{(i)})$$

Log trick to exchange product for sum

$$= \arg \max_W \sum_i \log f(x^{(i)}, W)$$

This will be our loss function!
Train with gradient descent

x_1

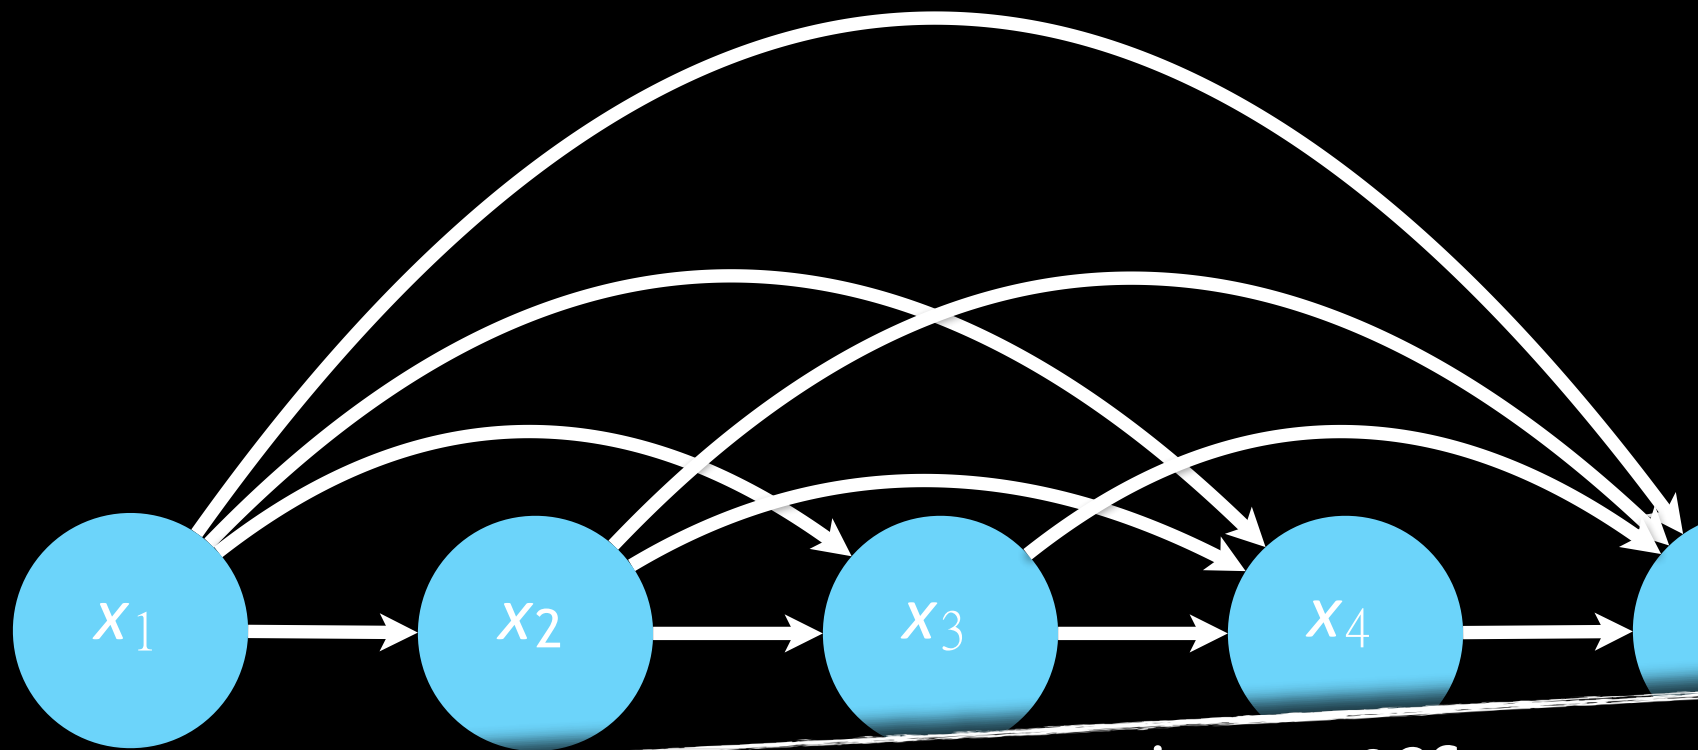
x_2

x_3

x_4



current pixel is dependent on all previous ones



current pixel is dependent on all previous ones

$p_{\text{model}}(\mathbf{x})$

use chain rule to decompose

$$p_{\text{model}}(\mathbf{x}) = \prod_{i=1}^N p(x_i | x_1, \dots, x_{i-1})$$

$$p_{\text{model}}(\mathbf{x}) = \prod_{i=1}^N p(x_i | x_1, \dots, x_{i-1})$$

probability of i th pixel given all previous pixels

$$p_{\text{model}}(\mathbf{x}) = \prod_{i=1}^N p(x_i | x_1, \dots, x_{i-1})$$

express distribution over pixel values as neural network

$$p_{\text{model}}(\mathbf{x}) = \prod_{i=1}^N p(x_i | x_1, \dots, x_{i-1})$$

**maximize likelihood of training data
with respect to network weights**

Maximum Likelihood estimate: $\mathit{arg\,max}_{\theta} \log f_{model}(x; \theta)$

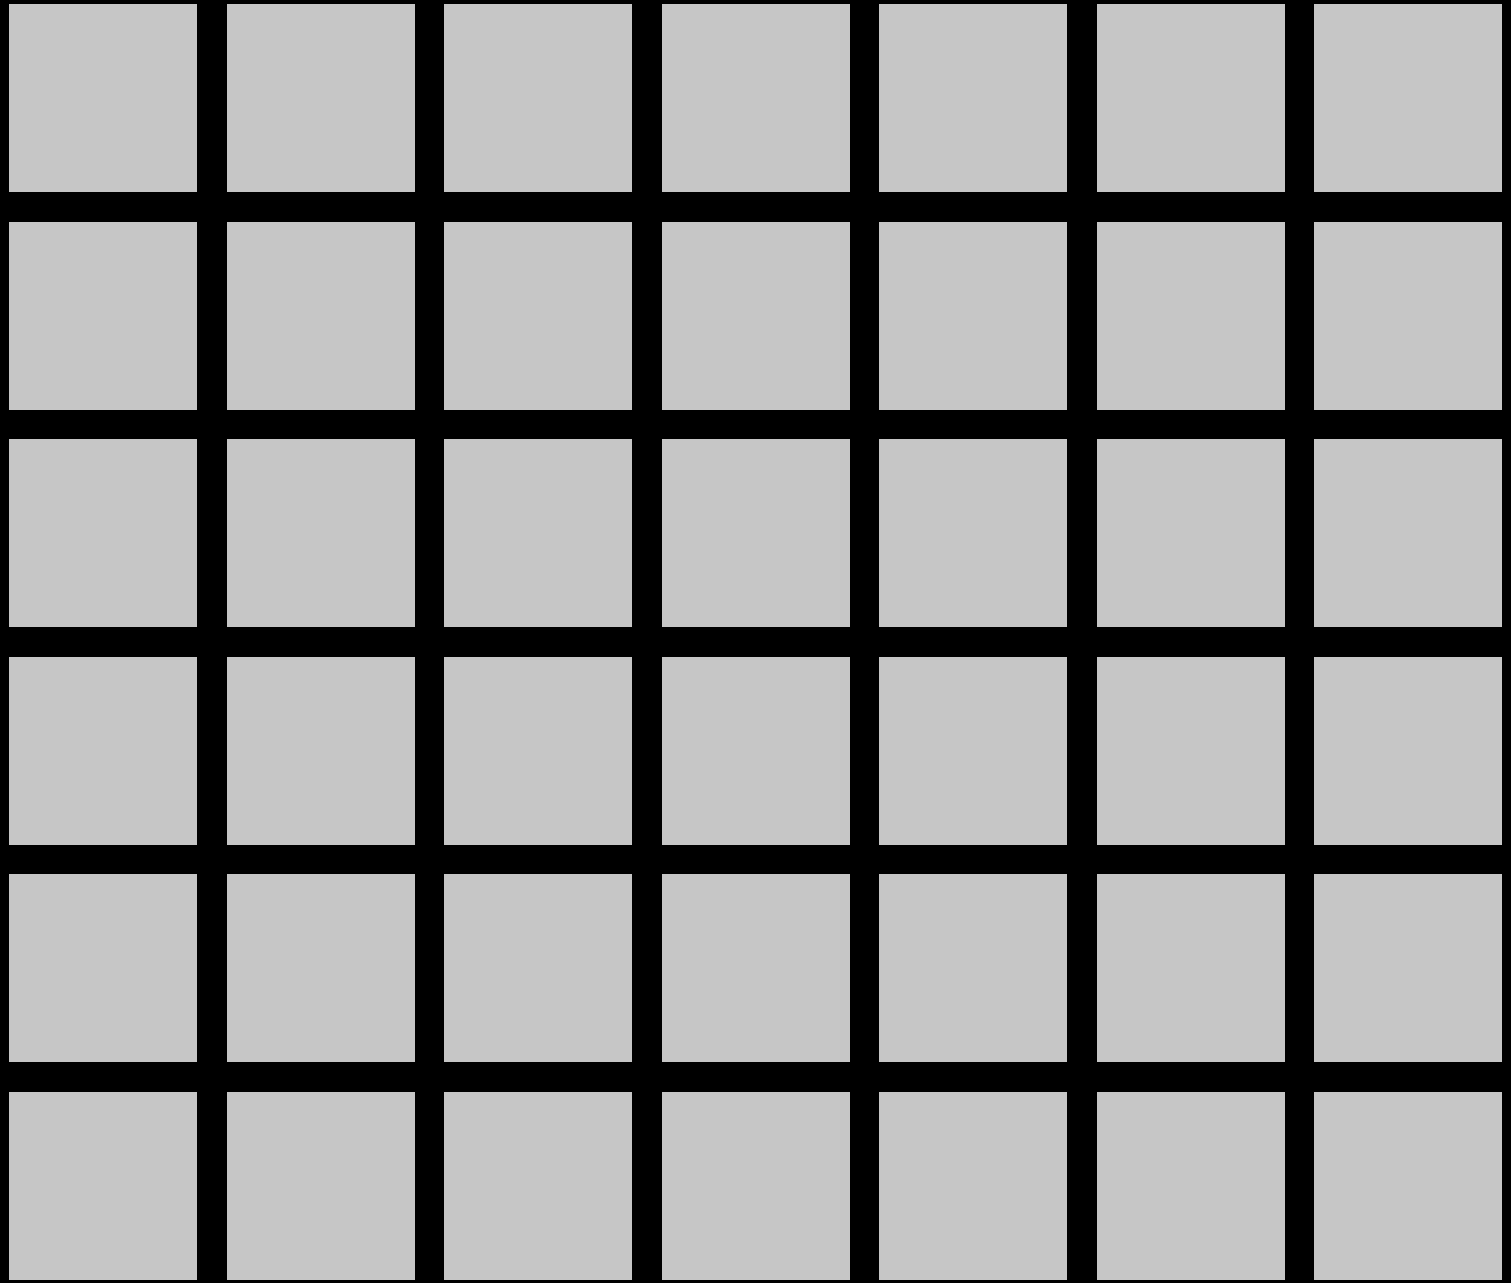
- Maximize this estimator w.r.t. neural network parameter θ .
- In practice we minimize negative of the above estimator as our loss function.

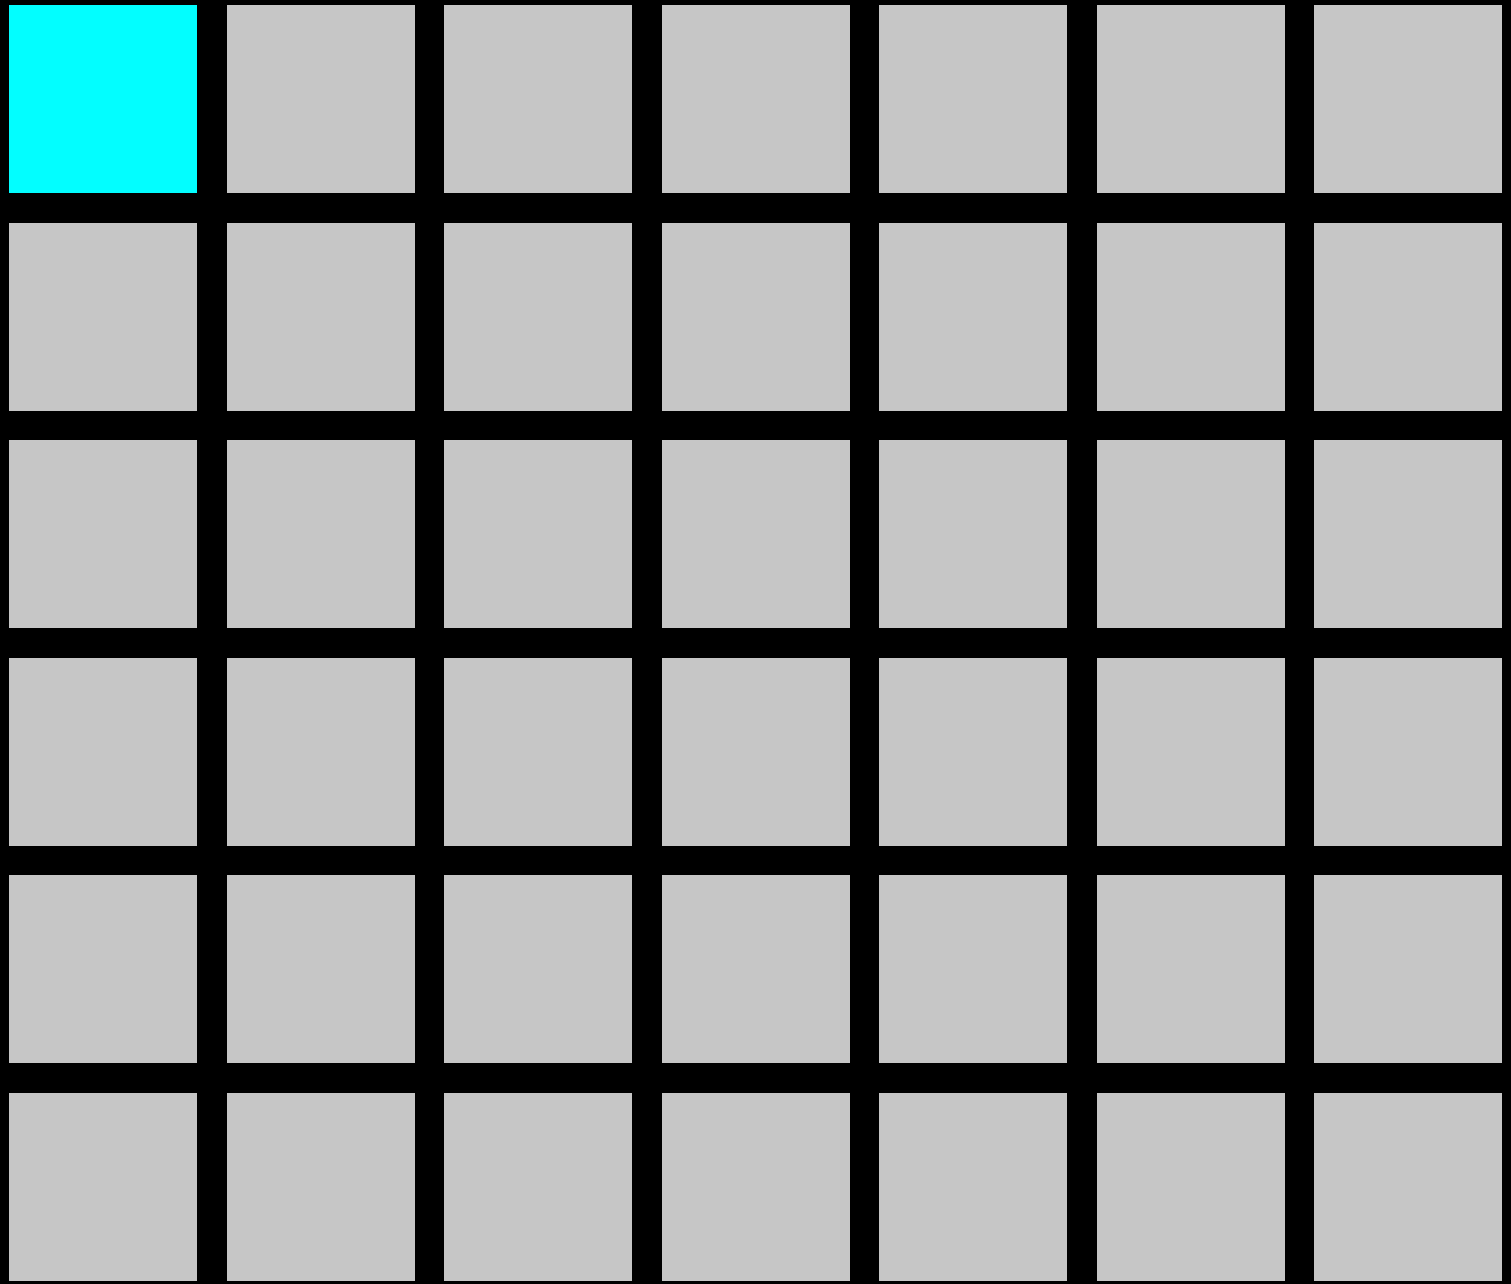
x_1	x_2					x_i
x_{i+1}						

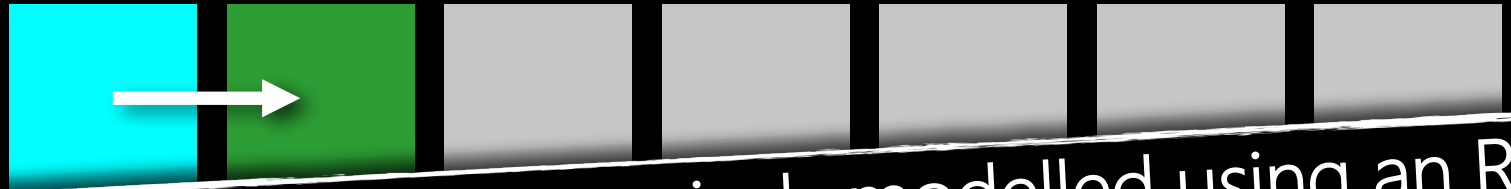
pixel values generated sequentially

PixelRNN

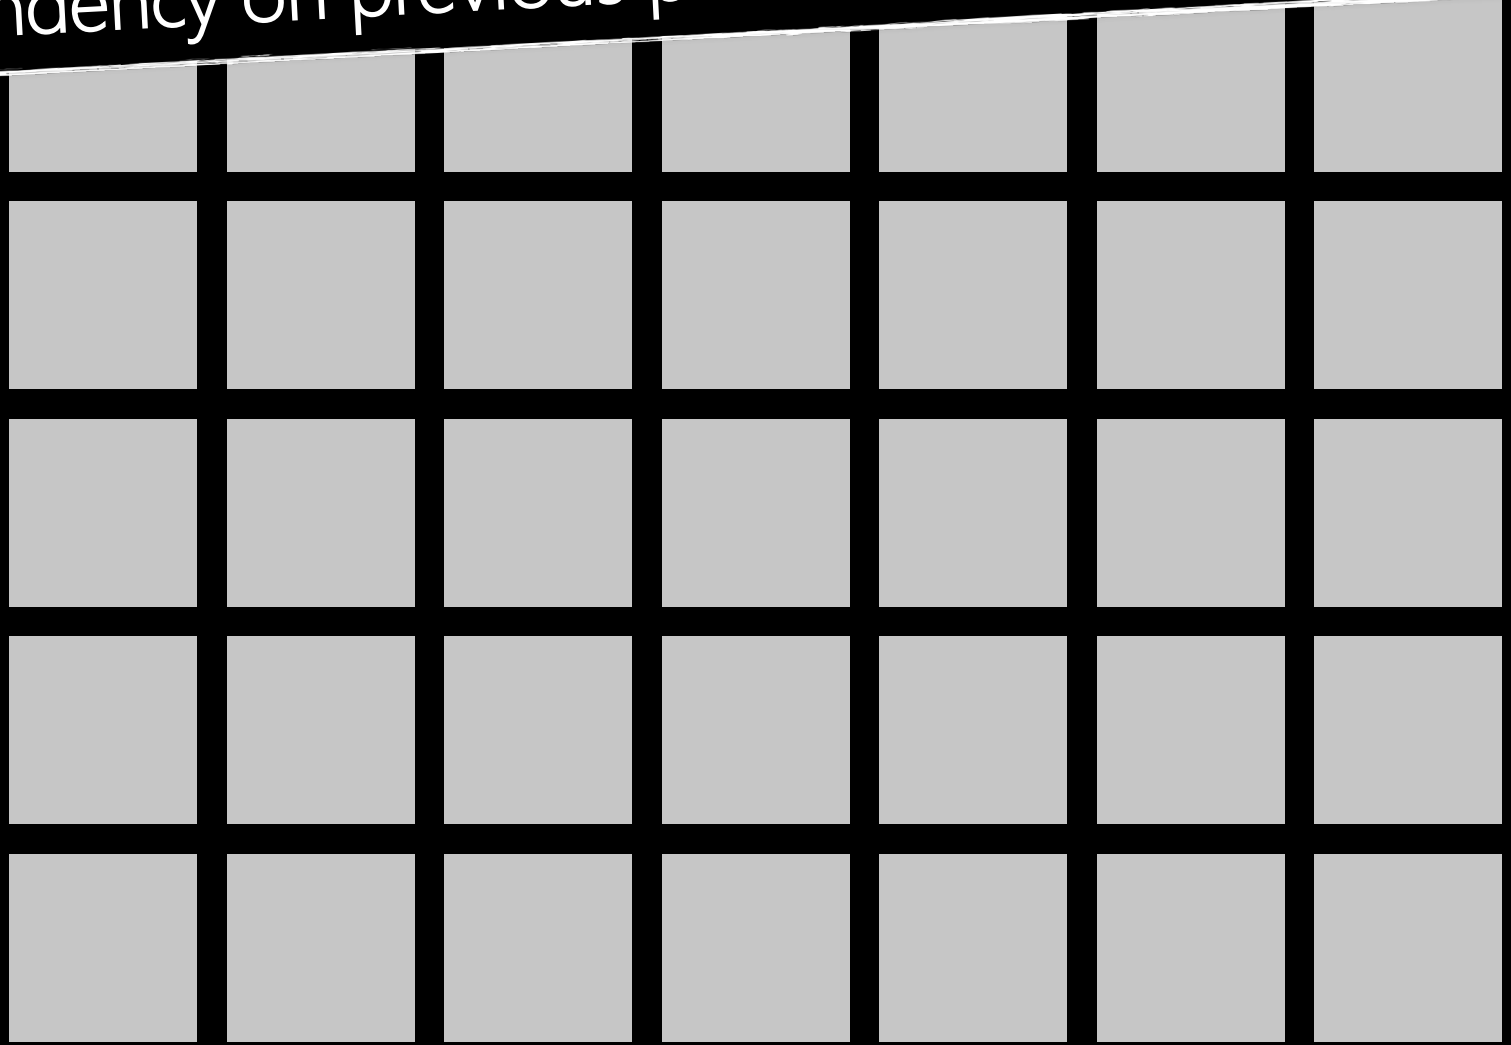
PixelRNN

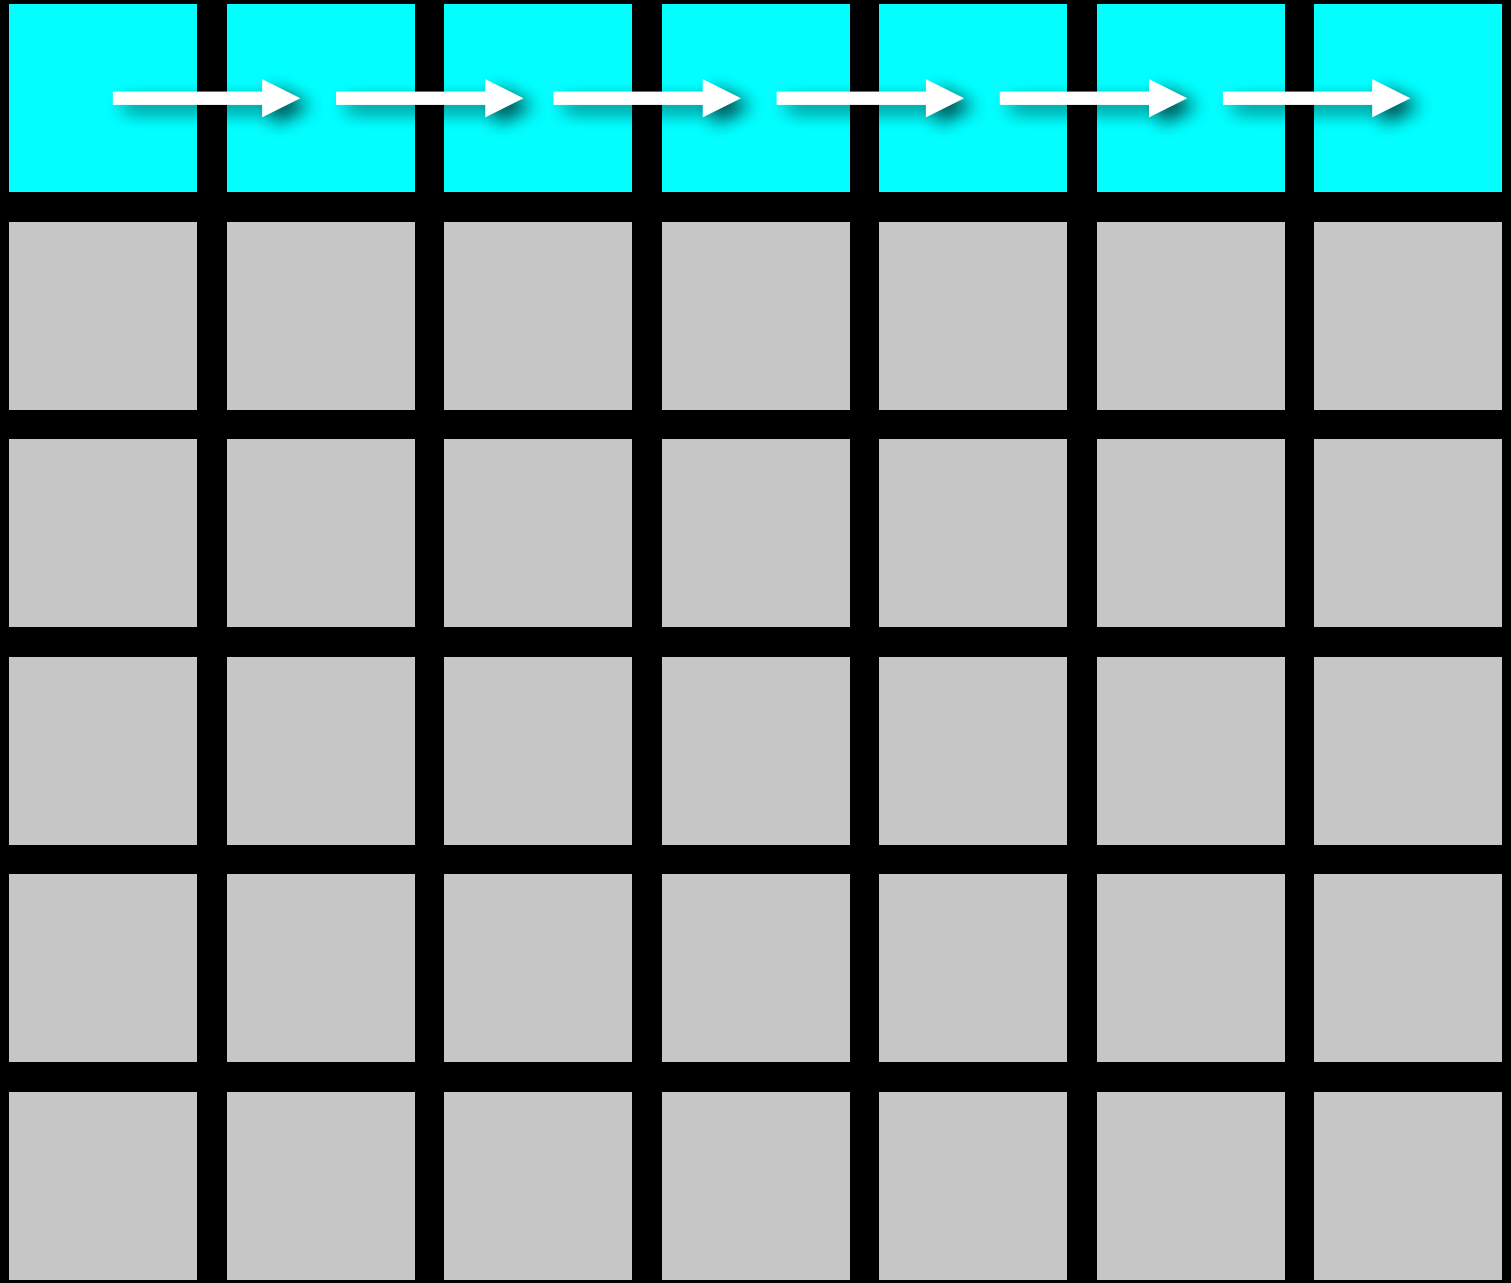


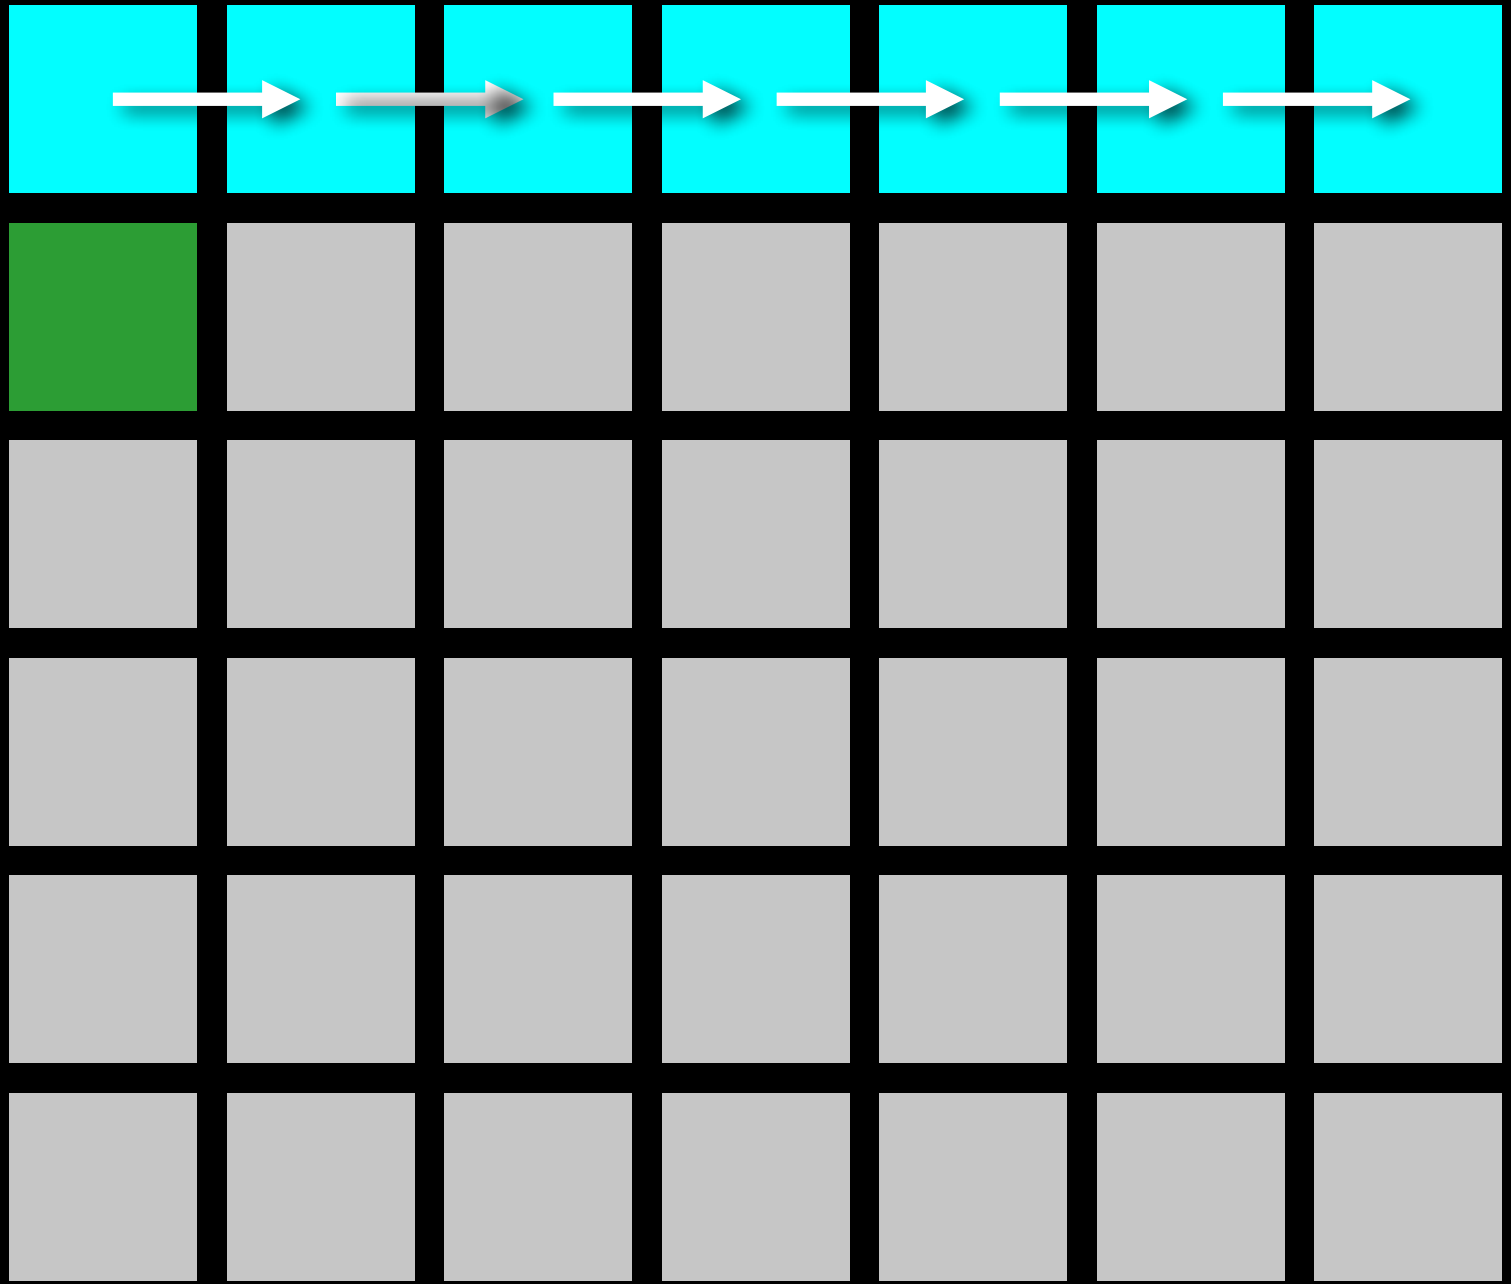


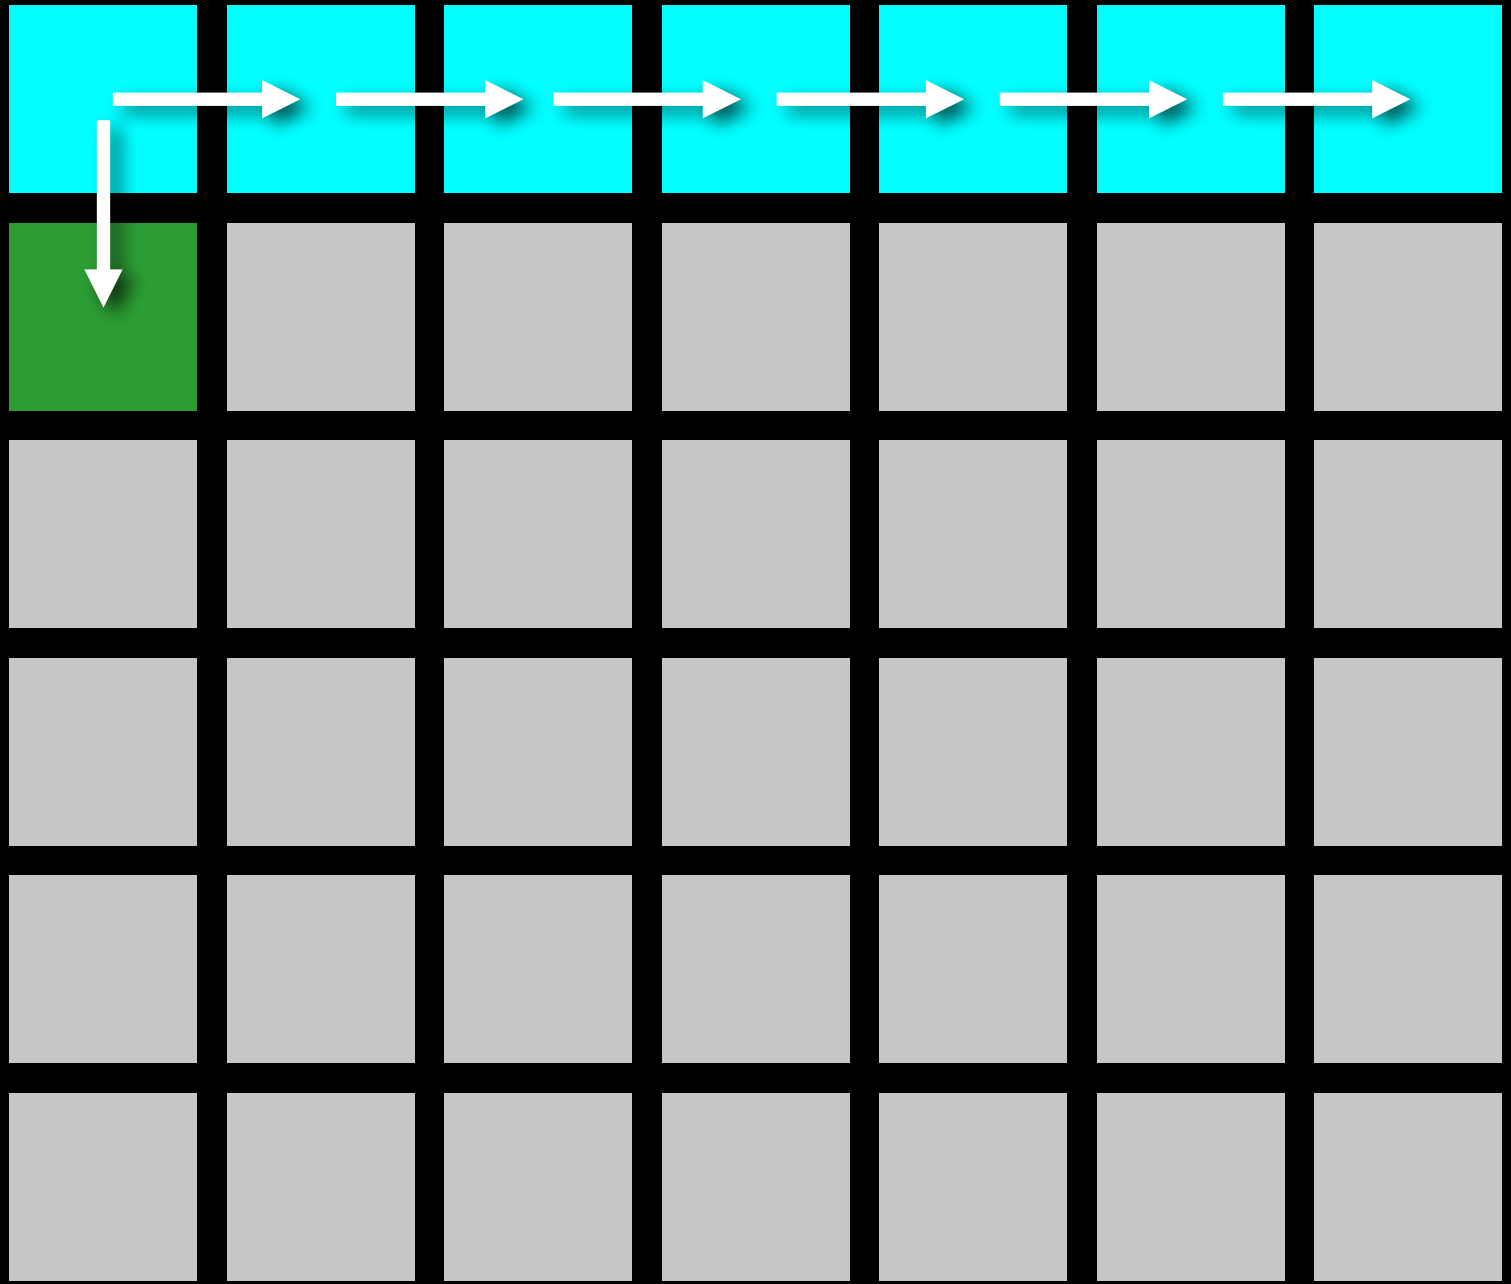


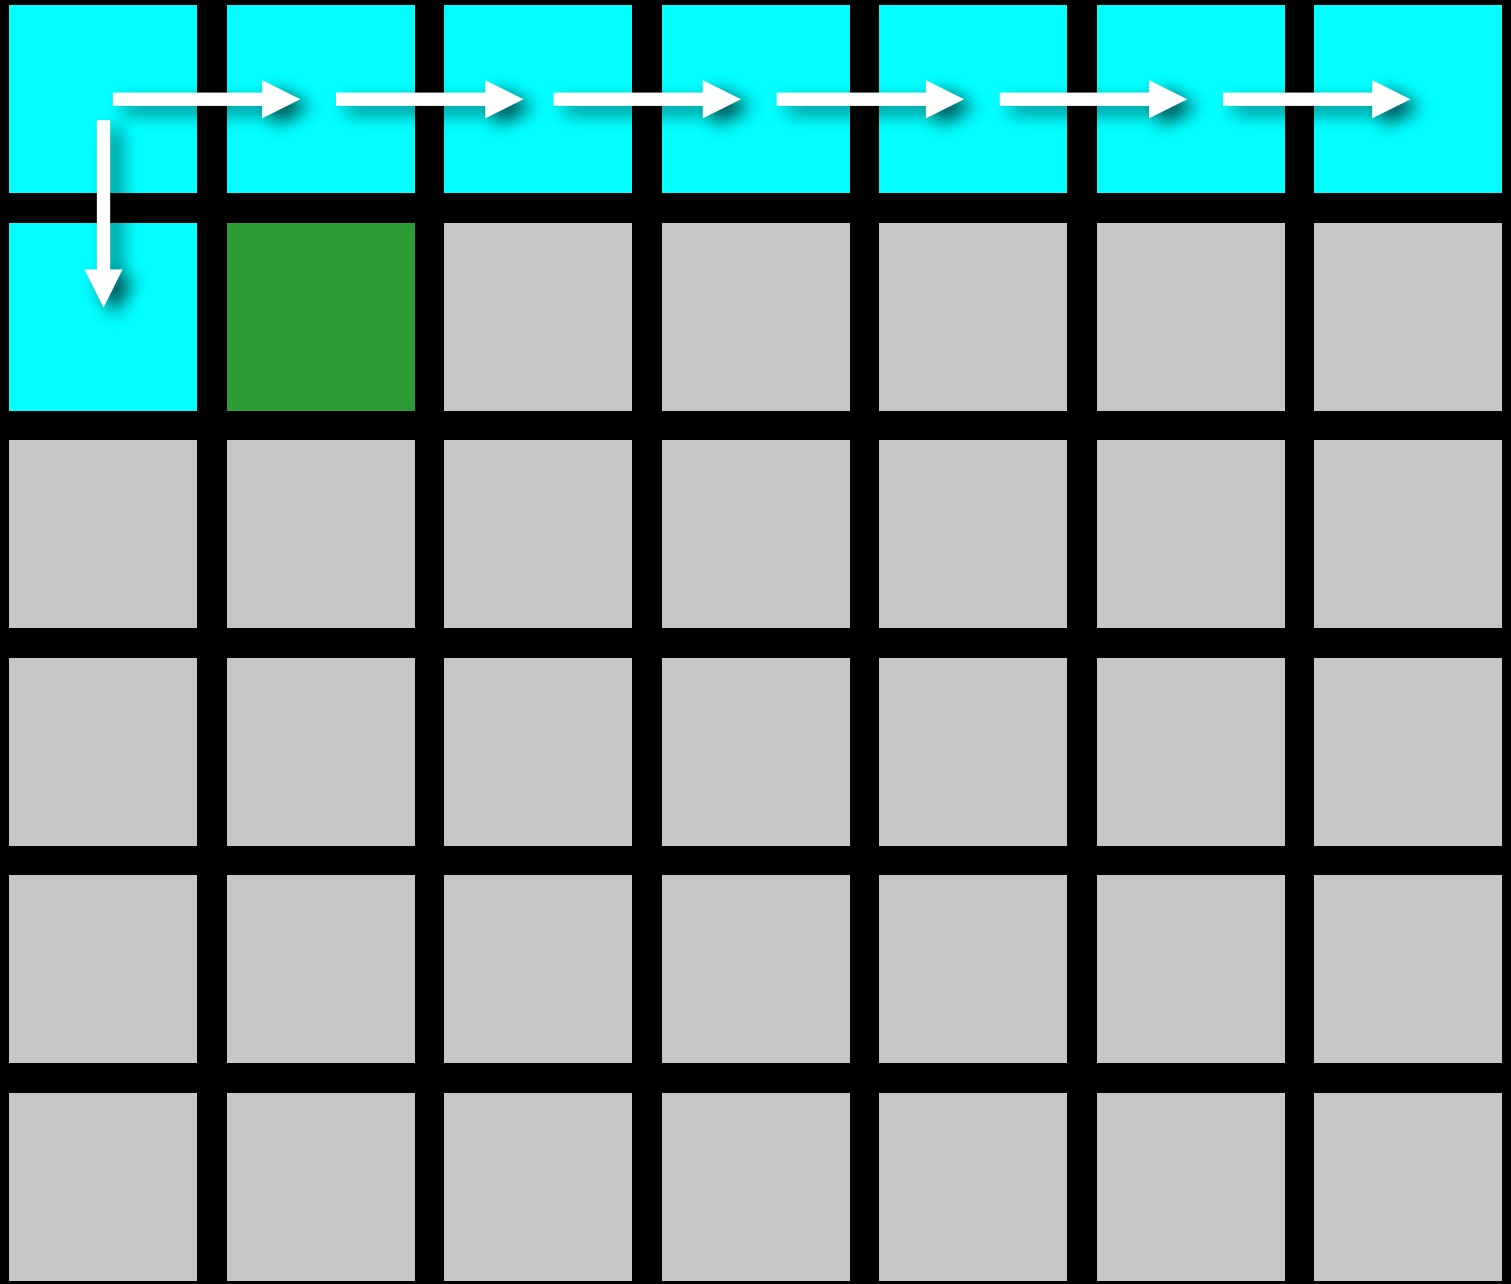
dependency on previous pixels modelled using an RNN

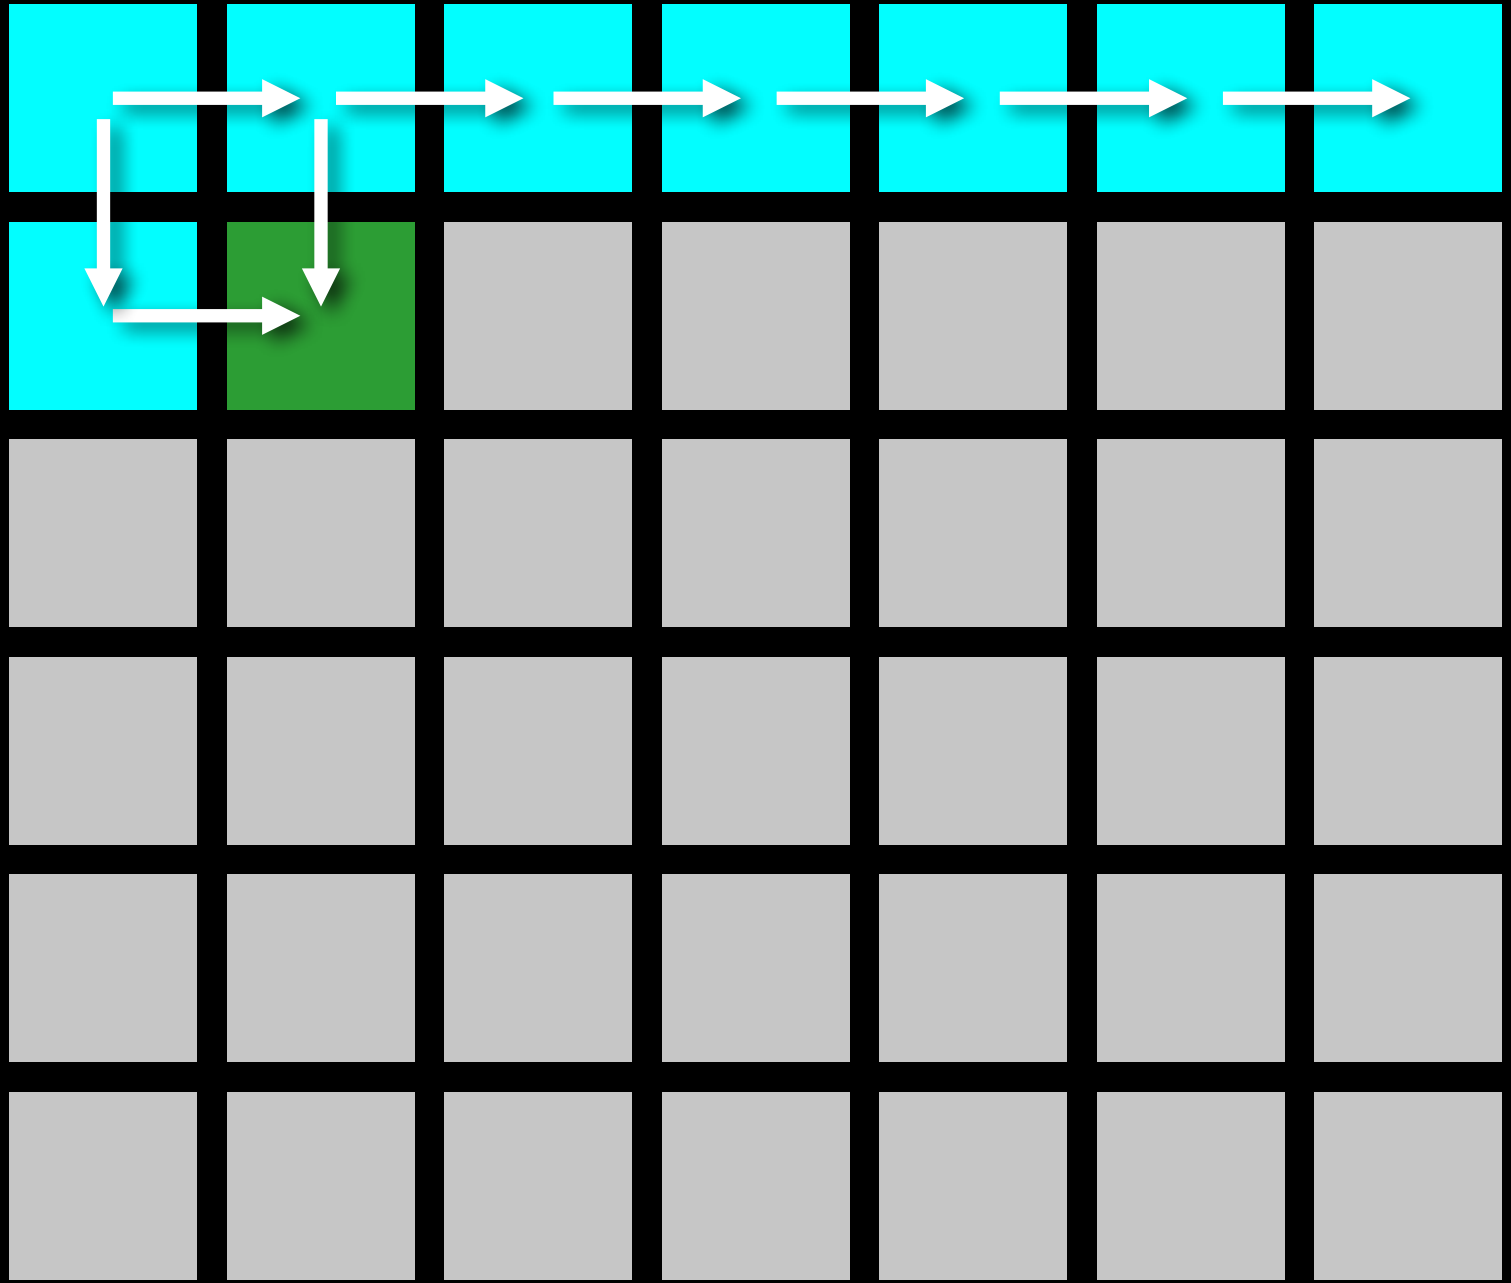




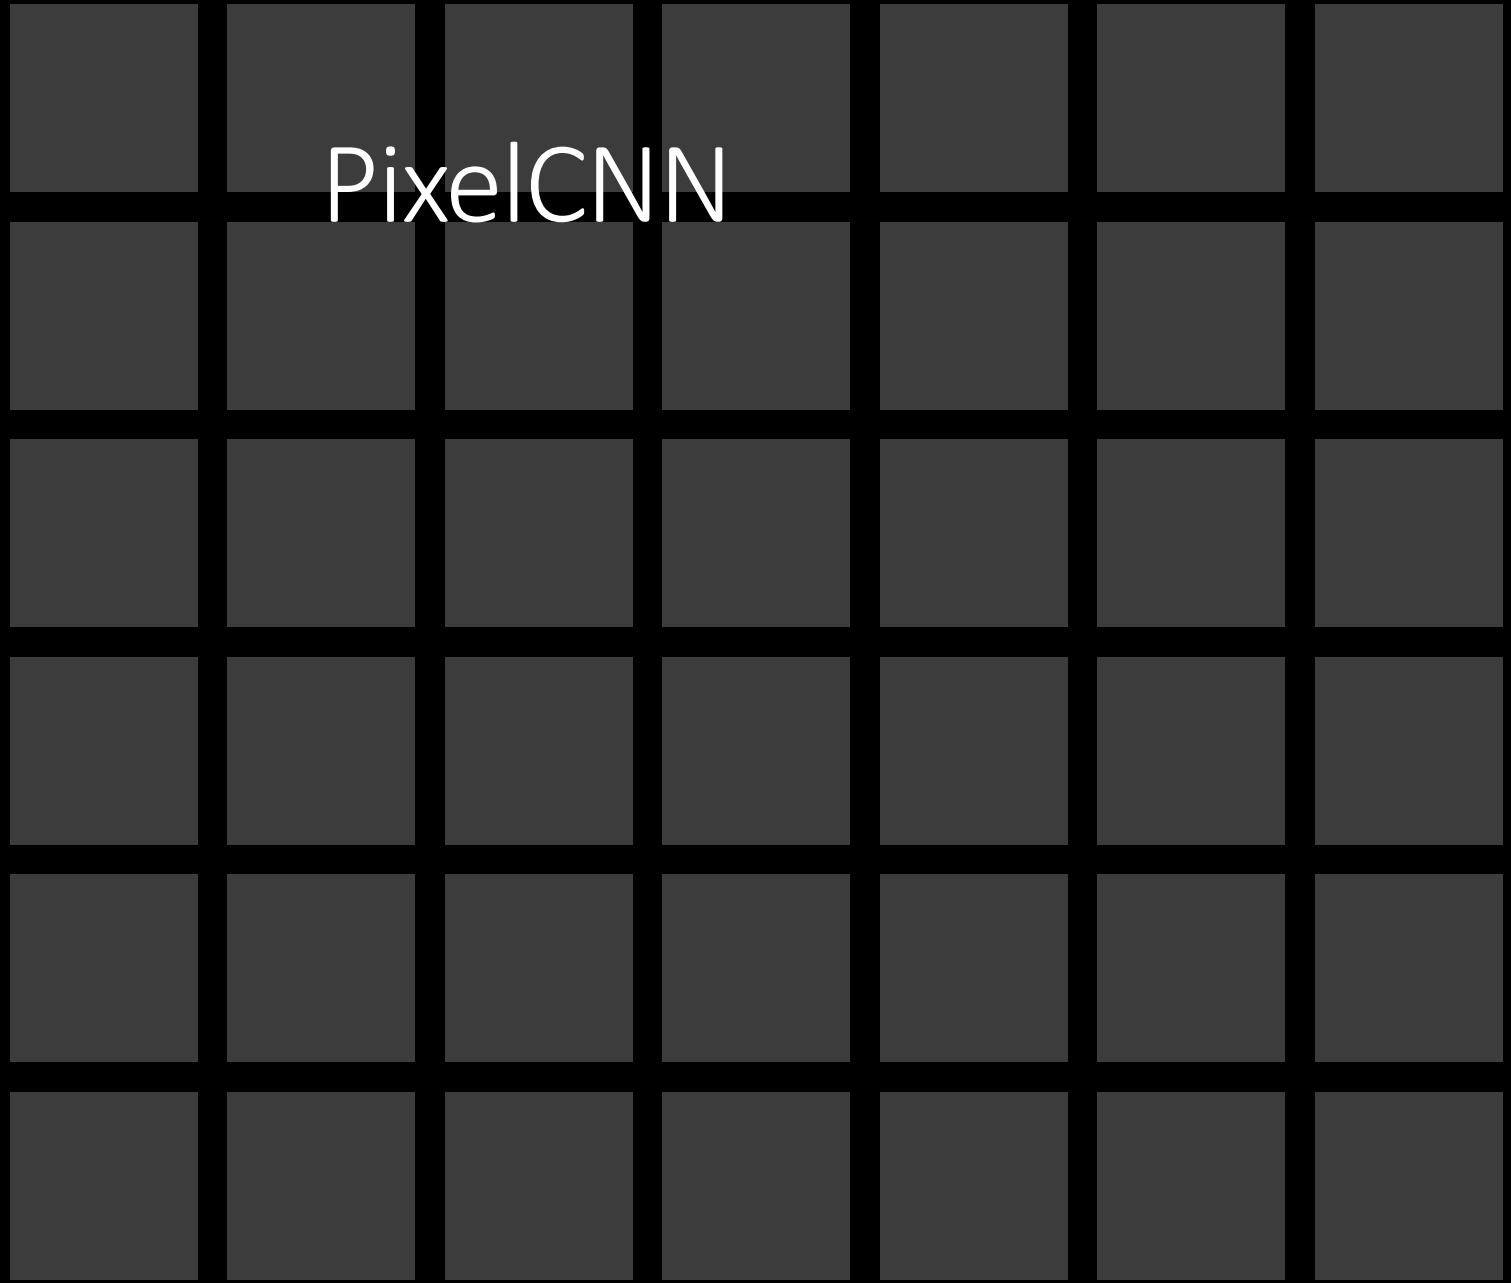


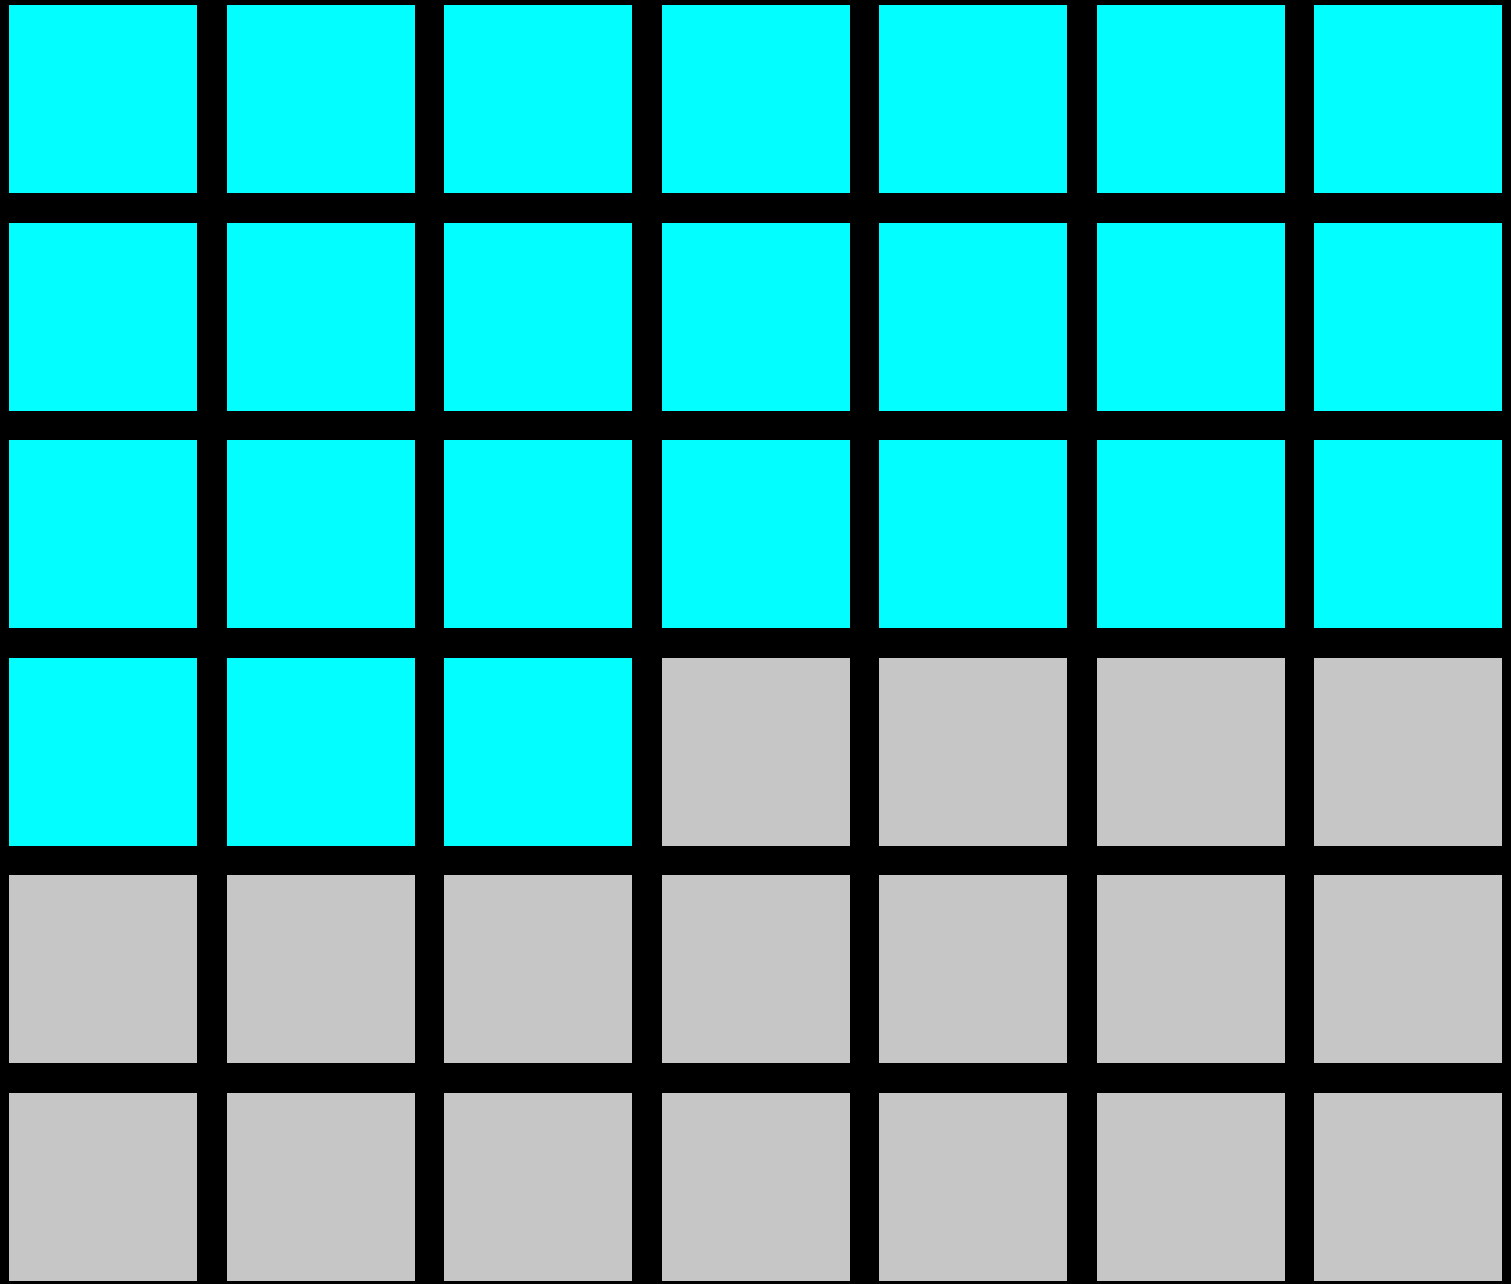


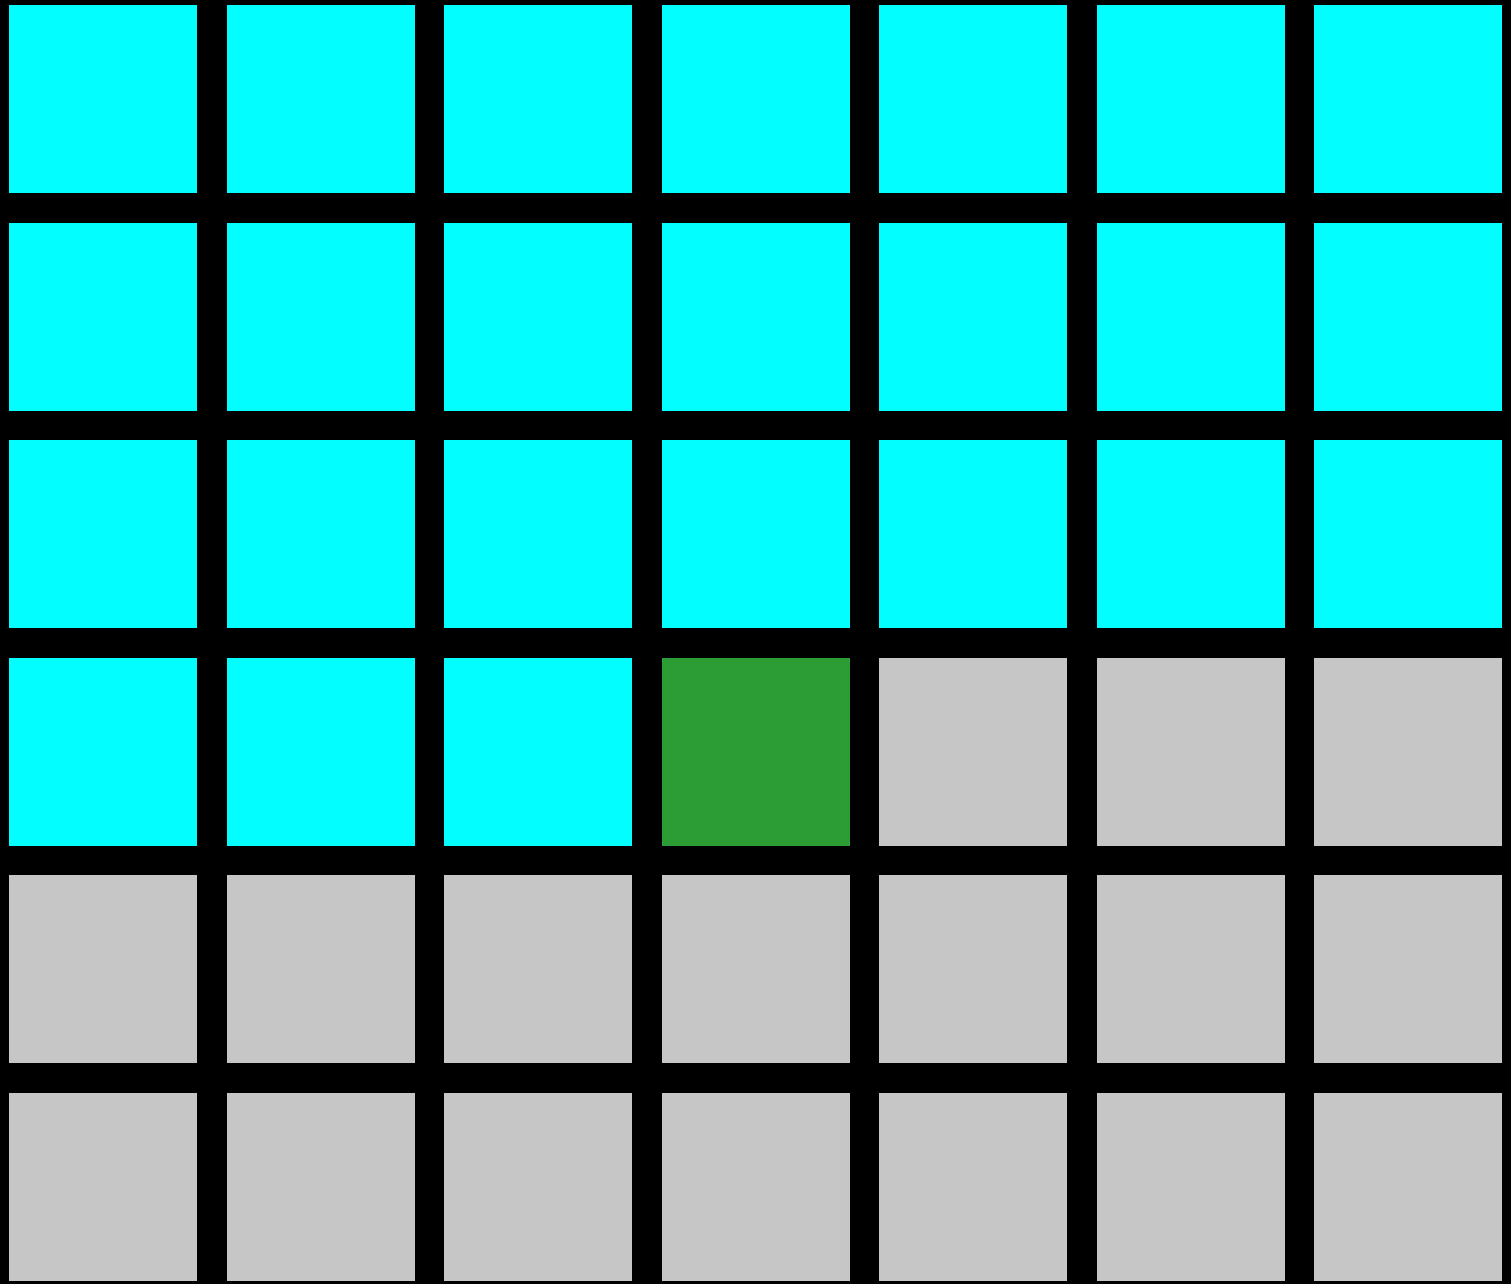


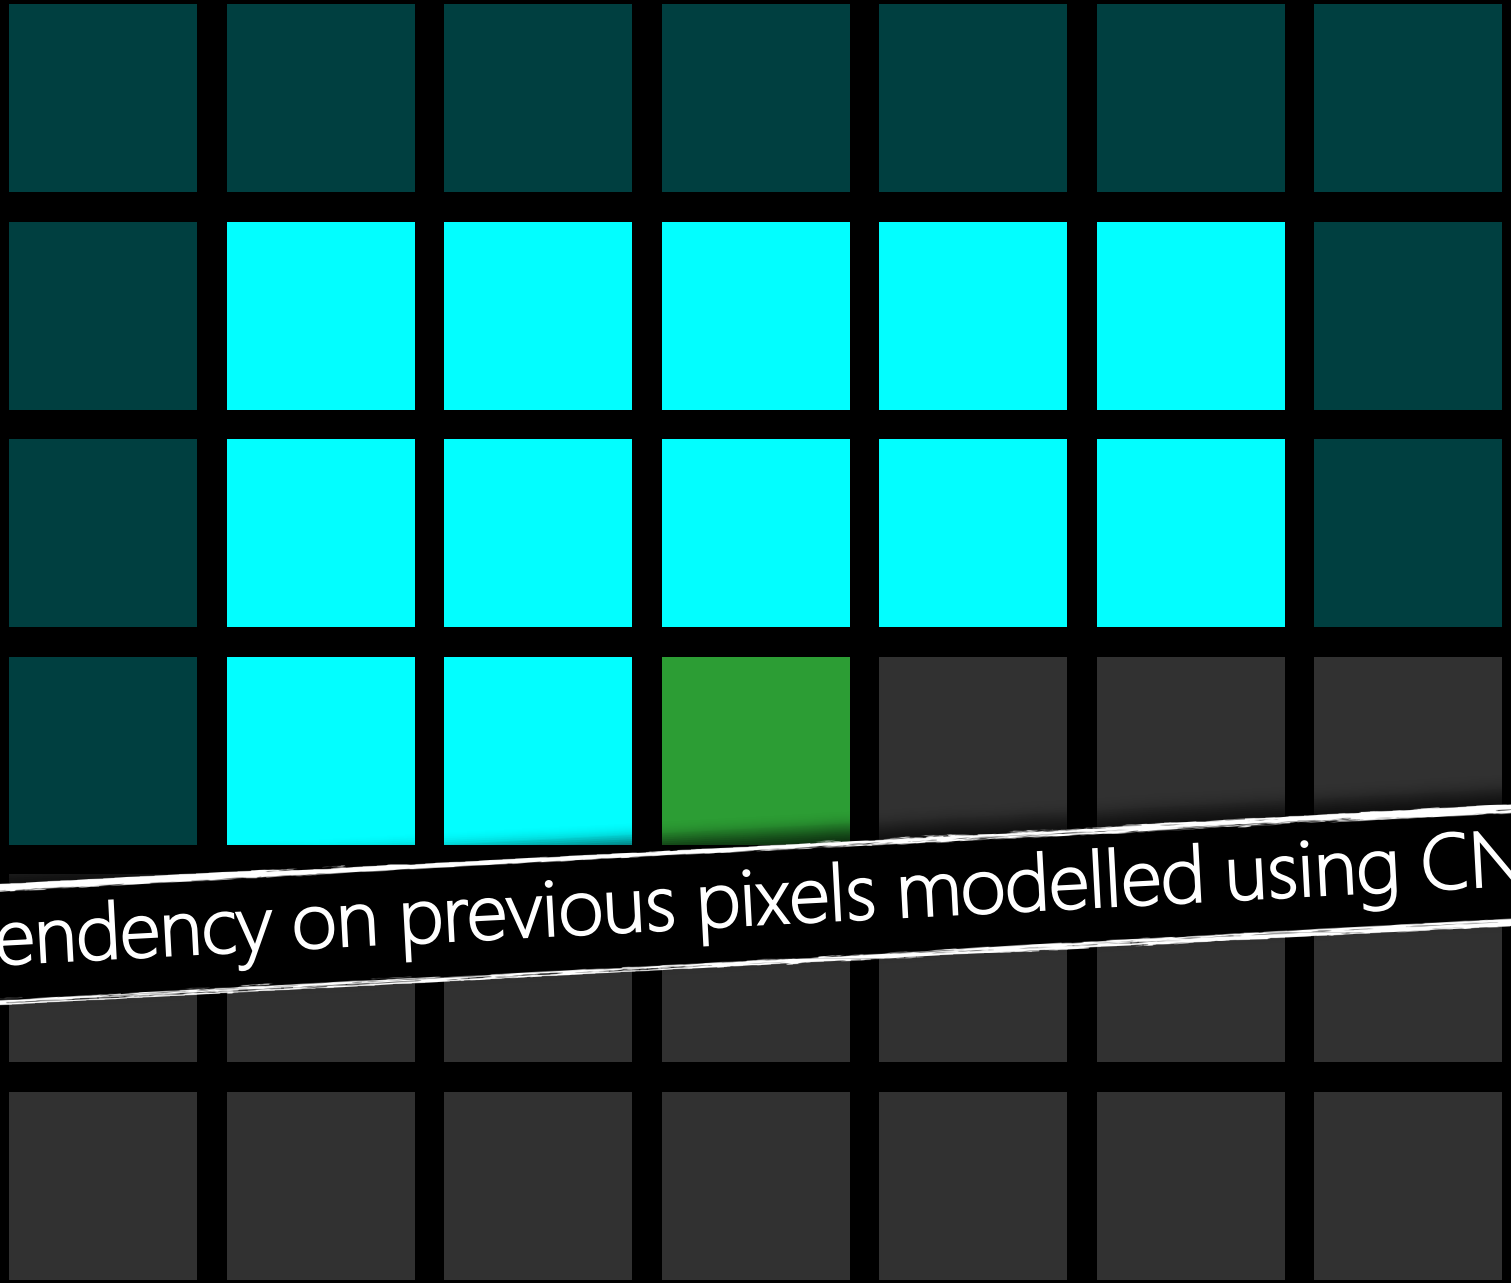


PixelCNN

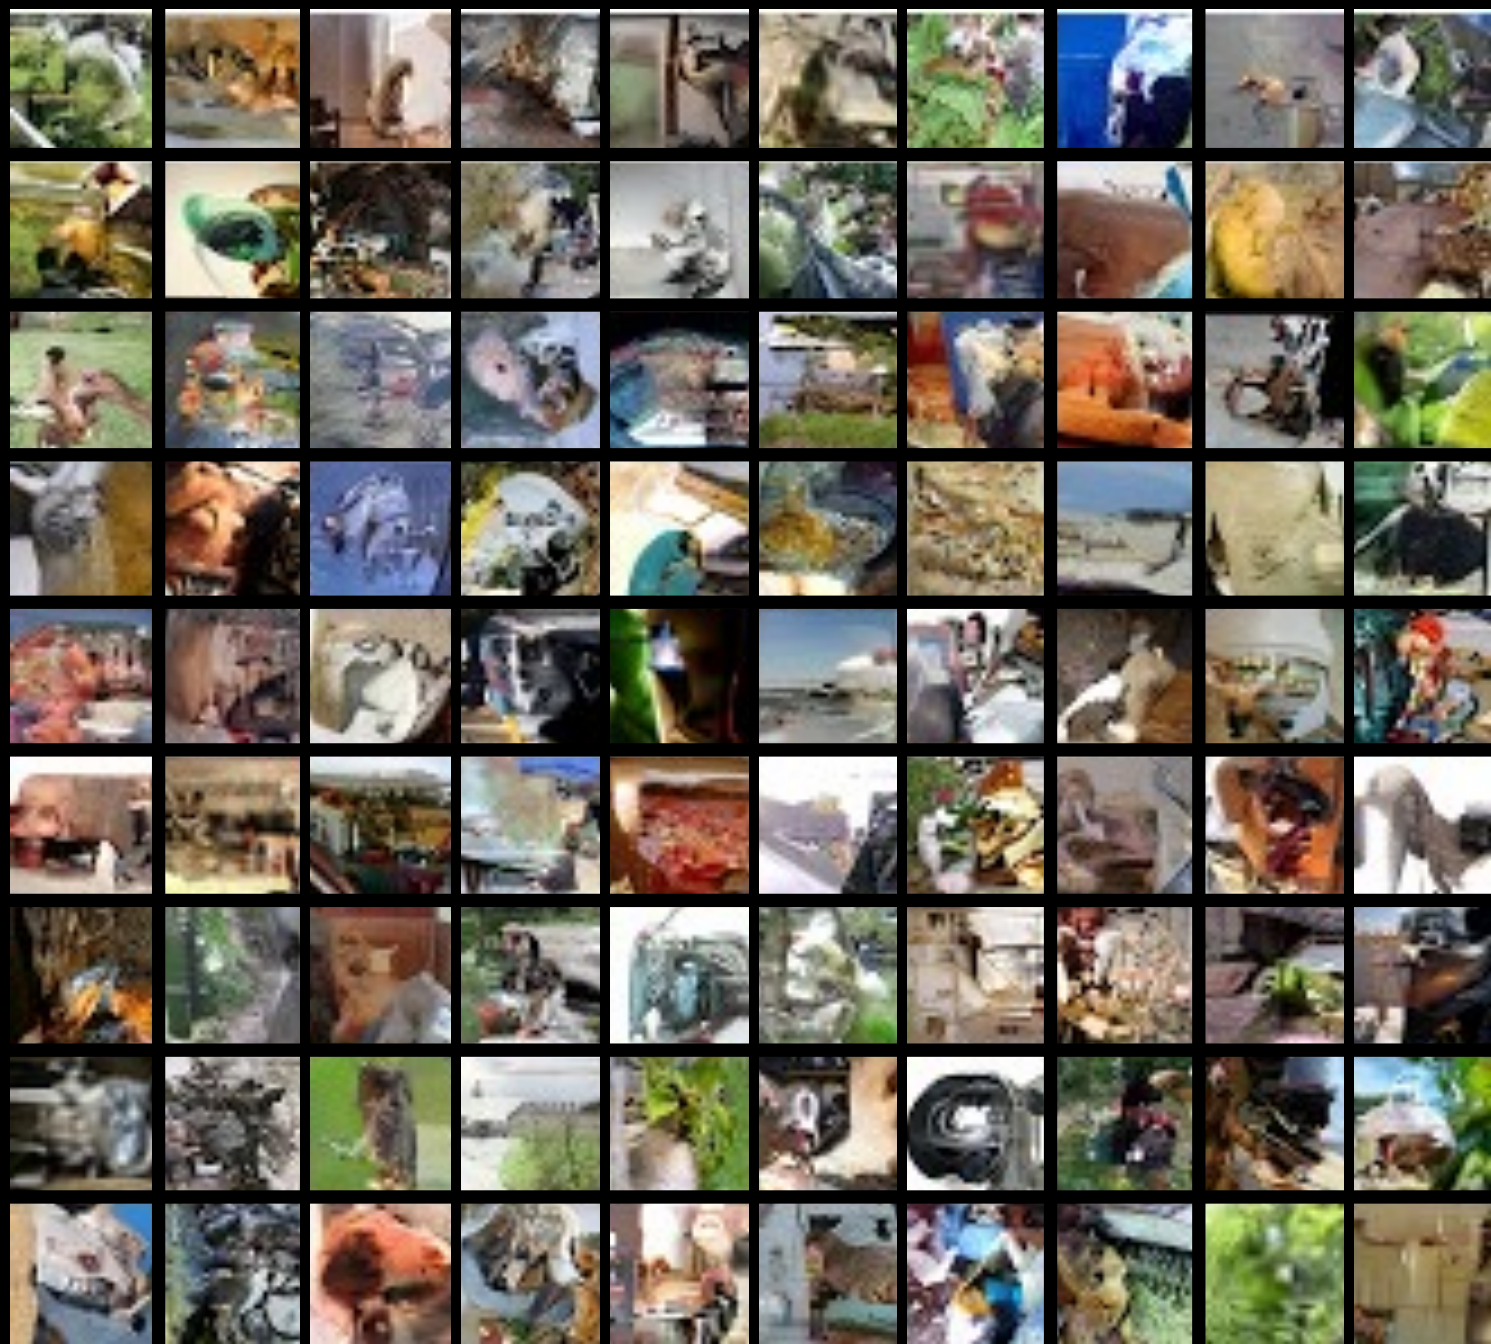


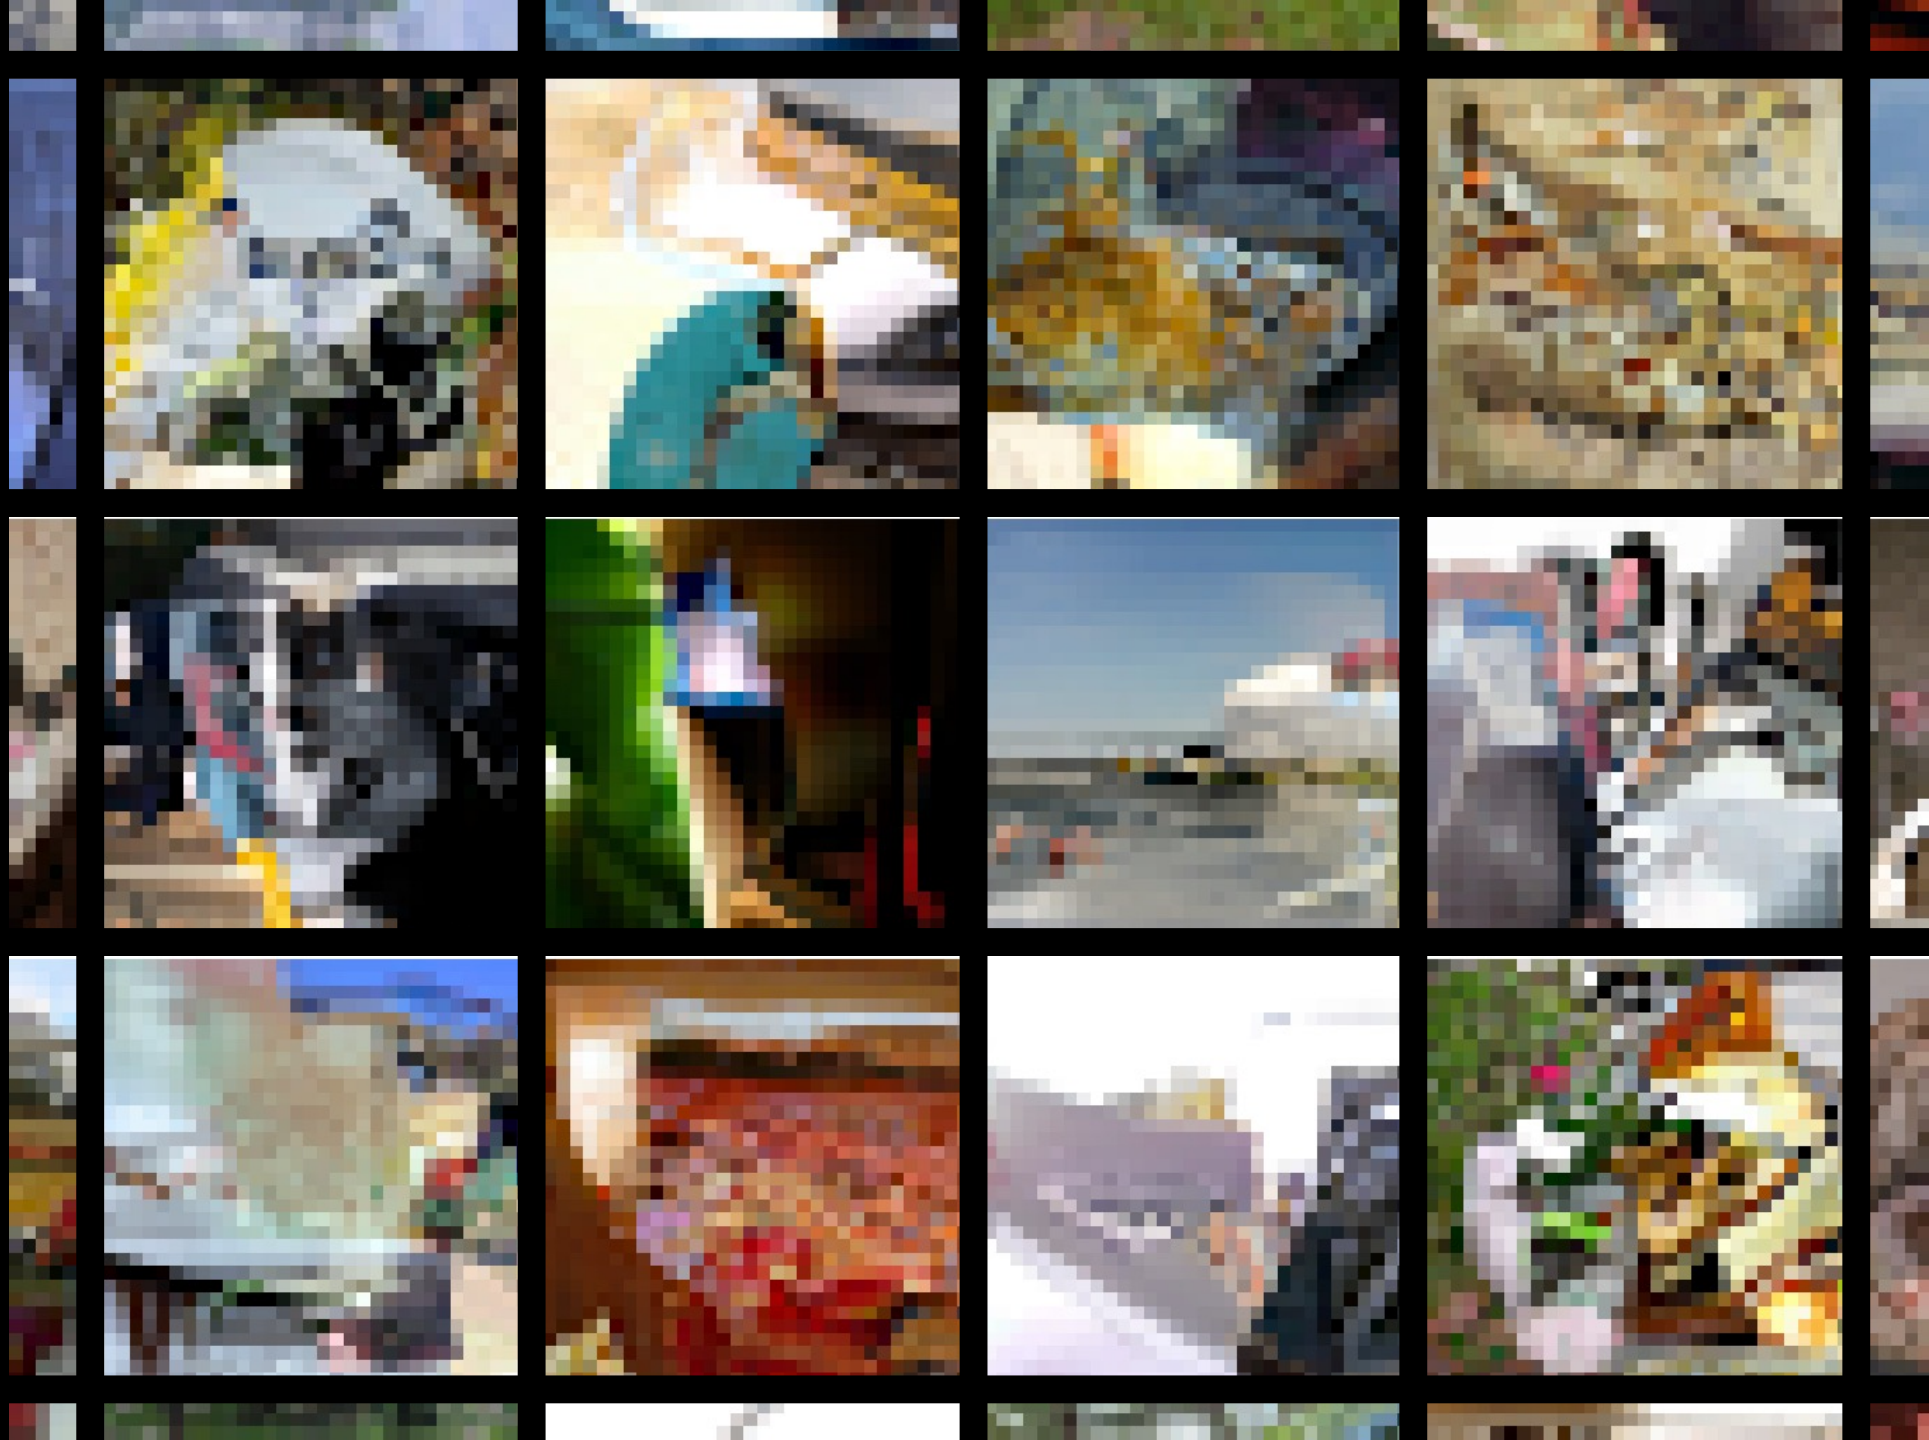




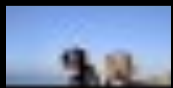
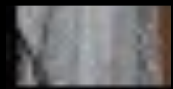


dependency on previous pixels modelled using CNN

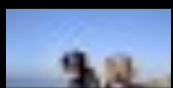




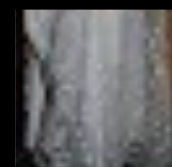
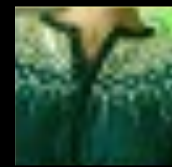
occluded



occluded



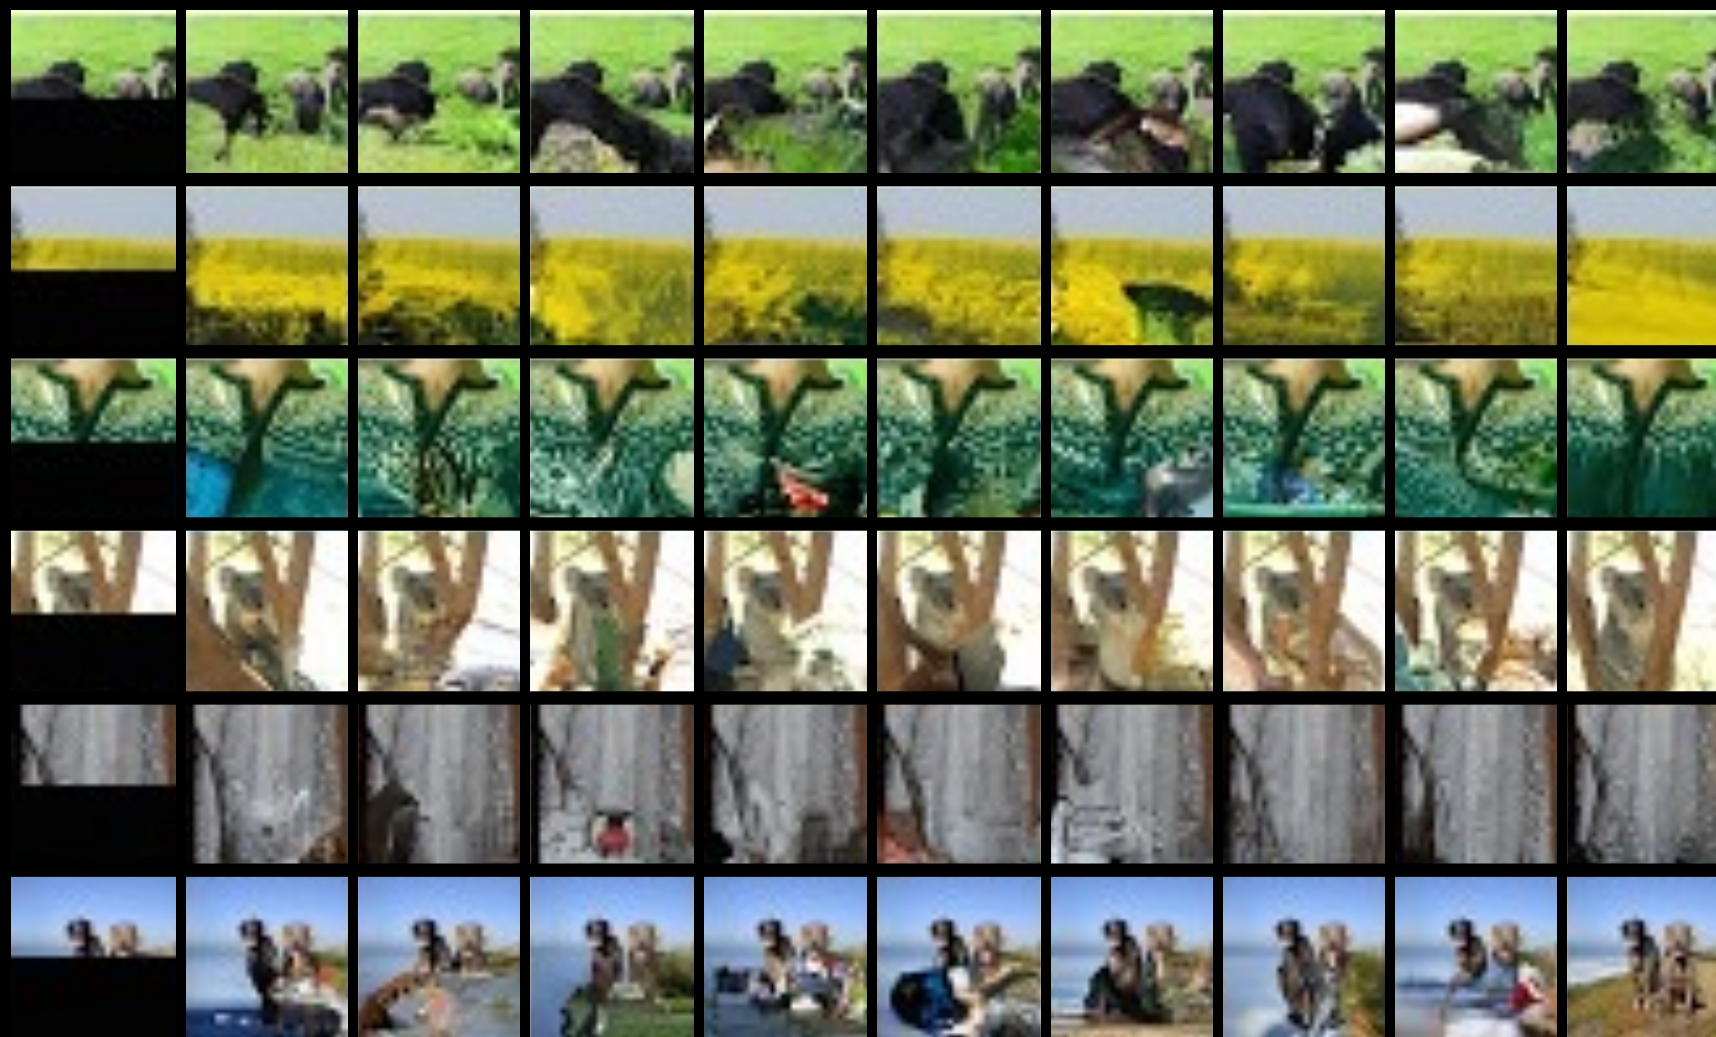
original



occluded

completions

original



sequential pixel generation is **SLOW**

pixel generation order **matters**

Explicit Density Estimation

Learn weights W

Goal: Write down an explicit function for $p(x) = f(x, W)$

Given dataset $x^{(1)}, x^{(2)}, \dots, x^{(N)}$, train the model by solving:

$$W^* = \arg \max_W \prod_i p(x^{(i)})$$

Maximize probability of training data
(Maximum likelihood estimation)

$$= \arg \max_W \sum_i \log p(x^{(i)})$$

Log trick to exchange product for sum

$$= \arg \max_W \sum_i \log f(x^{(i)}, W)$$

This will be our loss function!
Train with gradient descent

Explicit Density: Autoregressive Models

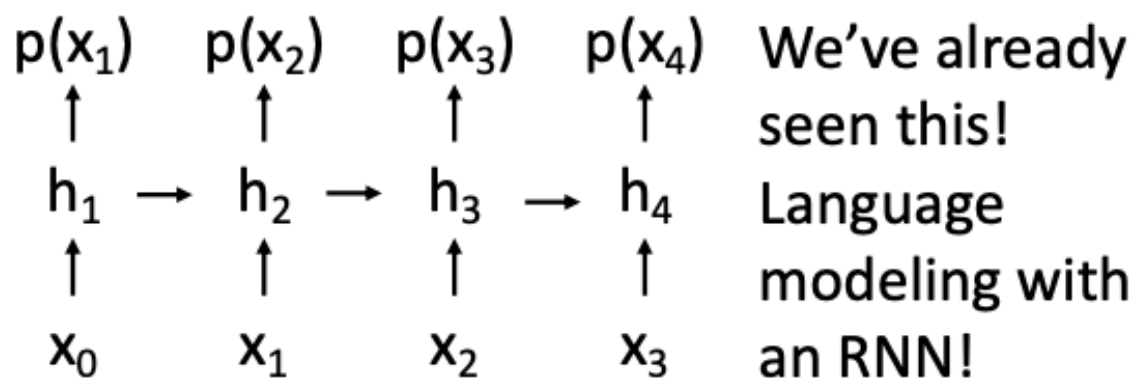
Goal: Write down an explicit function for $p(x) = f(x, W)$

Assume x consists of multiple subparts:

$$x = (x_1, x_2, x_3, \dots, x_T)$$

Break down probability using the chain rule:

$$\begin{aligned} p(x) &= p(x_1, x_2, x_3, \dots, x_T) \\ &= p(x_1)p(x_2 | x_1)p(x_3 | x_1, x_2) \dots \\ &= \prod_{t=1}^T p(x_t | x_1, \dots, x_{t-1}) \end{aligned}$$



Probability of the next subpart given all the previous subparts

Slide Credits

- EECS 6322 Deep Learning for Computer Vision, Kosta Derpanis (York University)
- EECS 498 Deep Learning for Computer Vision, Justin Johnson (U. Michigan)
- Many amazing research papers!

Extra Reading:

<https://towardsdatascience.com/auto-regressive-generative-models-pixelrnn-pixelcnn-32d192911173>