
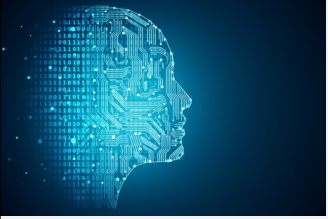


**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute



# Data Science for Researchers and Scholars

**Vasant G. Honavar**  
 Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence  
 Professor of Data Sciences, Informatics, Computer Science and Engineering, Bioinformatics & Genomics,  
 Public Health Sciences and Neuroscience  
 Director, Center for Artificial Intelligence Foundations and Scientific Applications  
 Associate Director, Institute for Computational and Data Sciences  
 Pennsylvania State University

vhonavar@psu.edu  
<http://faculty.ist.psu.edu/vhonavar>  
<http://ailab.ist.psu.edu>




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

1




**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## From Neural Networks to Deep Neural Networks

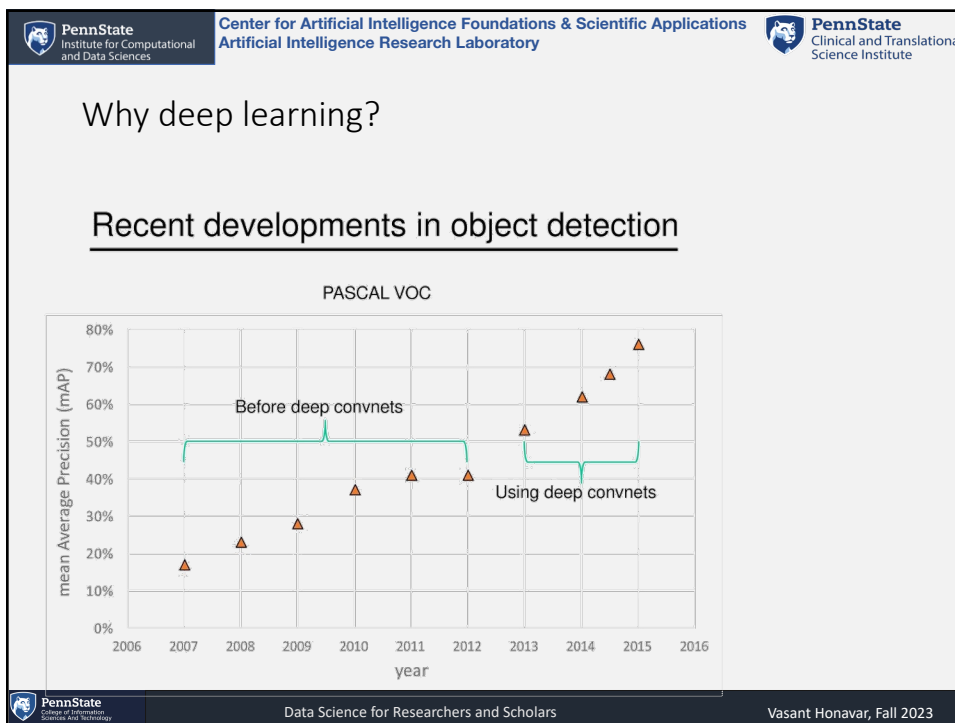


**PennState**  
College of Information  
Science and Technology

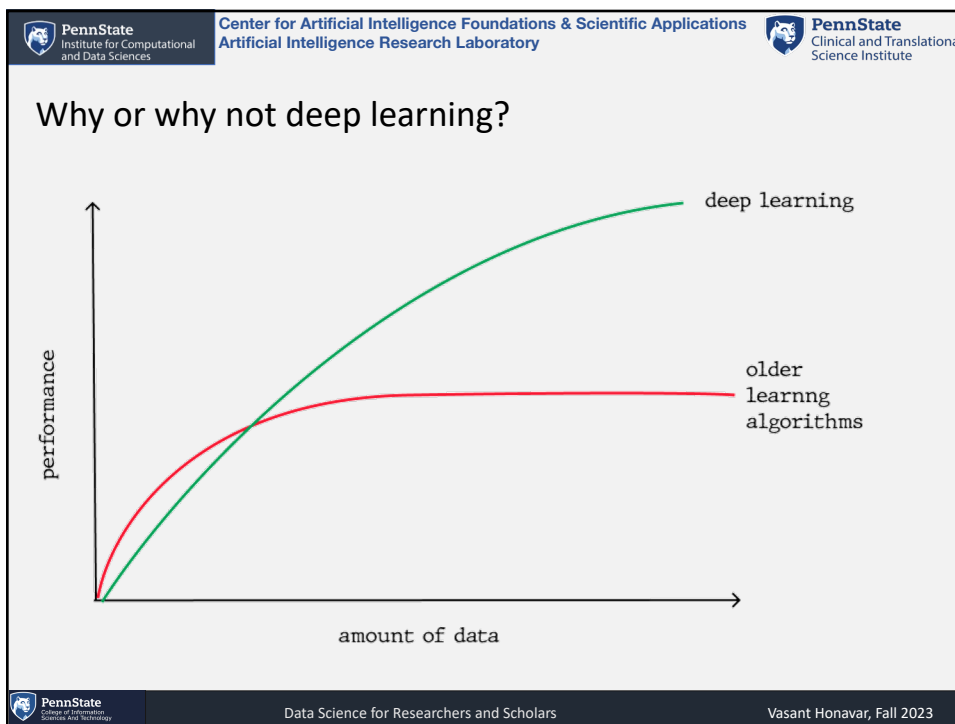
Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023


2



3




4



**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Why not deep learning?


### Carbon footprint comparison

Source: Strubell et al, 2019.

CO2 emissions (lbs)

Training Transformer (big) w/ neural architecture search	626,155
Car, avg incl. fuel, 1 lifetime	126,000
Human life, avg, 1 year	11,023
Air travel, 1 passenger, NY<-> SF	1,984
Training BERTbase on GPU	1,438

*Reconstructed from: <http://arxiv.org/abs/1906.02243>*




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

5




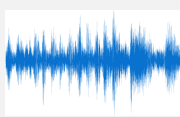

**PennState**  
Institute for Computational  
and Data Sciences


**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Typical machine learning scenarios

input		output
images/video		Image label Descriptive tags Similar images ...
audio		Speech recognition Music categories Speaker identity ...
text		Sentiment detection Text categories Translation ...

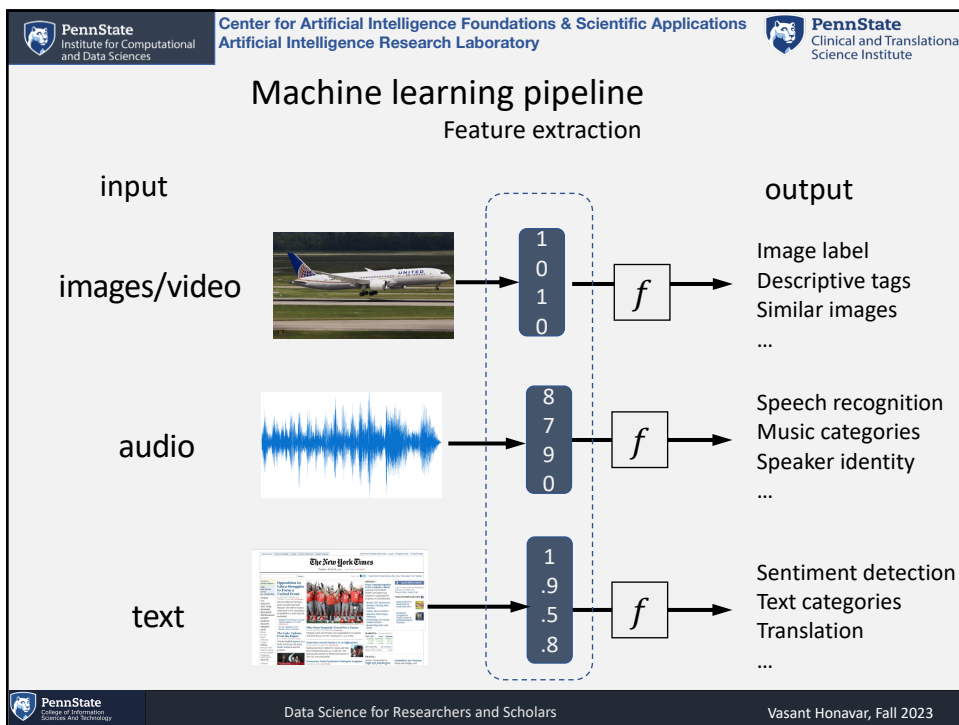


**PennState**  
College of Information  
Science and Technology

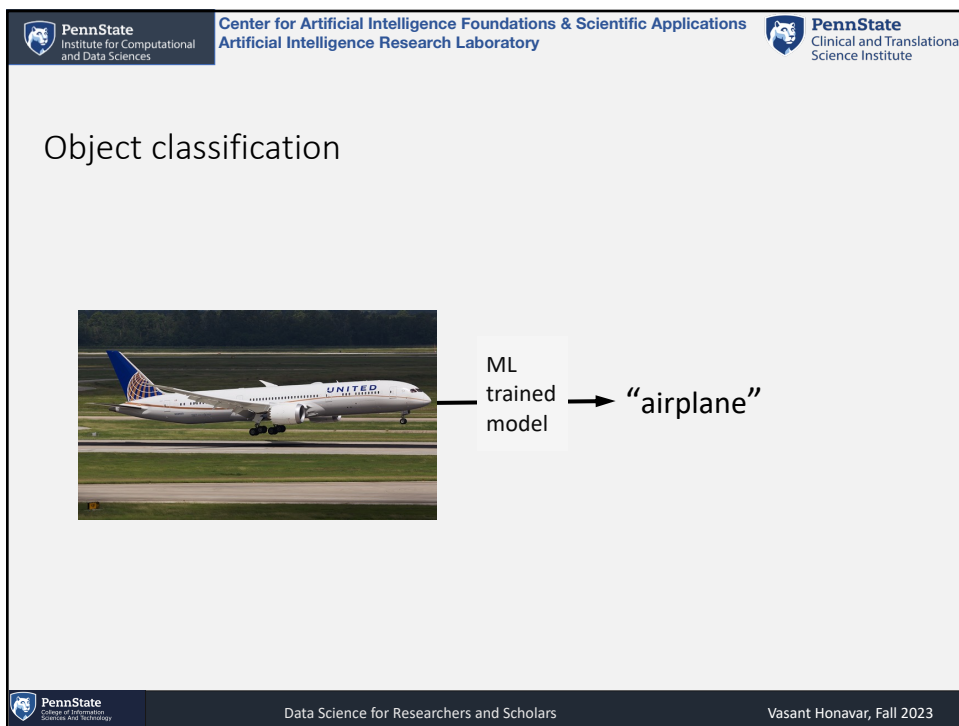
Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

6




7



8

**Why is object classification hard?**

**What you see**



**What the machine sees**


194	210	201	212	199	213	215	195	178	158	182	209
180	189	190	221	209	205	191	167	147	115	129	163
114	126	140	188	176	165	152	140	170	106	78	88
87	103	115	154	143	142	149	153	173	101	57	57
102	112	106	131	122	138	152	147	128	84	58	66
94	95	79	104	105	124	129	113	107	87	69	67
68	71	69	98	89	92	98	95	89	88	76	67
41	56	68	99	63	45	60	82	58	76	75	65
20	43	69	75	56	41	51	73	55	70	63	44
50	50	57	69	75	75	73	74	53	68	59	37
72	59	53	66	84	92	84	74	57	72	63	42
67	61	58	65	75	78	76	73	59	75	69	50

Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

9

**Pixel-based representation**


**Input**




**Learning algorithm**

**Raw image**


+




-



-



+



Encode each image by a vector of pixel values

Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

10

**Pixel-based representation**

Input

Learning algorithm

Raw image

- + Motorbikes
- "Non"-Motorbikes

pixel 1

pixel 2

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

11

**Feature based representation**

Input

Feature representation

Learning algorithm

E.g., Does it have Handlebars? Wheels?

Raw image

Features

- + Motorbikes
- "Non"-Motorbikes

pixel 1

pixel 2


Wheels

Handlebars

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute


Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

12



**PennState**  
Institute for Computational  
and Data Sciences


**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Traditional Machine learning

- Machine learning 30 years ago relied on feature engineering
- Feature engineering is hard!
- It would be nice if we can avoid ad hoc feature engineering
- Kernel machines offer one solution
- Deep learning offers another




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

13



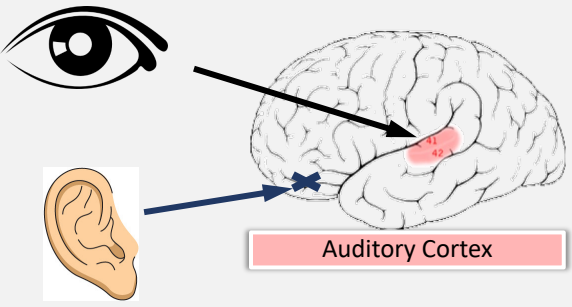
**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## The brain: inspiration for deep learning



Auditory cortex learns to see!

Sur, M., Garraghty, P.E. and Roe, A.W., 1988. Experimentally induced visual projections into auditory thalamus and cortex. *Science*, 242(4884), pp.1437-41.



**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

14


PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

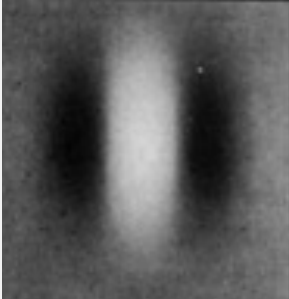
PennState Clinical and Translational Science Institute

## First stage of visual processing: V1

- V1 is the first stage of visual processing in the brain
- Neurons in V1 act as edge detectors



Neuron #1 of visual cortex (model)



Neuron #2 of visual cortex (model)

PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

15

PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

## Basic idea of deep learning

- Involves representation learning or unsupervised feature learning (with subtle distinctions)
- Learn features from data even without knowing the task to be performed?
- Then, stack such representation learning layers to obtain 'deep' networks


PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023


16





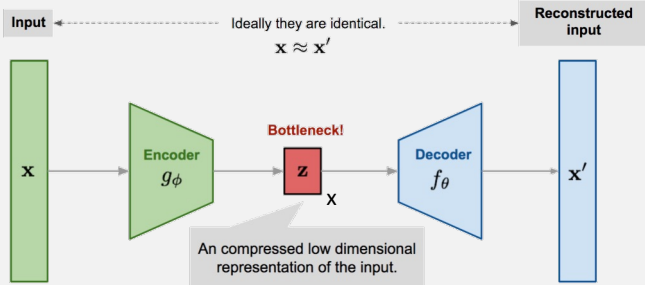
PennState  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



PennState  
Clinical and Translational  
Science Institute

## Autoencoder




Ideally they are identical.  
 $x \approx x'$

**Bottleneck!**  
 $z$

An compressed low dimensional representation of the input.

A feedforward neural network that learns an information preserving representation of its input

- The input-to-hidden part corresponds to an **encoder**
- The hidden-to-output part corresponds to a **decoder**
- Input and output are of the same dimension




PennState  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

17



PennState  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

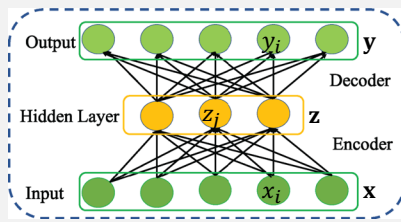


PennState  
Clinical and Translational  
Science Institute

## Autoencoders

- Training a 3-layer linear autoencoder with  $N$  inputs (plus the constant input  $x_{0p} = 1 \forall p$ ) and 1 hidden layer of size  $M$  (plus  $z_{0p} = 1 \forall p$ )


$$z_{jp} = \sum_i w_{ji} x_{ip}$$



$$y_{ip} = \sum_j u_{ij} z_{jp}$$

Reconstruction loss  $E = \sum_p E_p = \frac{1}{2} \sum_p \sum_i (x_{ip} - y_{ip})^2$

Use backpropagation algorithm to learn  $\mathbf{w}_j$  and  $\mathbf{u}_i$



PennState  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

18

PennState  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
 Artificial Intelligence Research Laboratory

PennState  
Clinical and Translational  
Science Institute

## Training Autoencoders

$$z_{jp} = \sum_i w_{ji} x_{ip}$$

$$y_{ip} = \sum_j u_{ij} z_{jp}$$

$$E = \sum_p E_p = \frac{1}{2} \sum_p \sum_i (x_{ip} - y_{ip})^2$$

$$u_{ij} \leftarrow u_{ij} - \frac{\partial E}{\partial u_{ij}}$$

$$\frac{\partial E}{\partial u_{ij}} = \sum_p \frac{\partial E}{\partial y_{ip}} \frac{\partial y_{ip}}{\partial u_{ij}} = - \sum_p (x_{ip} - y_{ip}) z_{jp}$$

$E_p$   
 $\downarrow$   
 $y_{ip}$   
 $\downarrow$   
 $u_{ij}$

PennState  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

19

PennState  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
 Artificial Intelligence Research Laboratory

PennState  
Clinical and Translational  
Science Institute

## Training Autoencoders

$$z_{jp} = \sum_i w_{ji} x_{ip}$$

$$y_{ip} = \sum_j u_{ij} z_{jp}$$

$$E = \sum_p E_p = \frac{1}{2} \sum_p \sum_i (x_{ip} - y_{ip})^2$$

$$w_{ji} \leftarrow w_{ji} - \frac{\partial E}{\partial w_{ji}}$$

$$\frac{\partial E}{\partial w_{ij}} = \sum_p \sum_i \frac{\partial E}{\partial y_{ip}} \frac{\partial y_{ip}}{\partial z_{jp}} \frac{\partial z_{jp}}{\partial w_{ij}} = - \sum_p \sum_i (x_{ip} - y_{ip}) u_{ij} x_{ip}$$


$E_p$   
 $\swarrow \quad \downarrow \quad \searrow$   
 $y_{1p} \quad \dots \quad z_{kp} \quad \dots \quad y_{Np}$   
 $\swarrow \quad \searrow$   
 $z_{jp}$   
 $\downarrow$   
 $w_{ji}$

PennState  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

20



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Autoencoders

- What is the point of learning a mapping that reproduces the input?
  - If the hidden layer has lower dimension than the input, the network is forced to learn a low-dimensional information preserving representation of the input
  - Once such a representation is learned, we can discard the decoder and use the encoder to extract features from the input data that can then be used to train a classification or regression model




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

22



**PennState**  
Institute for Computational  
and Data Sciences

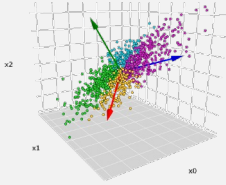
Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Autoencoders

- The autoencoder considered above is a linear autoencoder because it uses a linear function in the hidden layer
- A linear autoencoder learns an encoding that mimics **principal component analysis**
  - the number of hidden nodes correspond to the number of principal components
- We can obtain a non-linear autoencoder by replacing the linear function with a sigmoid function in the hidden layer
- We can use multiple hidden layers to learn complex autoencoder mapping






**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

23



**PennState**  
Institute for Computational  
and Data Sciences

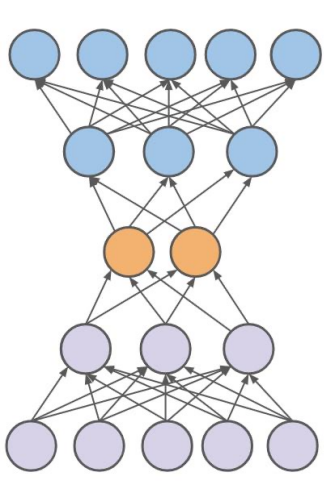
Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Stacked autoencoders

- We can learn hierarchical representations of input data using multi-layer nonlinear encoder
- But as we increase the number of layers, training becomes slow
- Solution: Learn a multi-layer encoder one layer at a time
  - How?






**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

24



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Stacked encoders

- Learn a multi-layer encoder one layer at a time
  - First learn an autoencoder  $\mathbf{x} \rightarrow \mathbf{z} \rightarrow \mathbf{x}$
  - Strip off the decoder  $\mathbf{z} \rightarrow \mathbf{x}$  and keep  $\mathbf{x} \rightarrow \mathbf{z}$
  - Now learn an autoencoder  $\mathbf{z} \rightarrow \mathbf{u} \rightarrow \mathbf{z}$
  - Strip off the decoder  $\mathbf{u} \rightarrow \mathbf{z}$  and keep  $\mathbf{z} \rightarrow \mathbf{u}$
  - Stack the encoders  $\mathbf{x} \rightarrow \mathbf{z}$  and  $\mathbf{z} \rightarrow \mathbf{u}$  to obtain the stacked encoder  $\mathbf{x} \rightarrow \mathbf{z} \rightarrow \mathbf{u}$
  - Repeat the preceding steps if more layers are desired



**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

25

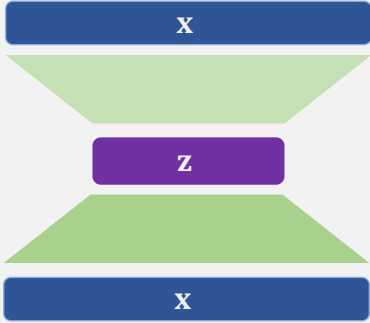
PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

## Learning stacked encoders

First learn an autoencoder  $\mathbf{x} \rightarrow \mathbf{z} \rightarrow \mathbf{x}$



PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

26

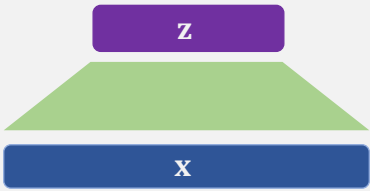
PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

## Learning stacked encoders

Strip off the decoder  $\mathbf{z} \rightarrow \mathbf{x}$  and keep the encoder  $\mathbf{x} \rightarrow \mathbf{z}$



PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

27

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

## Learning stacked encoders

Learn an autoencoder  $\mathbf{z} \rightarrow \mathbf{u} \rightarrow \mathbf{z}$

PennState College of Information Science and Technology | Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

28


PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

## Learning stacked encoders

Strip off the decoder  $\mathbf{u} \rightarrow \mathbf{z}$  and keep the encoder  $\mathbf{z} \rightarrow \mathbf{u}$


PennState College of Information Science and Technology | Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

29



**PennState**  
Institute for Computational  
and Data Sciences

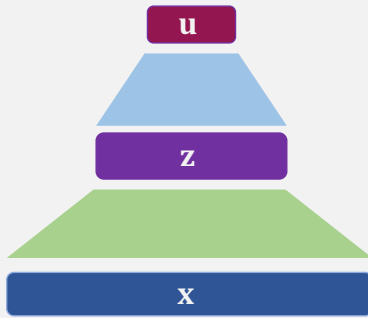
**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Learning stacked encoders

Stack the encoders  $\mathbf{x} \rightarrow \mathbf{z}$  and  $\mathbf{z} \rightarrow \mathbf{u}$  to obtain the stacked encoder  $\mathbf{x} \rightarrow \mathbf{z} \rightarrow \mathbf{u}$






**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

30



**PennState**  
Institute for Computational  
and Data Sciences

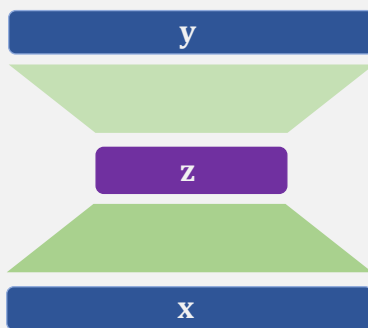
**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Multi-modal encoder-decoder architectures

- Given a data set of images annotated with tags
- Learn a network that
  - Given an image generates tags that describe image content
  - Given a set of tags generates corresponding images






**PennState**  
College of Information  
Science and Technology


Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023 31

31


 PennState  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
 Artificial Intelligence Research Laboratory

 PennState  
Clinical and Translational  
Science Institute

## Autoencoders


- Autoencoders are only able to compress or extract useful features from data that are similar to training data
- The decoded output will be a noisy reconstruction of input
- Can be trained in an unsupervised fashion on large data sets
- The resulting encoding can be used to extract useful features from for supervised training of classifiers from much smaller data sets
- Examples of successful applications of autoencoders
  - Foundation models for computer vision

 PennState  
College of Information  
Science and Technology


Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

32


 PennState  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
 Artificial Intelligence Research Laboratory


 PennState  
Clinical and Translational  
Science Institute

## Denosing autoencoders

- The basic autoencoder minimizes the loss between  $\mathbf{x}$  and the reconstruction  $g(f(\mathbf{x}))$  where  $f$  is the encoder and  $g$  the decoder.
- Denoising autoencoders minimizes the loss between  $\mathbf{x}$  and  $g(f(\mathbf{x} + \mathbf{w}))$ , where  $\mathbf{w}$  is gaussian random noise
- Input and output of a denoising autoencoder



- Noise added to the input forces a denoising autoencoder to learn a mapping from noise perturbed training data to noise-free training data


 PennState  
College of Information  
Science and Technology

[//blog.keras.io/building-a-denosing-autoencoder.html](https://blog.keras.io/building-a-denosing-autoencoder.html) Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023


33





**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Sparse autoencoders

- Add a term to the loss function that forces sparse encoding of training data, e.g., sum of absolute values of the activations of the nodes in the hidden layer
- As a result, different subsets of the hidden nodes are activated by different inputs
- This is often combined by penalties for large weights (e.g., the square of the norm of the weight vectors).




**PennState**  
College of Information Science and Technology

[www.jeremyjordan.me](http://www.jeremyjordan.me) Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

34



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Contractive autoencoders

- Contractive autoencoders make the *feature extraction function* (ie. encoder) robust in the presence of small perturbations of the input
- How?
- Instead of minimizing the loss, minimize the loss plus a term proportional to the magnitude of  $\nabla_{\mathbf{x}}E$  i.e., the gradient of the loss with respect to the input  $\mathbf{x}$




**PennState**  
College of Information Science and Technology

[/f6266h17\\_files/words/](#) Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

35



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Stochastic Encoders and Decoders

- Modern autoencoders use stochastic mappings
- We can generalize the notion of the encoding and decoding functions to encoding and decoding distributions
- The resulting encoders are called variational autoencoders




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

36



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Deep Convolutional networks

- Images have features at different spatial scales
- Need to be able to recognize objects regardless of where they appear in an image
- How can we modify neural networks to accommodate these considerations?
- Extract local features from images



“beak” detector





**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023


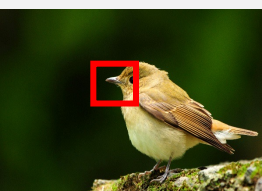
37

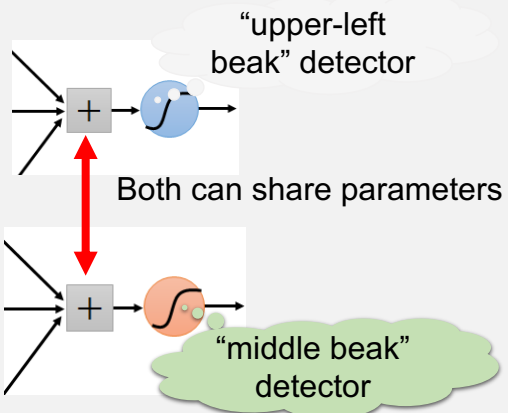
 PennState  
Institute for Computational and Data Sciences
 

 PennState  
Clinical and Translational Science Institute


Center for Artificial Intelligence Foundations & Scientific Applications  
 Artificial Intelligence Research Laboratory


## Translation invariance




Both can share parameters


 PennState  
College of Information Science and Technology
 

 PennState  
Clinical and Translational Science Institute

Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

38

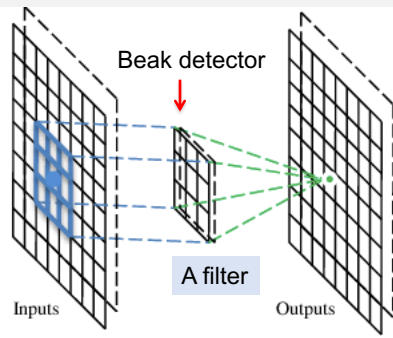
 PennState  
Institute for Computational and Data Sciences
 


 PennState  
Clinical and Translational Science Institute

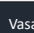
Center for Artificial Intelligence Foundations & Scientific Applications  
 Artificial Intelligence Research Laboratory

## Solution: convolutional neural network

- A CNN is a neural network with one or more convolutional layers
- A convolutional layer performs convolutional operations




 PennState  
College of Information Science and Technology
 

 PennState  
Clinical and Translational Science Institute


Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

39




**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Convolutional Neural networks




Training set: Labeled images of faces.

**Early work:**

- Uhr and students (recognition cones)
- Fukushima (neocognitron)
- Tanimoto
- Levisldi
- Rosenfeld

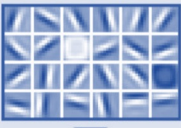
Scaled up and popularized by LeCun

Deep-learning neural networks use layers of increasingly complex rules to categorize complicated shapes such as faces.




Layer 1: The computer identifies pixels of light and dark.

↓



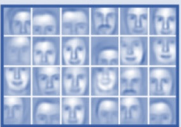
Layer 2: The computer learns to identify edges and simple shapes.

↓




Layer 3: The computer learns to identify more complex shapes and objects.

↓



Layer 4: The computer learns which shapes and objects can be used to define a human face.




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

40



**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Convolution

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

These are the network parameters to be learned.

1	-1	-1
-1	1	-1
-1	-1	1

Filter 1


↓

-1	1	-1
-1	1	-1
-1	1	-1

Filter 2

⋮

Each filter detects a small pattern (3 x 3).




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

41



**PennState**  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Convolution

If stride = 1

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

Dot  
product →


3

-1


1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

6 x 6 image




Data Science for Researchers and Scholars




Vasant Honavar, Fall 2023

42



**PennState**  
Institute for Computational  
and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Convolution

If stride = 2

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0


3

-3


1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

6 x 6 image





Data Science for Researchers and Scholars



Vasant Honavar, Fall 2023

43

 PennState  
Institute for Computational and Data Sciences
 

 PennState  
Clinical and Translational Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## Convolution

Stride = 1

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image


1	-1	-1
-1	1	-1
-1	-1	1


Filter 1

3	-1	-3	-1
-3	1	0	-3
-3	-3	0	1
3	-2	-2	-1

Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

44

 PennState  
Institute for Computational and Data Sciences
 

 PennState  
Clinical and Translational Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## Convolution

Stride = 1

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

-1	1	-1
-1	1	-1
-1	1	-1

Filter 2

Repeat this for each filter

-1	-1	-1	-1
-1	-1	-1	-1
-1	-1	-2	1
-1	0	-4	3

Two 4 x 4 images  
Forming 2 x 4 x 4 tensor

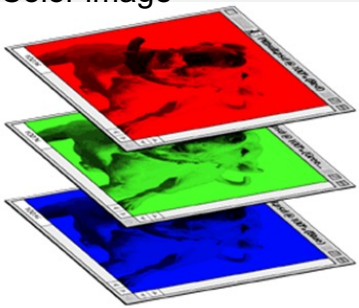
Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

45

PennState  
Institute for Computational  
and Data Sciences
PennState  
Clinical and Translational  
Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## Color image: RGB 3 channels



Color image

1	-1	-1
-1	1	-1
-1	-1	1

Filter 1

-1	1	-1
-1	1	-1
-1	1	-1

Filter 2

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

46

PennState  
Institute for Computational  
and Data Sciences
PennState  
Clinical and Translational  
Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## Convolution v.s. Fully Connected

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

image

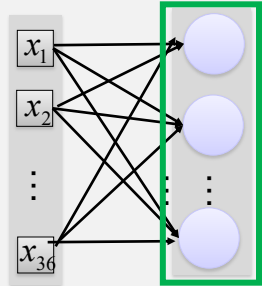
1	-1	-1
-1	1	-1
-1	-1	1

convolution

-1	-1	-1	-1
-1	-1	-2	1
-1	-1	-2	1
-1	0	-4	3

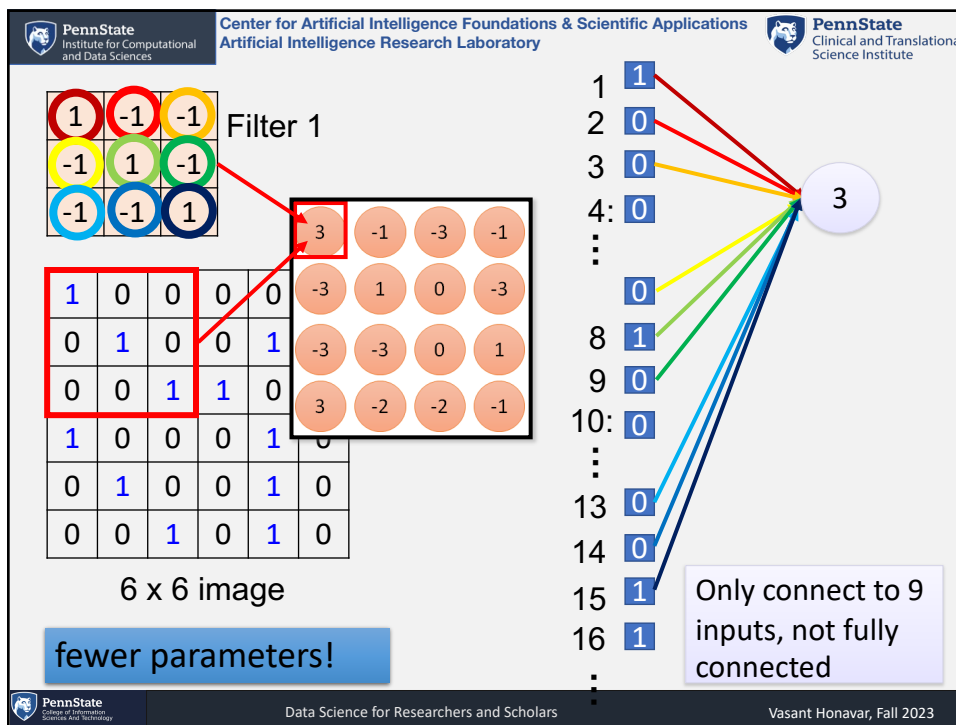
Fully-connected

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

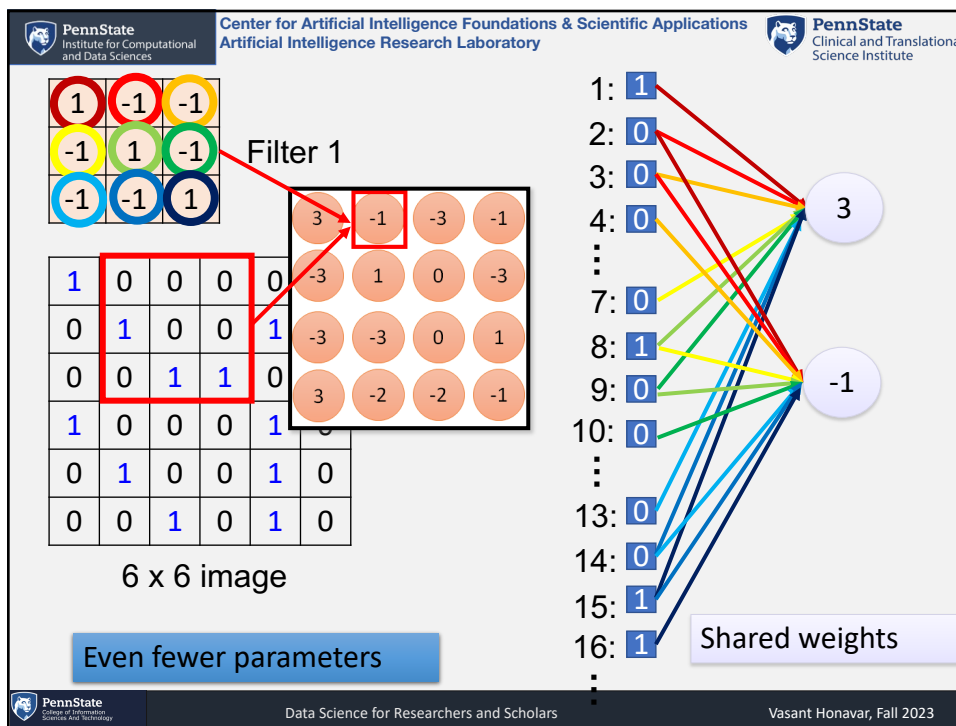


Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

47





48



49

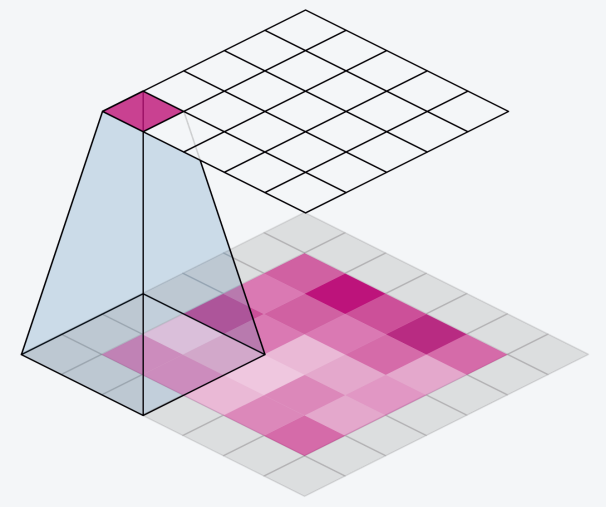




**PennState**  
Institute for Computational and Data Sciences


**PennState**  
Clinical and Translational Science Institute

**Center for Artificial Intelligence Foundations & Scientific Applications**  
**Artificial Intelligence Research Laboratory**

## Convolution at work






**PennState**  
College of Information Science and Technology

**Data Science for Researchers and Scholars**

Vasant Honavar, Fall 2023

50


**PennState**  
Institute for Computational and Data Sciences


**PennState**  
Clinical and Translational Science Institute

**Center for Artificial Intelligence Foundations & Scientific Applications**  
**Artificial Intelligence Research Laboratory**

## Pooling

Max

3	1	1	3
2	5	0	2
1	4	2	1
4	7	2	4

=

5	3
7	4

Avg


3	1	1	3
2	5	0	2
1	4	2	1
4	7	2	4

=

2.75	1.5
4	2.25

- Down-samples the feature maps
- Reduces resolution
- Summarizes the features in a region
- Reduces the number of parameters to be learned
- Combines simpler features into more complex ones

Pooling can be traced back to early work on computer vision, e.g., recognition cones (Uhr, Li, Honavar), pyramids (Rosenfeld, Dyer, Tanimoto, Tsotsos)


**PennState**  
College of Information Science and Technology

**Data Science for Researchers and Scholars**

Vasant Honavar, Fall 2023

51

**PennState**  
Institute for Computational and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory

**PennState**  
Clinical and Translational Science Institute

## Max Pooling

1	0	0	0	0	1
0	1	0	0	1	0
0	0	1	1	0	0
1	0	0	0	1	0
0	1	0	0	1	0
0	0	1	0	1	0

6 x 6 image

Conv

↓

Max Pooling

New image but smaller

-1

1

0

3

2 x 2 image

Each filter is a channel

**PennState**  
College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

52

**PennState**  
Institute for Computational and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory

**PennState**  
Clinical and Translational Science Institute

## The structure of CNN

↓

Convolution

↓

Max Pooling

↓

Convolution

↓

Max Pooling

↓

A new image

Smaller than the original image  
The number of channels is the number of filters

Can repeat many times

-1

1

0

3

**PennState**  
College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

53

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

## Flattening

The diagram illustrates the process of flattening a 2D input into a 1D vector. On the left, a 2x2 grid of 4x4 pixel values is shown. The top row contains values -1 and 1, and the bottom row contains 0 and 3. A large black arrow labeled "Flattened" points to a vertical column of 16 circles containing the values: 3, 0, 1, 3, -1, 1, 0, 3. A second black arrow points from this column to a diagram of a "Fully Connected Feedforward network" with multiple layers of nodes.

Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

54

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

## Complete convolutional neural network (CNN)

The diagram shows the complete pipeline of a CNN. It starts with an "Input Image" of a cartoon character. A "Kernel" is applied through "Convolution" to produce "Feature Maps". These are then processed by "Pooling" to create "Pooled Feature Maps". This sequence of convolution and pooling is repeated. The final "Pooled Feature Maps" are passed through a "Flatten layer" and then a "Fully connected layer" to produce the "Output". The output values are 0.2 for Donald, 0.1 for Goofy, and 0.7 for Tweety.

- Can be repeated many times
- Modern deep neural nets can have tens or hundreds of feature maps as well as layers

Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

55

PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

## A CNN compresses a fully connected network

- Reduces the number of connections
- Shared weights on the edges
- Max pooling further reduces the complexity

PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

56

PennState Institute for Computational and Data Sciences

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

PennState Clinical and Translational Science Institute

## CNN Applications

- Image classification
- Speech recognition
- Biomolecular sequence classification
- Text classification

PennState College of Information Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023

57

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

### Speech recognition using CNN

Convolution happens along the frequency axis

Frequency

Image Spectrogram Time

CNN

PennState College of Information Science and Technology | Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

58

PennState Institute for Computational and Data Sciences | Center for Artificial Intelligence Foundations & Scientific Applications Artificial Intelligence Research Laboratory | PennState Clinical and Translational Science Institute

### Text classification using CNN

sentence matrix  $S \in \mathbb{R}^{d \times |s|}$

convolutional feature map  $C \in \mathbb{R}^{n \times |s| - m + 1}$

pooled representation  $c_{\text{pool}} \in \mathbb{R}^{1 \times n}$

softmax

embedding dimension

$F \in \mathbb{R}^{d \times m}$

I love my new iphone :)

PennState College of Information Science and Technology | Data Science for Researchers and Scholars | Vasant Honavar, Fall 2023

59

PennState  
Institute for Computational  
and Data Sciences
PennState  
Clinical and Translational  
Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## RNA secondary structure prediction using deep CNN

(Sequence)

CGUGUCAGGUCCGGAAGGA  
AGCAGCACUAAC

Representation

(Input tensor)

**Input tensor legend**

- Invalid pairing
- Potential unpaired
- Potential GC pairing
- Potential CG pairing
- Potential UG pairing
- Potential GU pairing
- Potential UA pairing
- Potential AU pairing

(Sequence + its structure)

(Target matrix)

PennState  
College of Information  
Science and Technology
PennState  
Clinical and Translational  
Science Institute

Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

60

PennState  
Institute for Computational  
and Data Sciences
PennState  
Clinical and Translational  
Science Institute

Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory

## Recurrent neural networks

- Good for learning over sequences of data, e.g., a sentence of words
- LSTM (Long Short Term Memory) a popular architecture

Input: a Word

Output: Most likely next word

Memory of previous words influence next prediction


Output so far: Machine

Image source: Adam Gettgey

PennState  
College of Information  
Science and Technology
PennState  
Clinical and Translational  
Science Institute


Data Science for Researchers and Scholars Vasant Honavar, Fall 2023

61



**PennState**  
Institute for Computational  
and Data Sciences

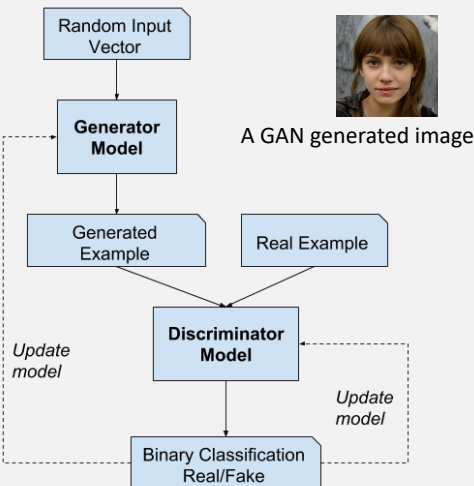
**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute


## Generative adversarial networks (GAN)

- System of two neural networks generator and discriminator compete against each other in a zero-sum game
  - Generator attempts to fool the discriminator
  - Discriminator attempts to defeat the generator
- Provides a kind of unsupervised learning that improves the network
- The result is a generative model whose output is indistinguishable from real data



A GAN generated image

Goodfellow, et al. (2014). Generative Adversarial Networks. NIPS 2014. pp. 2672–2680.  
Special case of an earlier model proposed by Jurgen Schmidhuber in 1991




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023 62

62



**PennState**  
Institute for Computational  
and Data Sciences

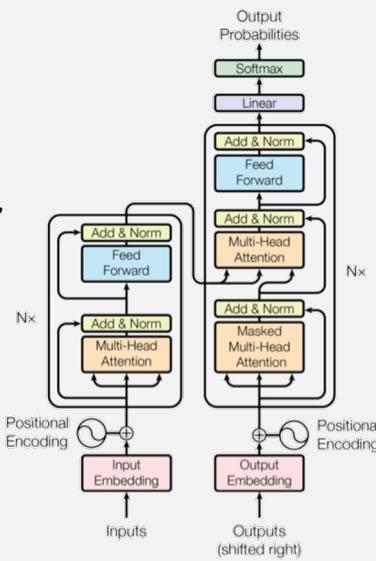
**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Transformers

- Initially introduced for natural language processing (NLP)
- "Transform" input text into output text
  - E.g., translation, text summarization, question answering
- Uses encoder-decoder architecture
- Since extended for text to image, sequence to structure, and many other applications
- Popular pretrained models available, e.g. BERT (encoder only) and GPT (decoder only)



Vaswani, A. et al. (2017). Attention is all you need. In NIPS 2017.




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023 63

63



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Transformers compared to older models

- The **RNN** and **LSTM** neural models were designed to process language and perform tasks like classification, summarization, translation, and sentiment detection
  - RNN: Recurrent Neural Network
  - LSTM: Long Short Term Memory
- In both models, layers get the next input word and have access to some previous words, allowing it to use the word's left context
- They used word embeddings where each word was encoded as a vector of 100-300 real numbers representing its meaning




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

64



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Transformers compared to older models

- Transformers extend this to allow the network to process a word input given the context: words both to the left and right
- Transformers add additional features, like attention which identifies the important words in context
- And break the problem into two parts:
  - An encoder (e.g., BERT)
  - A decoder (e.g., GPT)




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023


65





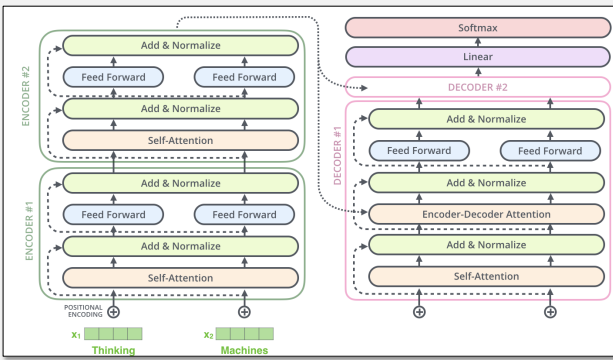
**PennState**  
Institute for Computational  
and Data Sciences

**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Transformer model



Encoder (e.g., BERT)
Decoder (e.g., GPT)




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

66



**PennState**  
Institute for Computational  
and Data Sciences

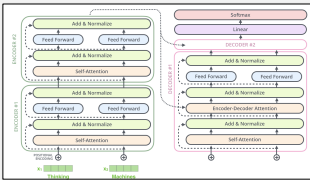
**Center for Artificial Intelligence Foundations & Scientific Applications**  
Artificial Intelligence Research Laboratory




**PennState**  
Clinical and Translational  
Science Institute

## Transformers, GPT, and BERT

- A transformer uses an **encoder stack** to model input, and uses **decoder stack** to model output (using input information from encoder side)
- If we do not have input, we just want to model the “next word”, we can get rid of the encoder side of a transformer and output “next word” one by one. This gives us **GPT**
- If we are only interested in training a language model for the input for some other tasks, then we do not need the decoder of the transformer, that gives us **BERT**






**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

67



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Training a Transformer

- Transformers typically use semi-supervised learning with
  - Unsupervised pretraining over a very large dataset of general text
  - Followed by supervised **fine-tuning** over a focused data set of inputs and outputs for a particular task
- Tasks for pretraining and fine-tuning commonly include:
  - language modeling
  - next-sentence prediction
  - question answering
  - sentiment analysis
  - paraphrasing




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023

68



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Hype versus reality of GPT

- Transformer models such as GPT are fundamentally generative models trained on large data sets
  - Very sophisticated next word, phrase, sentence predictors
- Many ethical concerns – copyright, plagiarism, intellectual property
- Great as muses or sources of inspiration
- Great for generating increasingly realistic art, poetry, etc.
- Not necessarily great for applications where accuracy matters
  - Known to hallucinate – make up "facts", citations, etc.
- Many potential downsides
  - Misinformation generation at scale
  - Phishing and cyberattacks at scale




**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars


Vasant Honavar, Fall 2023 69

69



**PennState**  
Institute for Computational  
and Data Sciences


Center for Artificial Intelligence Foundations & Scientific Applications  
Artificial Intelligence Research Laboratory



**PennState**  
Clinical and Translational  
Science Institute

## Summary

- Unsupervised representation learning offers a powerful mechanism for learning useful features from complex data
- Multi-layer deep networks, given sufficient data, can be trained to perform many tasks that were once beyond the reach of machine learning
- Powerful tools Keras, Pytorch allow rapid prototyping of deep learning solutions
- But deep learning, as it is practiced today, is an environmental nightmare – need better approaches
- Deep learning systems can be easily fooled by adversarial data samples - need learning methods that are robust to adversarial attacks
- Deep learning produces black boxes that are hard to understand and explain –need better tools to explain the results of deep learning
- Advanced deep learning models e.g., LLMs are susceptible to hallucination, catastrophic forgetting, and other problems



**PennState**  
College of Information  
Science and Technology

Data Science for Researchers and Scholars

Vasant Honavar, Fall 2023