


# Deep Learning in Comp. Photo

With slides from Phil Isola, James Hays, and Andrea Vedaldi



# Visual similarity via labels



“Penguin”

?  
==



“Penguin”

# Machine Learning as data association




image  $X$



"Penguin"

label  $Y$

At test time...



black box  
classifier



?

image  $X$


# Quick Background on Supervised Learning

Jitendra Malik

3 6 8 1 7 9 6 6 9 1  
6 7 5 7 8 6 3 4 8 5  
2 1 7 9 7 1 2 8 4 5  
4 8 1 9 0 1 8 8 9 4  
7 6 1 8 6 4 1 5 6 0  
7 5 9 2 6 5 8 1 9 7  
2 2 2 2 2 3 4 4 8 0  
0 2 3 8 0 7 3 8 5 7  
0 1 4 6 4 6 0 2 4 3  
7 1 2 8 7 6 9 8 6 1

Fig. 4. Size-normalized examples from the MNIST database.



# Warm-up Example: Binary Digit Classification



vs.

# Learning Approach to Object Recognition

- Collect Training Images

- Positive: 
- Negative: 

- Training Time

- Compute **feature vectors** for positive and negative example images
- Train a **classifier**

- Test Time

- Compute feature vector on new test image:
- Evaluate classifier





Let us take an example...


Let us take an example...

In feature space, positive  
and negative examples are  
just points...

How do we classify a new  
point?

Nearest neighbor rule  
“transfer label of nearest  
example”

# Linear classifier rule






Classification units




PIT/AIT



V4/PIT



V2/V4



V1/V2



# Basic idea




Brain/Machine



“clown fish”

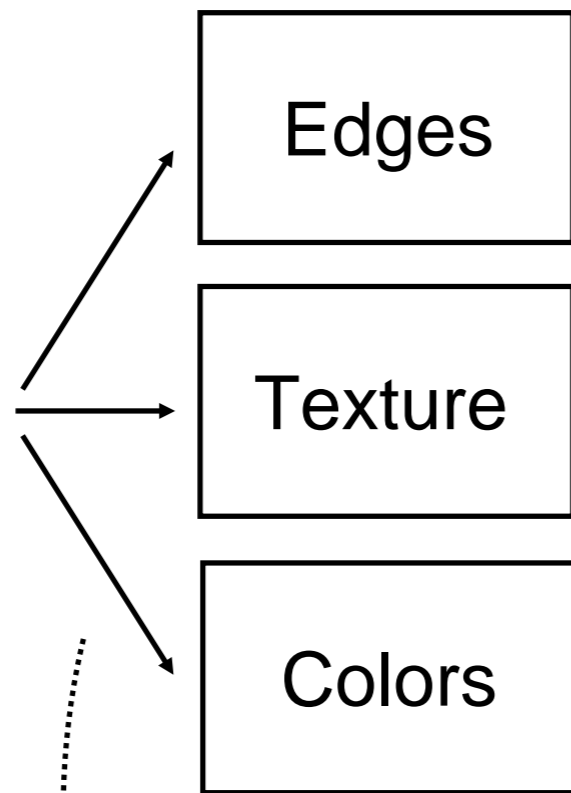



# Object recognition

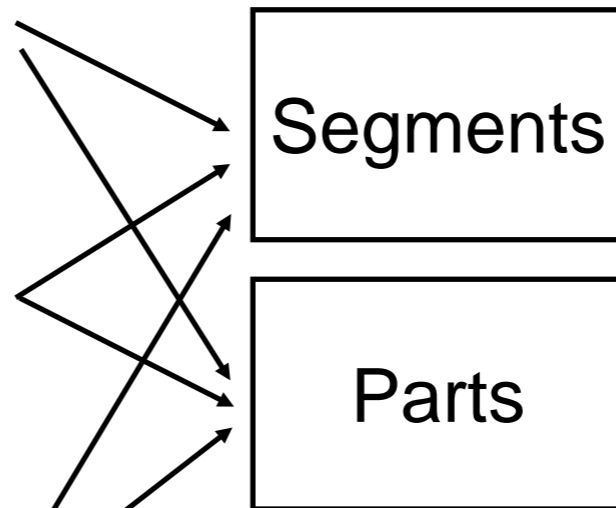


# Object recognition

Learned



Feature extractors



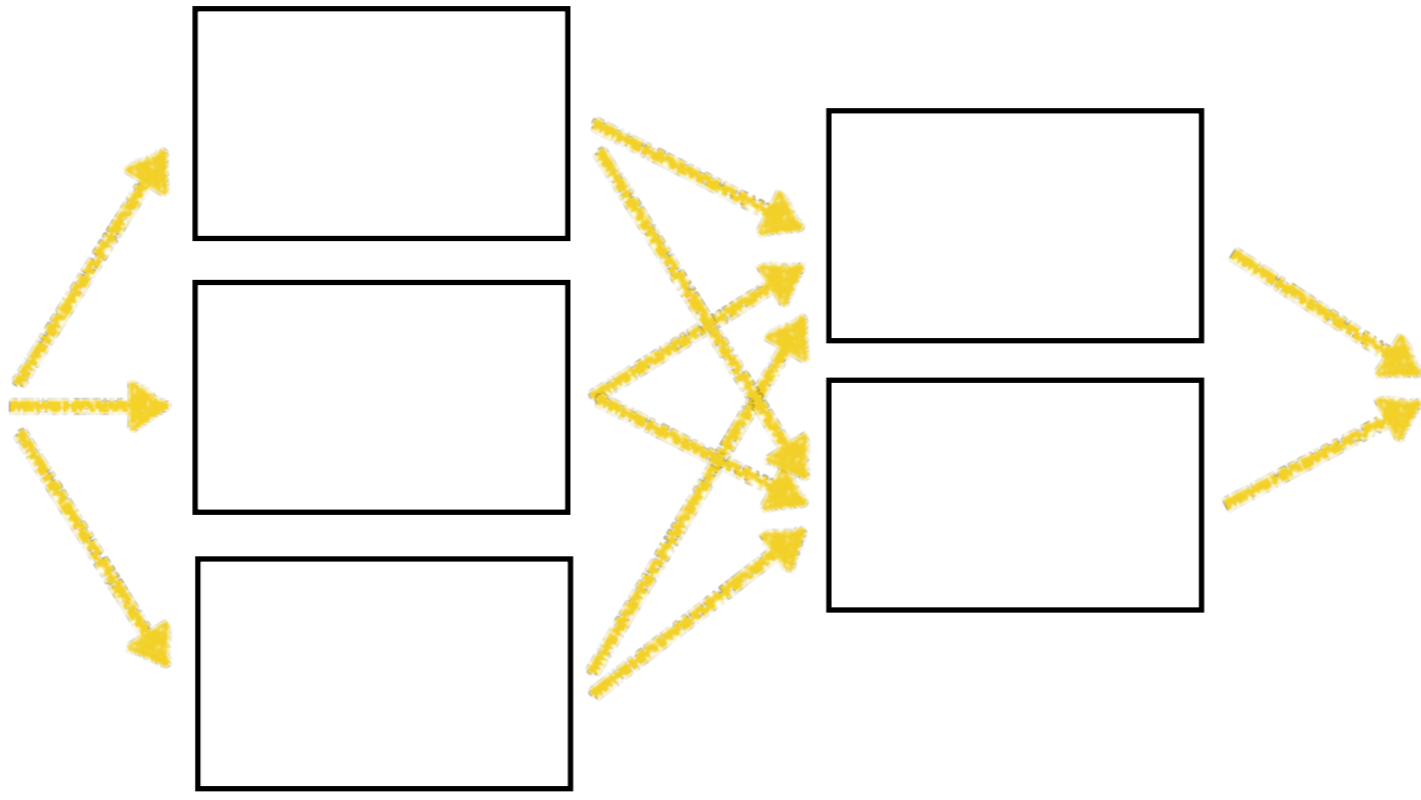
Classifier



“clown fish”

# Neural network


Learned



“clown fish”


# Neural network

Learned



# Deep neural network

Learned



# Computation in a neural net


Input  
representation



Output  
representation




# Computation in a neural net




$$y_j = \sum_i w_{ij} x_i$$

# Computation in a neural net




Rectified linear unit (ReLU)







# Convolutional Neural Nets




# Convolutional Neural Nets




## Convolution



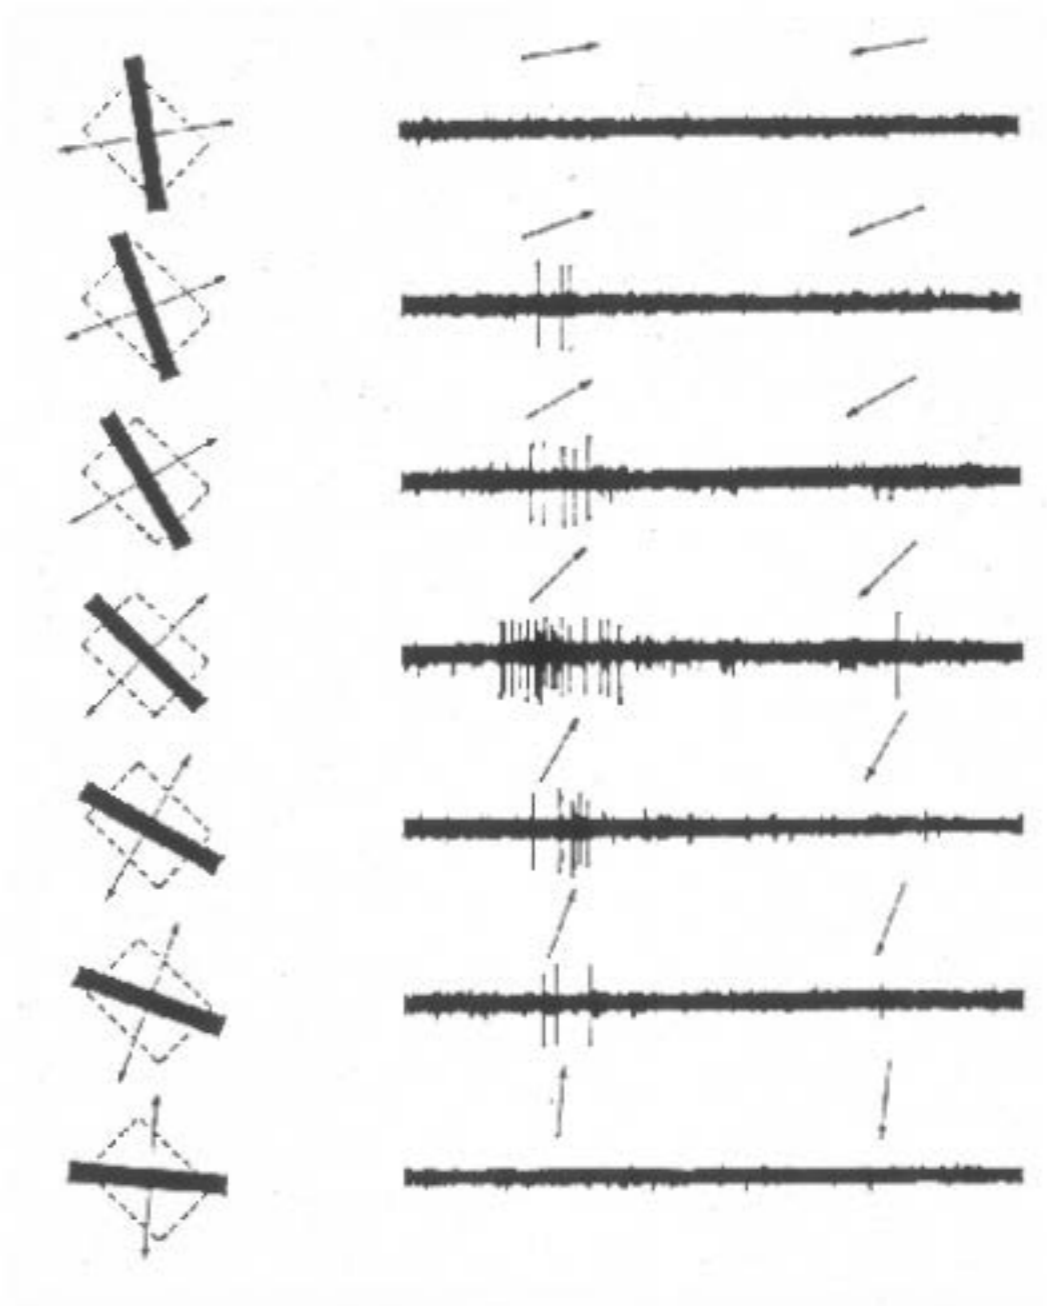
filter




# [Hubel and Wiesel 59]



oriented filter



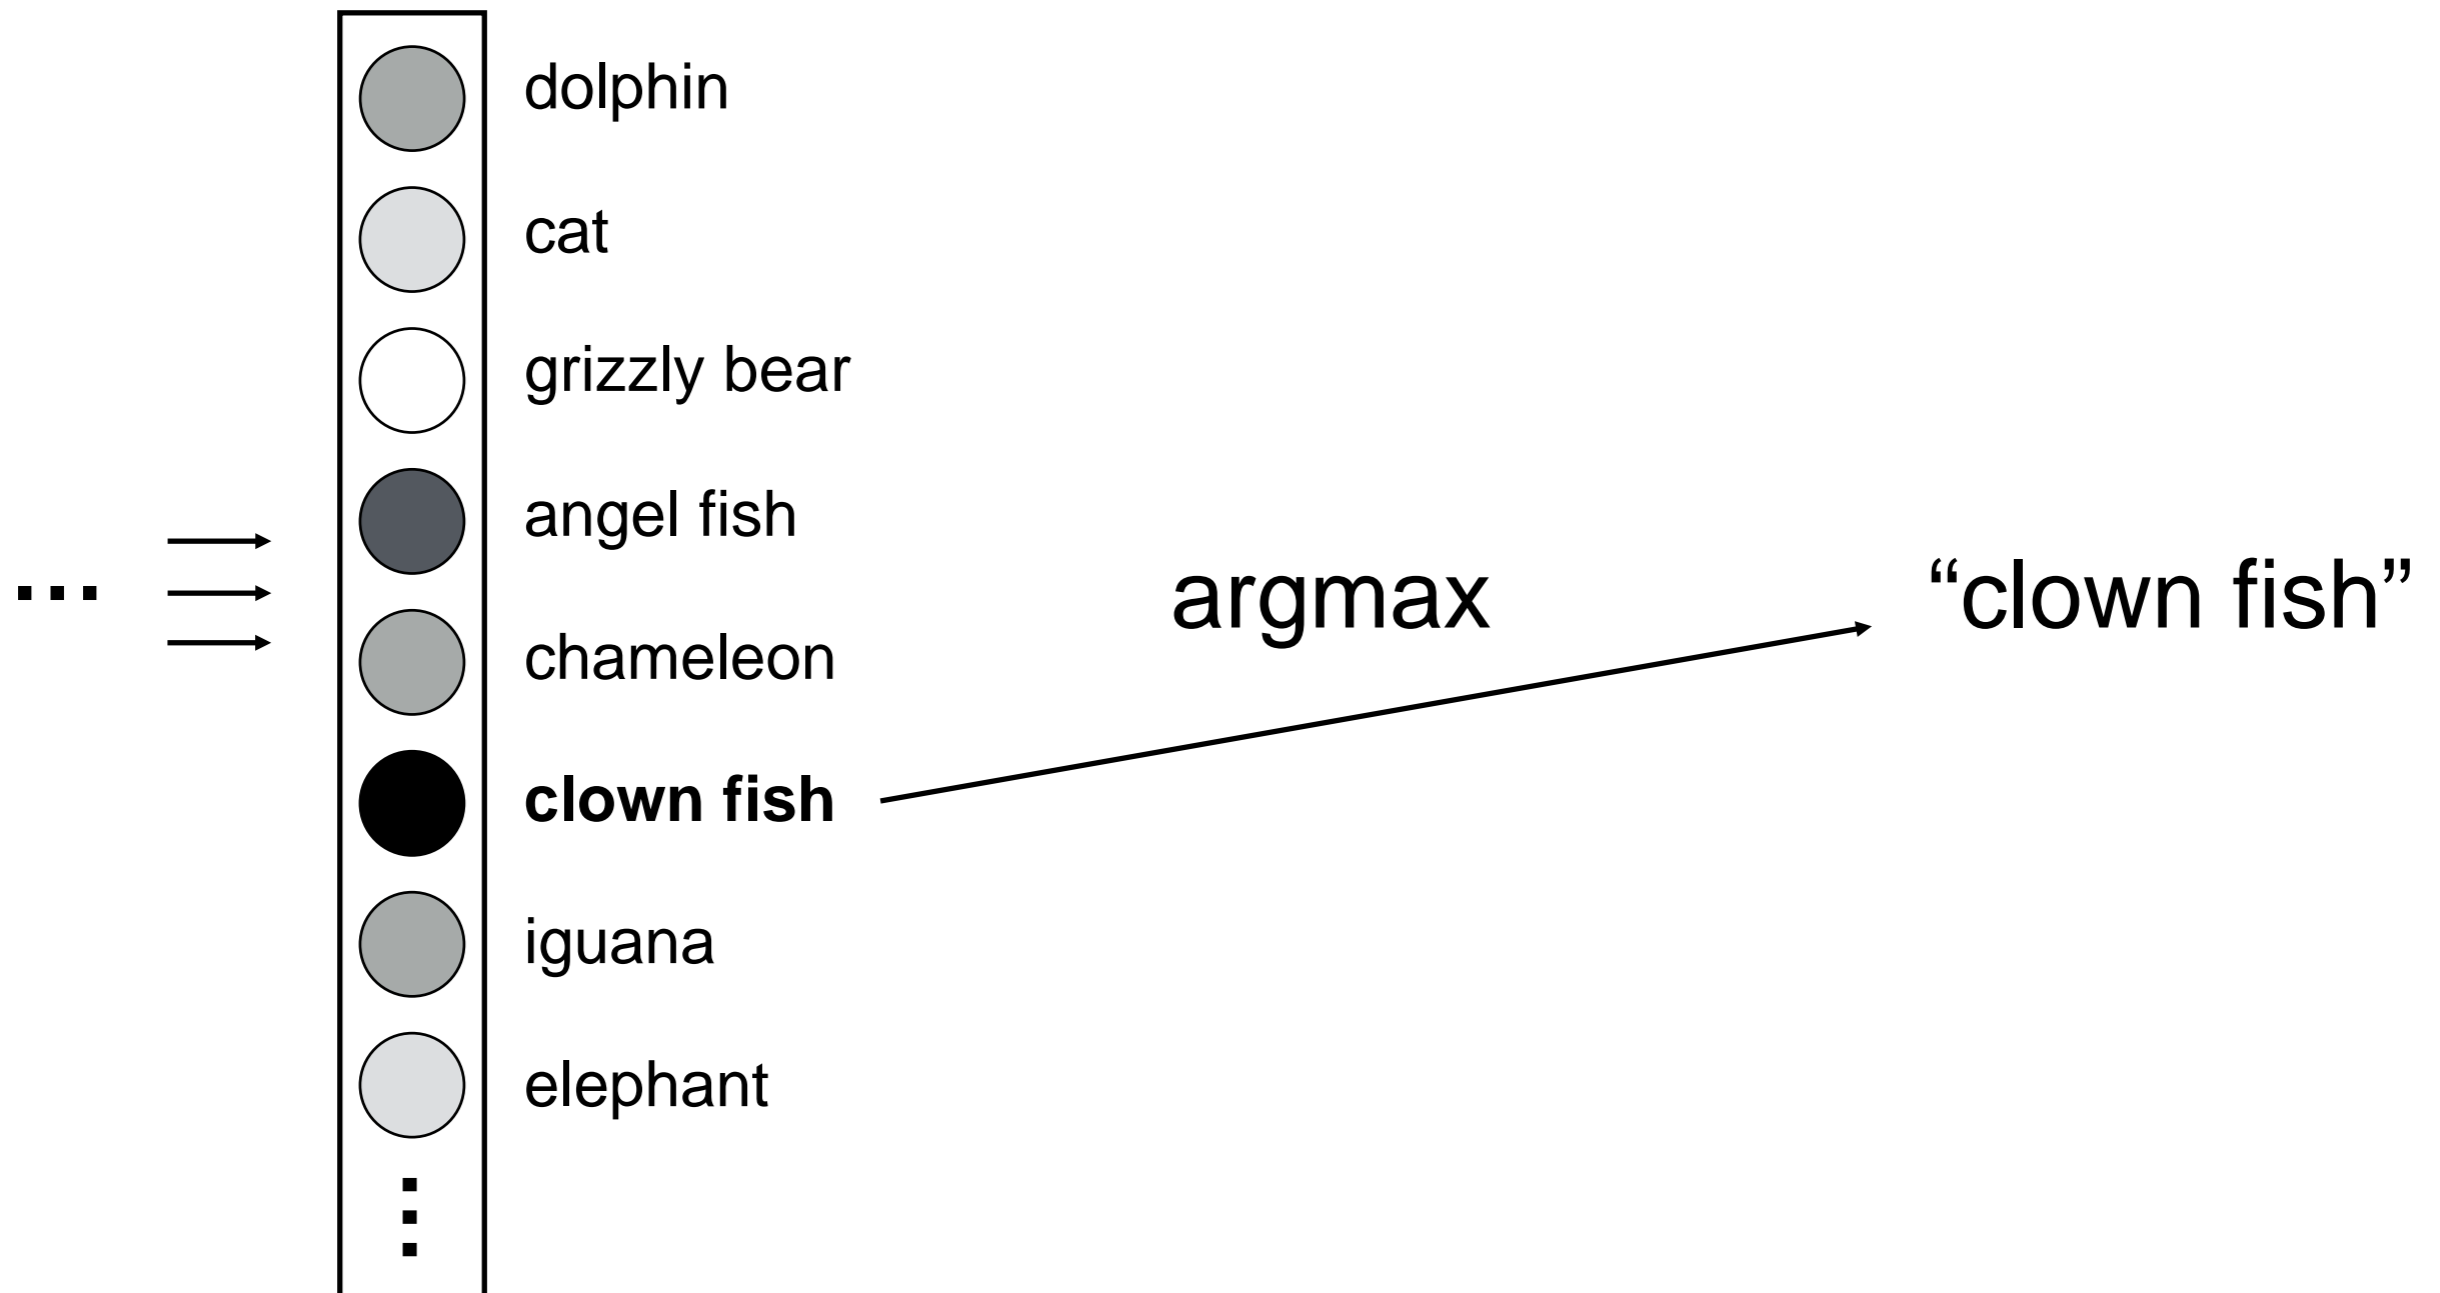
# Computation in a neural net



$$f(\mathbf{x}) = f_L(\dots f_2(f_1(\mathbf{x})))$$

# Computation in a neural net

Last layer






Classification units



PIT/AIT




V4/PIT



V2/V4




V1/V2



# Learning with deep nets

Learned



# Learning with deep nets



→ “clown fish”



→ “grizzly bear”



→ “chameleon”



Train network to  
associate the right  
label with each image



# Learning with deep nets

Learned

“clown fish”




$$\underset{\mathbf{w}}{\operatorname{argmin}} L(w_1, \dots, w_6)$$

# Loss function

Network output


Ground truth label



# Loss function

Network output

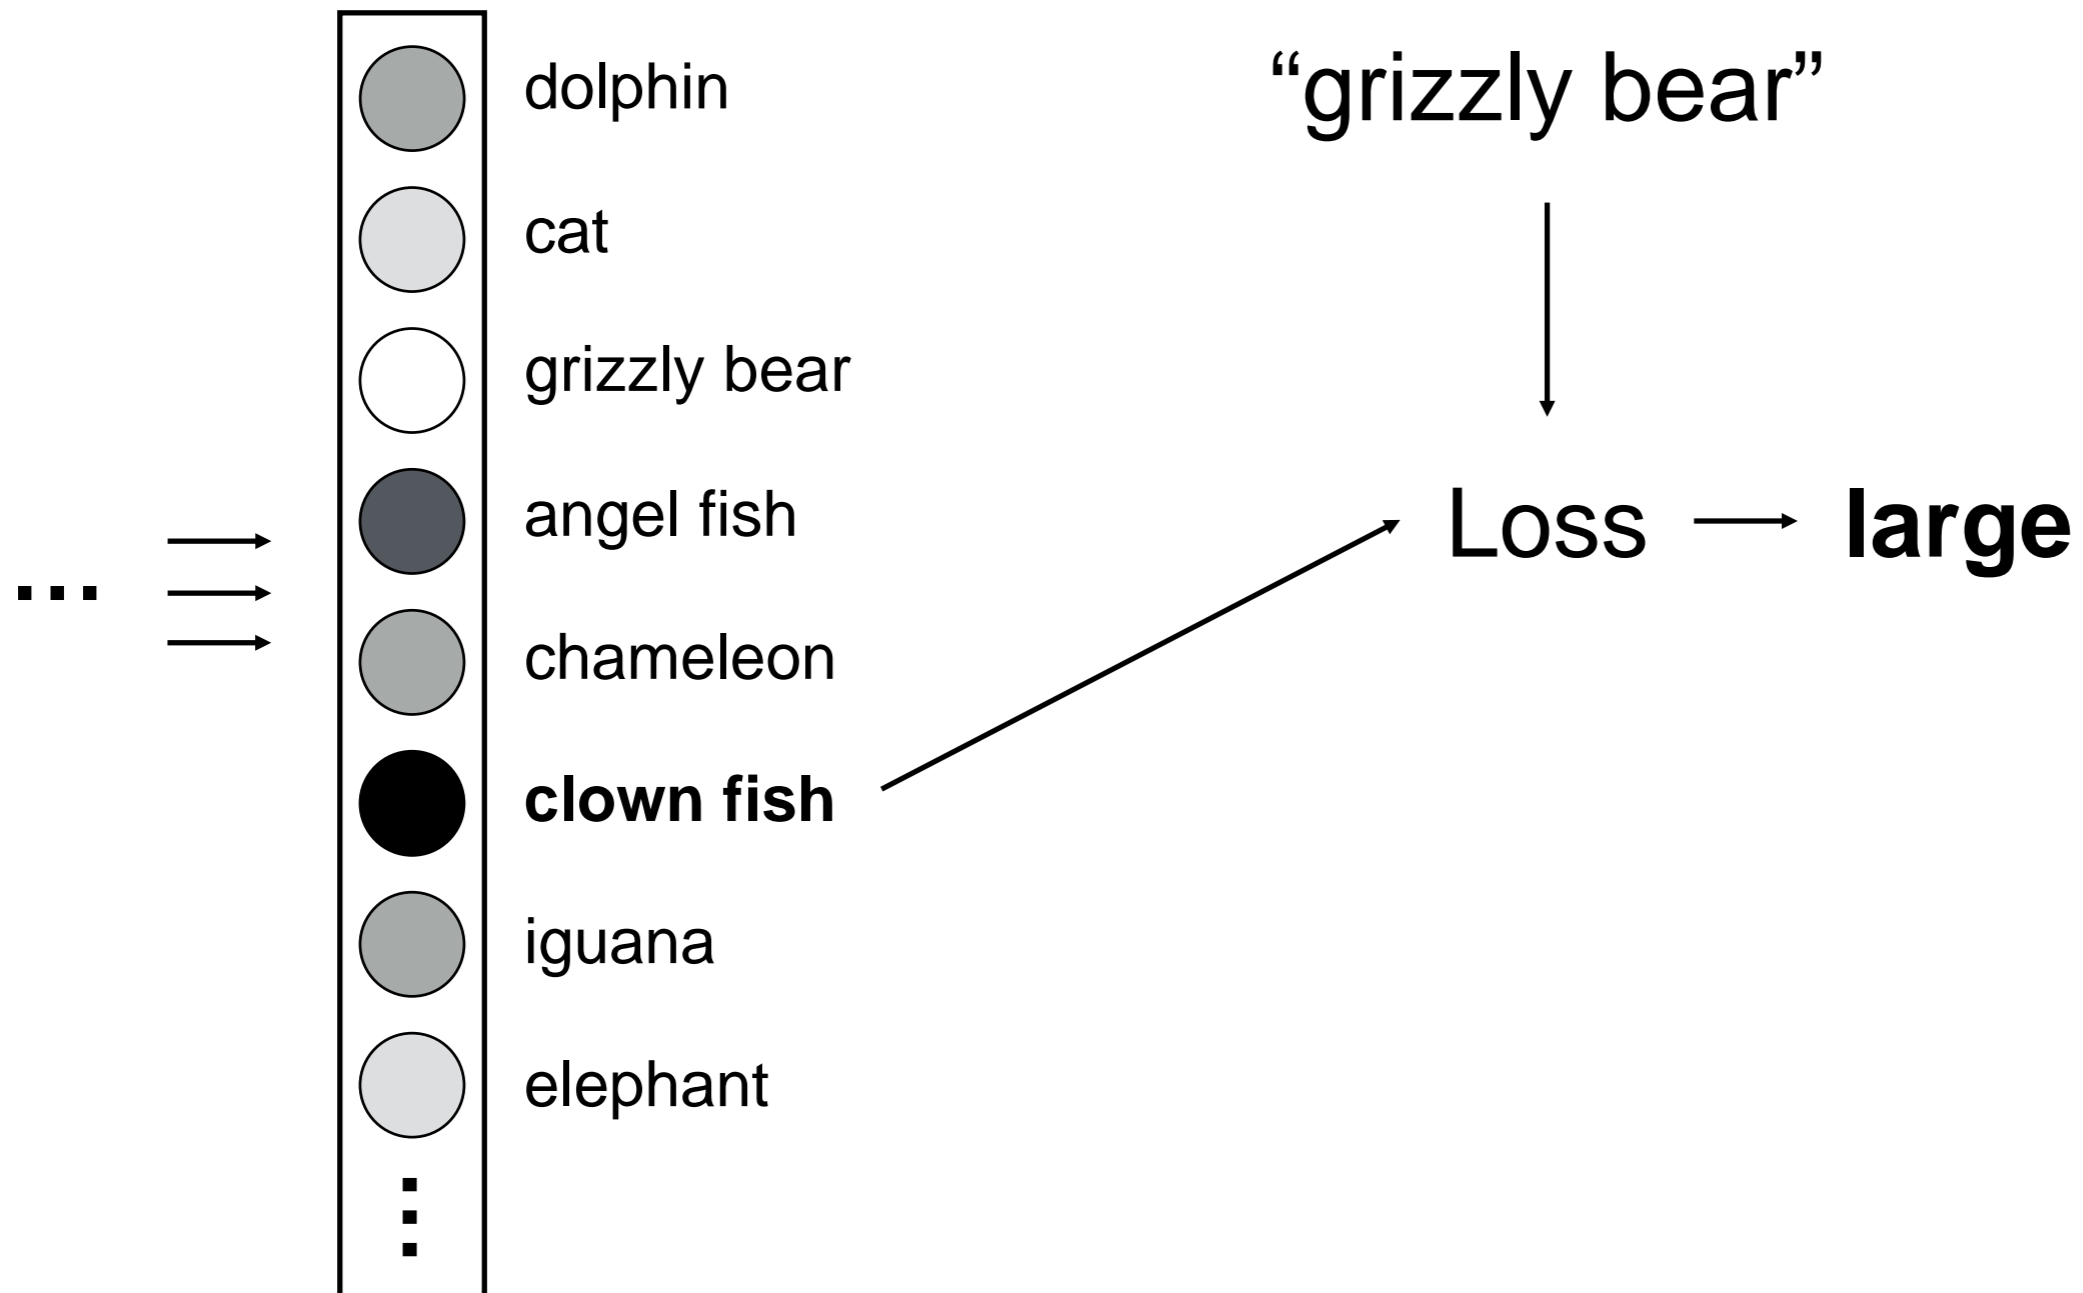
Ground truth label



# Loss function

Network output

Ground truth label



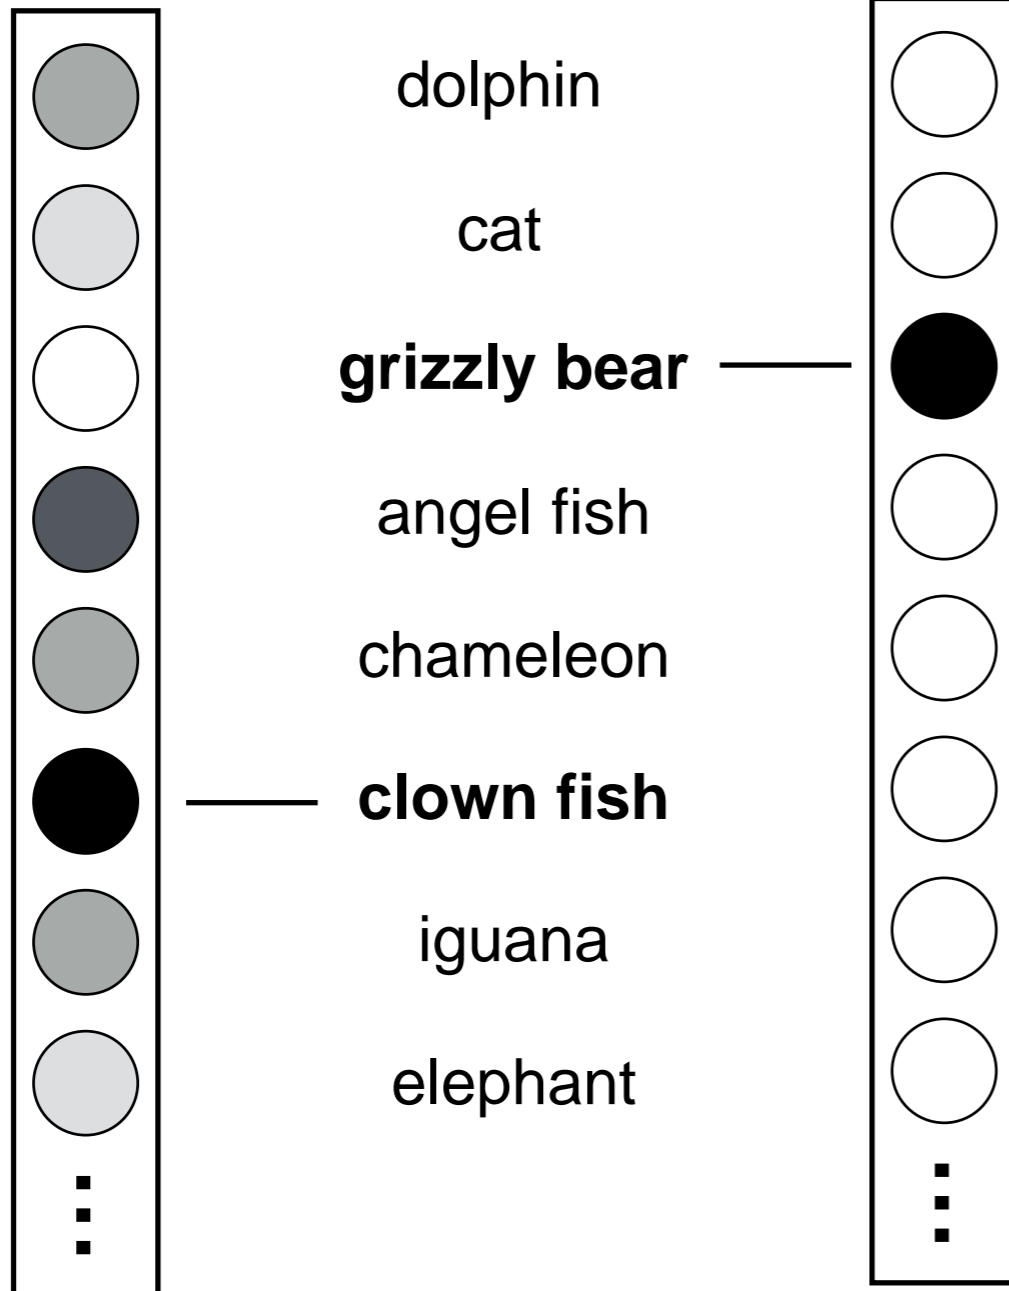
# Loss function for classification

Network output

Ground truth label

$\hat{\mathbf{z}}$

$\mathbf{z}$



**Probability of the  
observed data under  
the model**

$$H(\hat{\mathbf{z}}, \mathbf{z}) = - \sum_c \hat{\mathbf{z}}_c \log \mathbf{z}_c$$



*Results in learning a  
probability model  $p(c|\mathbf{x})!$*

# Learning with deep nets

Learned

$\mathbf{z}_1$   
“clown fish”

$\mathbf{x}_1$




# Learning with deep nets

Learned

$\mathbf{z}_2$   
“grizzly bear”




$\mathbf{x}_2$



# Learning with deep nets

Learned

$\mathbf{z}_i$   
“chameleon”



$$\operatorname{argmin}_{\mathbf{w}} \sum_i l(\mathbf{z}_i, f(\mathbf{x}_i; \mathbf{w}))$$



# Gradient descent

$$\operatorname{argmin}_{\mathbf{w}} \sum_i \ell(\mathbf{z}_i, f(\mathbf{x}_i; \mathbf{w})) = L(\mathbf{w})$$


One iteration of gradient descent:

$$\mathbf{w}^{t+1} = \mathbf{w}^t - \eta_t \frac{\partial L(\mathbf{w}^t)}{\partial \mathbf{w}}$$

learning rate

# Gradient descent

$L(\mathbf{w})$



$$p(c|\mathbf{x})$$



**mite**

**container ship**

**motor scooter**

**leopard**



**grille**



**mushroom**




**cherry**



**Madagascar cat**




# Texture synthesis by non-parametric sampling



Models  $P(p|N(p))$


[Efros & Leung 1999]

# Texture synthesis with a deep net



[van der Oord et al. 2016]


# Texture synthesis with a deep net



[van der Oord et al. 2016]

# Sampling



## Network output



$$P(p_i | p_1, \dots, p_{i-1})$$

# Sampling



## Network output







# Sampling

## Network output




# Sampling

## Network output




# Sampling

## Network output

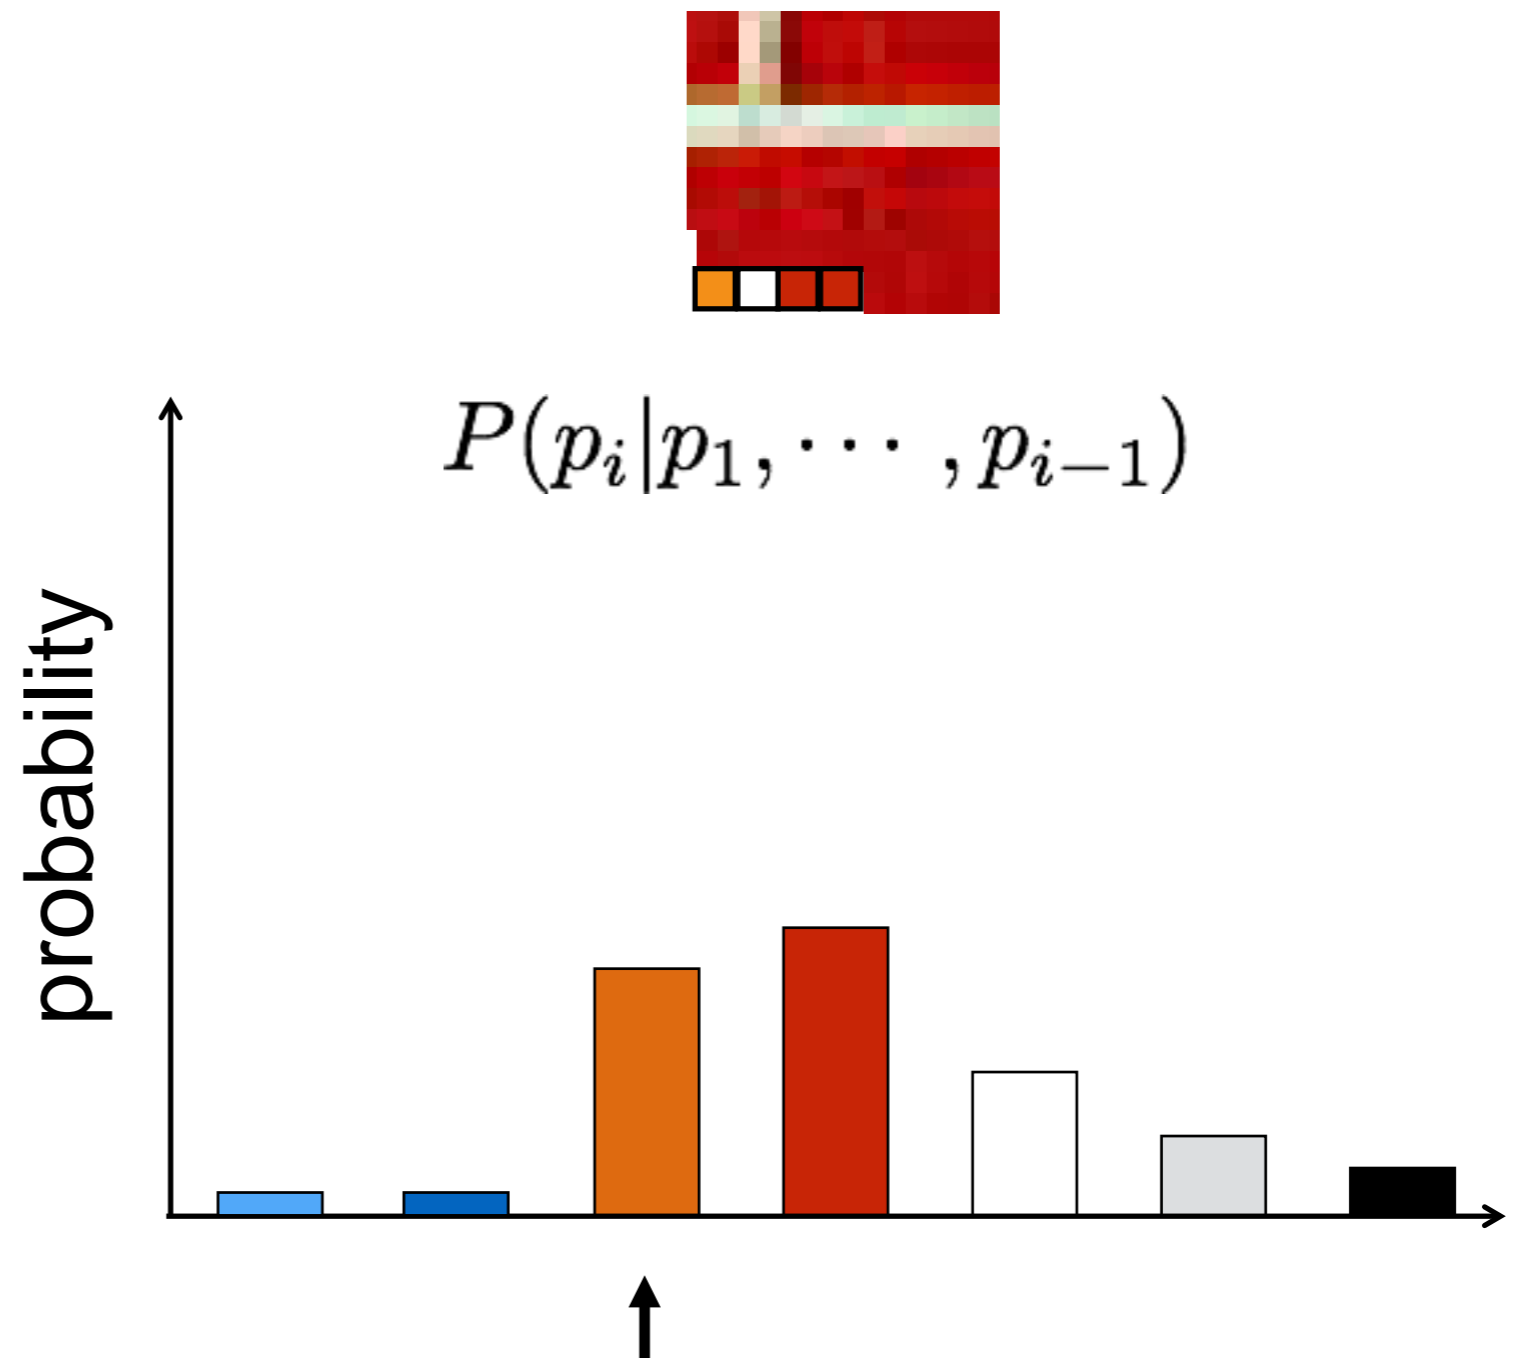



$$P(p_i | p_1, \dots, p_{i-1})$$



# Sampling

## Network output










# Inpainting

occluded

completions

original





■ “yellow”



...