


An abstract graphic on the left side of the slide, featuring a sphere-like shape composed of a dense grid of intersecting red, green, and blue lines. The lines are curved and follow the contour of the sphere, creating a complex, woven pattern. The sphere is set against a dark gray background.


Demystifying AI

Feature Learning

Instructor: Pat Virtue

Listen Learner

Chris Harrison, CMU



<https://chrisharrison.net/index.php/Research/ListenLearner>

Games to warm up

1. https://www.sporcle.com/games/MrChewypoo/minimalist_disney
2. <https://www.sporcle.com/games/Stanford0008/minimalist-cartoons-slideshow>
3. <https://www.sporcle.com/games/MrChewypoo/minimalist>

Dimensionality Reduction

(raw) input features



1,000,000



low-dimensional
features



30



Jasmine

Dimensionality Reduction



30



1,000,000

Dimensionality Reduction





Dimensionality Reduction

<https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html>



Exercise: Human-defined Feature Space

Step1: Write a bunch of digits 0-9 on post-it notes



Exercise: Human-defined Feature Space


Step2: In groups, students try to organize digits on a 2-D coordinate plot



Optional: label the extreme ends of both coordinate axes

Exercise: Human-defined Feature Space

Step2: In groups, try to organize digits on a 2-D coordinate plot




Optional: label the extreme ends of both coordinate axes

Exercise: Human-defined Feature Space

Step 3: Prediction!


1. Select three students: A,B,C
2. Student A draws a new digit and hands it to student B
3. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
4. Student C looks at the coordinate and the plot (but not the drawing from A) and predicts the digit, 0-9



Exercise: Human-defined Feature Space

Step 4: Creation!


1. Select three students: A,B,C
2. Student A draws a new digit and hands it to student B
3. Student B thinks about where to plot it and comes up with a 2-D coordinate, (x, y)
4. Student C looks at the coordinate and the plot (but not the drawing from A) and **draws a new digit**




Learning to Organize Data

Neural networks can learn to organization too!

Image $\rightarrow (z_1, z_2) \rightarrow$ Image



z_2




z_1

<https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html>

Reminder: Image Classification

Demo of (Fully-connected) neural network to classify images of hand-written digits


<https://www.cs.ryerson.ca/~aharley/vis/fc/>



Reminder: Image Classification

Demo of (Fully-connected) neural network to classify images of hand-written digits

<https://www.cs.ryerson.ca/~aharley/vis/fc/>




```
def relu(z)
    return
    max(z, 0)
```

Poll 1


When we pass a single image through this network how many times does $\text{ReLU}(z)$ get called (assume $\text{ReLU}(z)$ can only take a single number not many)?

- A. 1
- B. 10
- C. 100
- D. 300
- E. 400**
- F. 30,000



Digit Autoencoder

<https://cs.stanford.edu/people/karpathy/convnetjs/demo/autoencoder.html>




Poll 2

We talked about using an autoencoder to learn a 2-dimensional feature vector $[feature_1, feature_2]$.

Where in this network can we find those two numbers?

- A. Layer 3 input
- ☒ B. Layer 3 output
- C. Layer 3 weights
- D. Layer 3 bias (offset)(intercept)
- E. Layer 4 weights




Dimensionality Reduction with Deep Learning

Hinton, Geoffrey E., and Ruslan R. Salakhutdinov.


"Reducing the dimensionality of data with neural networks."

Science 313.5786 (2006): 504-507.



Listen Learner


<https://chrisharrison.net/index.php/Research/ListenLearner>



Autoencoder Demo




Zhuoyue Lyu, Safinah Ali, and
Cynthia Breazeal. EAAI 2022.

<https://colab.research.google.com/github/ZhuoyueLyu/5046225a9ae3675cf633c1df5f63be06/digits-interpolation-notebook-eaai.ipynb>




Autoencoder Demo

Feature space interpolation



Exploring Feature Space

<https://experiments.withgoogle.com/ai/melody-mixer/view/>



Exploring Feature Space

<https://experiments.withgoogle.com/ai/beat-blender/view/>

