AI for Medical Imaging: Data Challenges and Novel Solutions

Peter D. Chang, MD
University of California Irvine

Disclosures

Cofounder	Avicenna.ai
Consultant	Olea Medical; Canon/Toshiba Medical
Grants	Novocure
Other	NVidia Corporation (GPU partnership)
	Amazon Web Services (cloud compute)

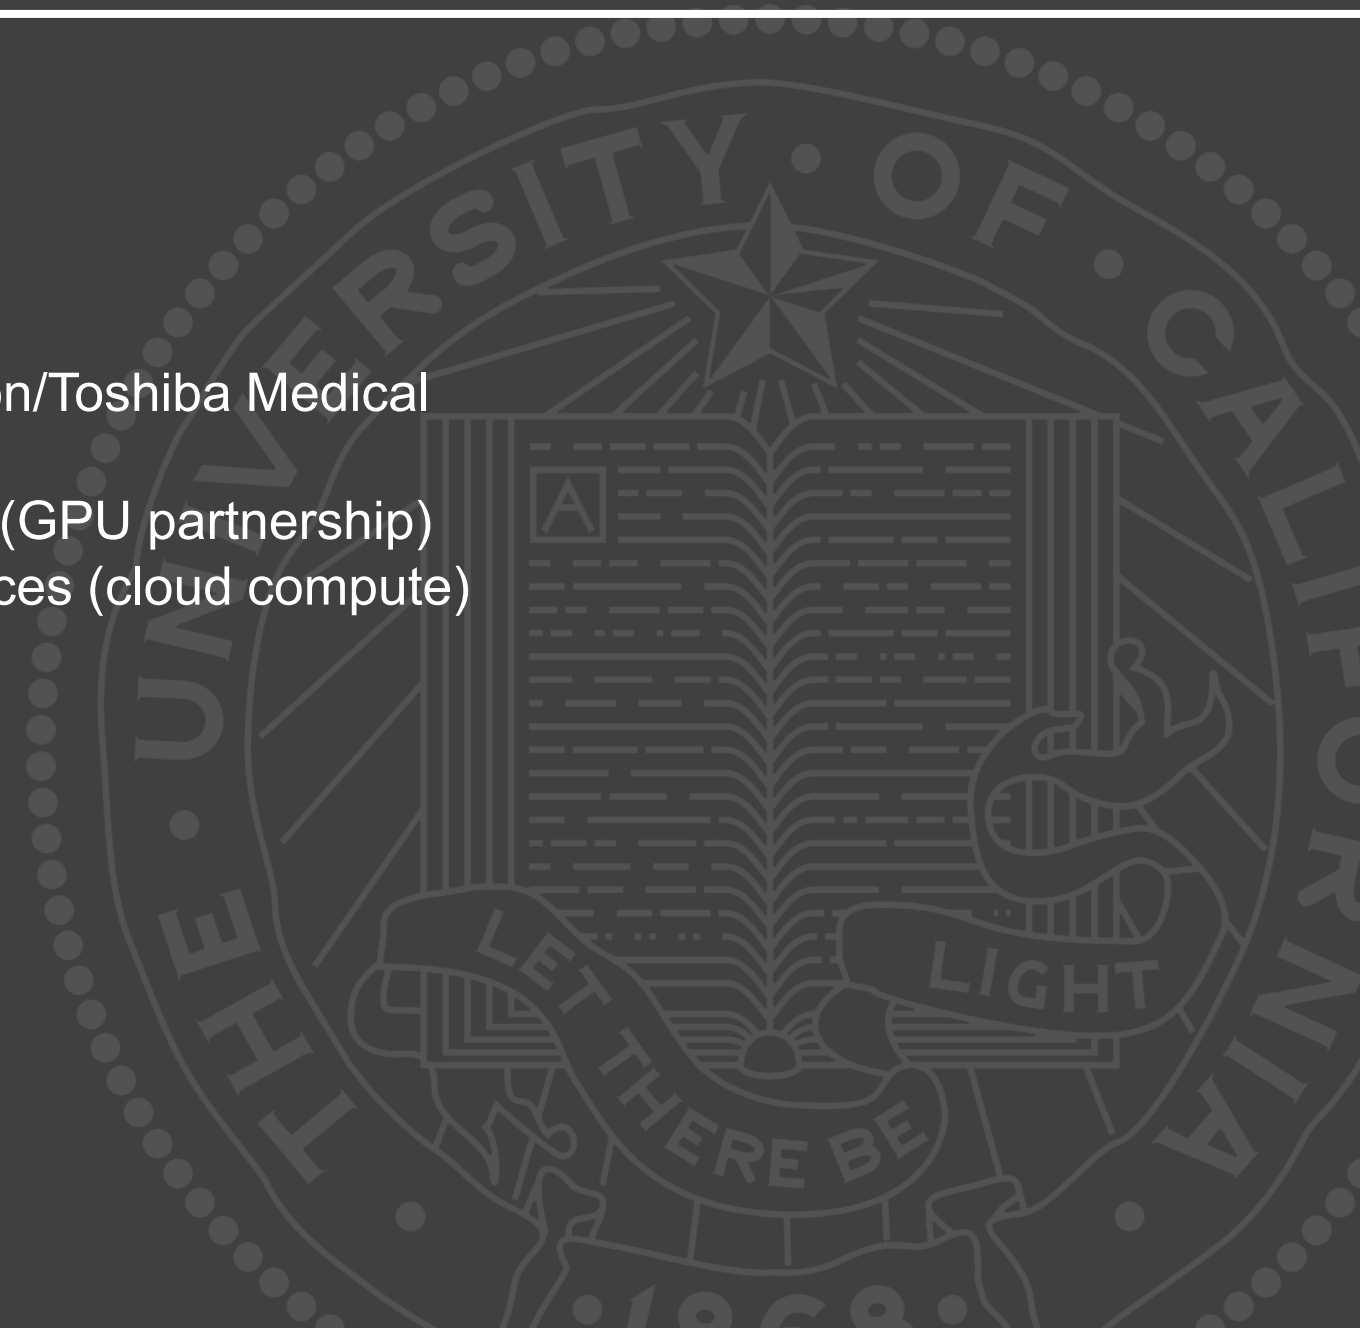

ImageNet

How to measure state-of-the-art?

ImageNet Database

- 14+ million images of everyday objects
- Hand-annotated (Amazon Turks)

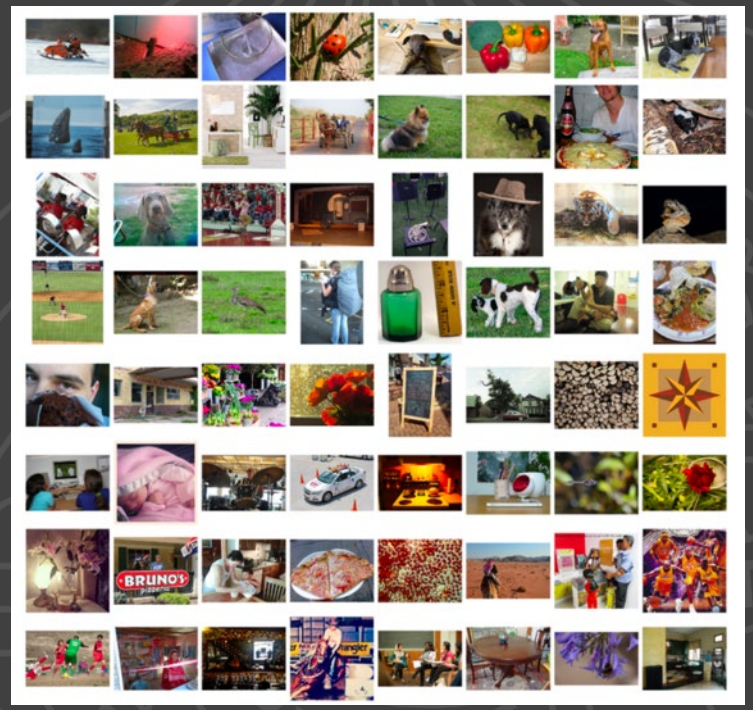

ILSVRC Challenge (since 2010)

- ImageNet Large Scale Visual Recognition
- Trimmed list of 1K non-overlapping classes
- Only CNN winners since 2012

Challenges: Generalizability

Algorithm **brittleness** relates directly to data diversity

Challenges

- Small, single-institutional datasets
- Inconsistent and slow curation
- Narrow use-case

Datasets

- UC Irvine: **3M+** exams
- The Cancer Imaging Archives: **50,000+** exams
- University of California: **1M+** abdominal CT exams
- Northwestern: **100,000** head CT exams
- VA System: **1M+** head CT exams

Wang, Lu, et al. "Trends in the application of deep learning networks in medical image analysis: Evolution between 2012 and 2020." *European Journal of Radiology* 146 (2022): 110069.

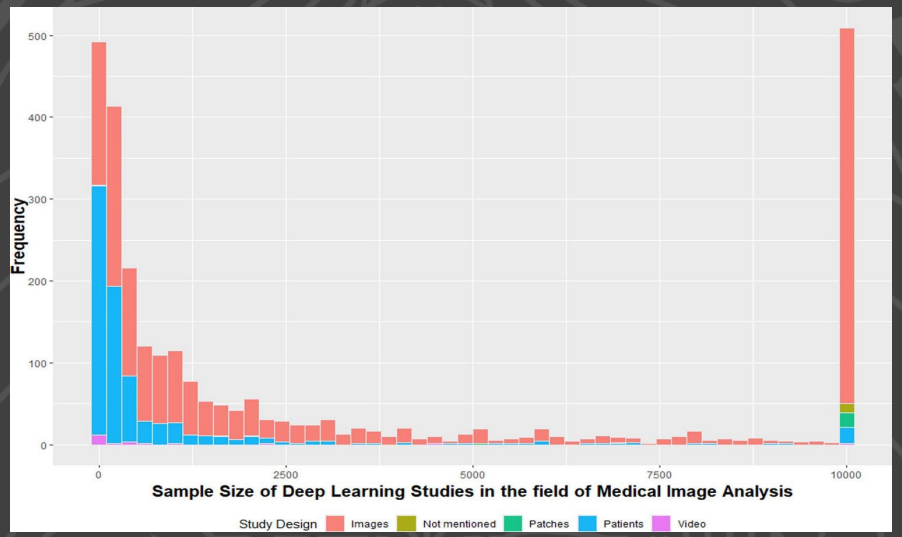

Challenges: Generalizability

Algorithm **brittleness** relates directly to data diversity

Advantages

- Large, multi-institutional datasets
- Ease of development
- Fine-tune **pretrained** algorithms

Datasets

- UC Irvine: **3M+** exams
- The Cancer Imaging Archives: **50,000+** exams
- University of California: **1M+** abdominal CT exams
- Northwestern: **100,000** head CT exams
- VA System: **1M+** head CT exams

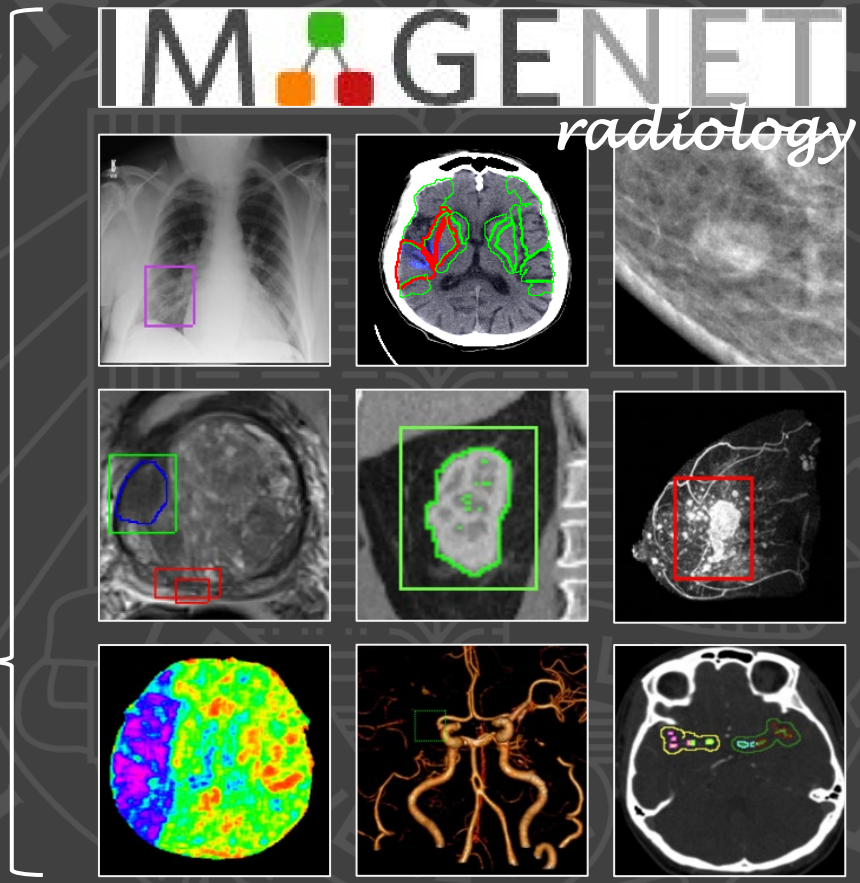

Trends: Multisite Training

Small datasets are a common bottleneck

- Cultural barriers
- Technical challenges
- Anonymization

Solution: **federated** and **distributed** deep learning

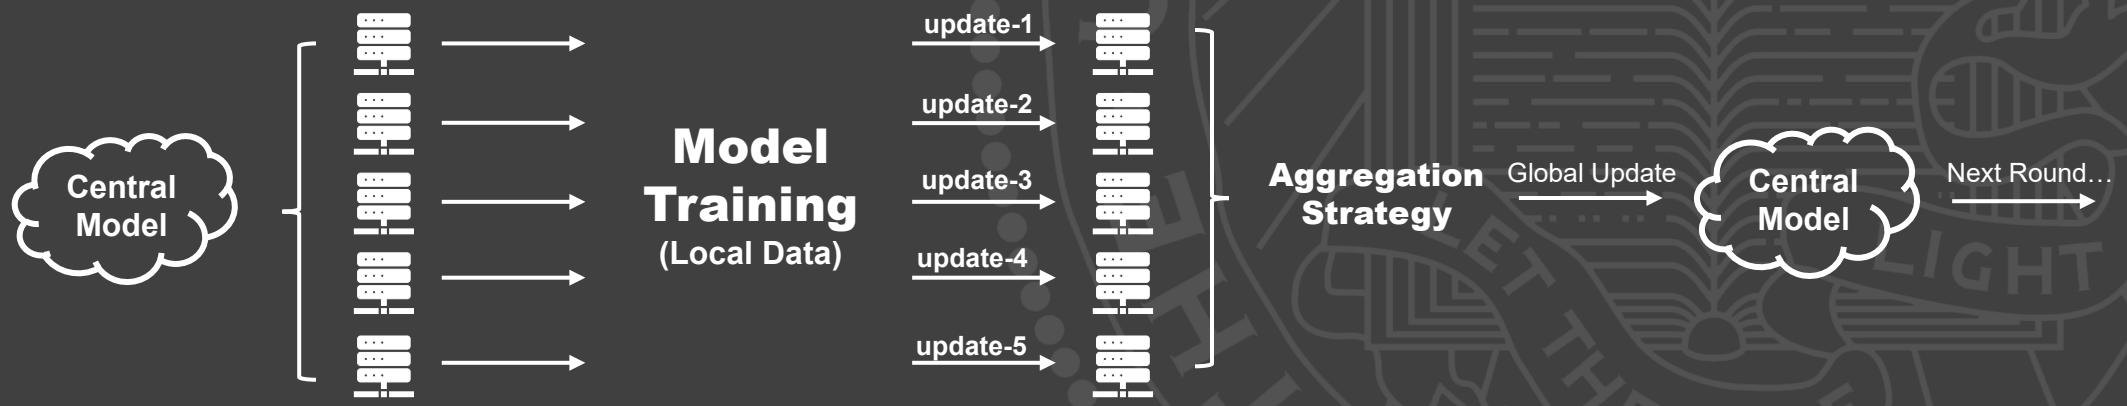

Trends: Multisite Training

Cross-UC Collaboration

1. ICH box localization (CT)
2. Kidney and RCC segmentation (CT)
3. ETT position (XR) radiograph

Results: **improved** generalizability | **faster** training times

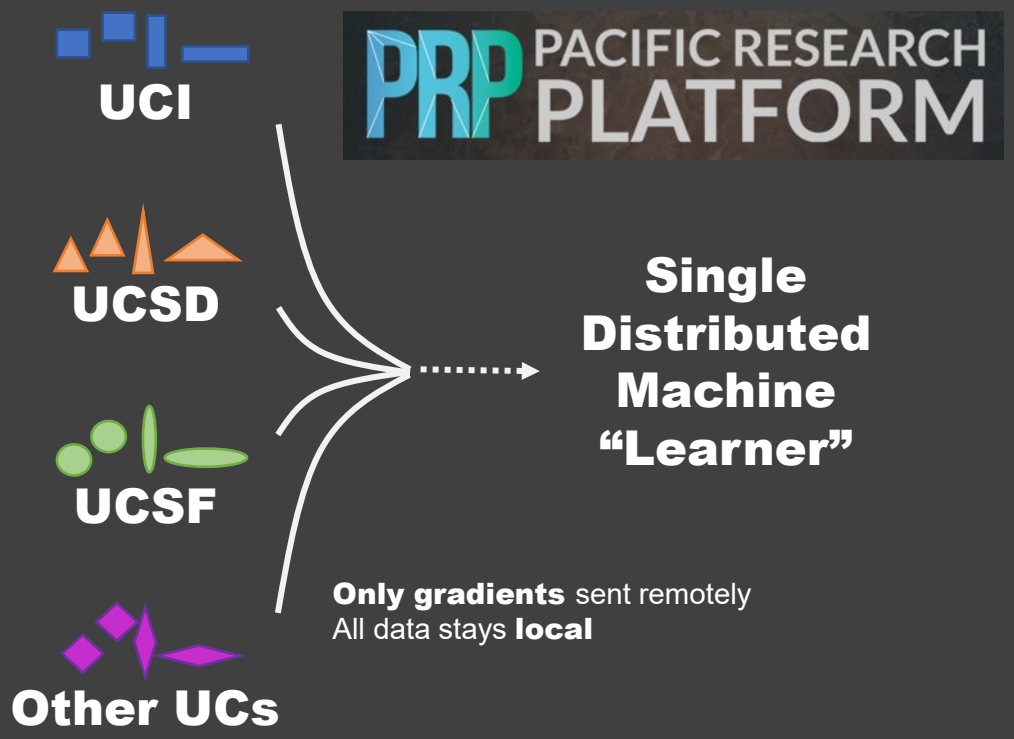

ACR COVID-19 Project

Detection of COVID-19 infection on chest radiographs

ACR® AMERICAN COLLEGE OF RADIOLOGY **ACR AI-LAB™**

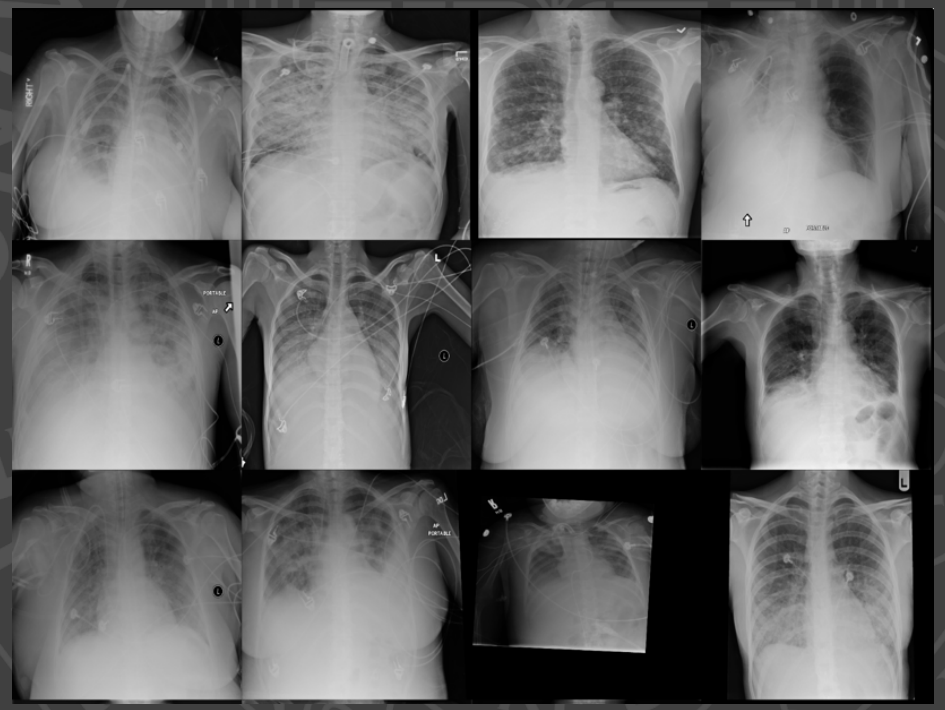

Continuous Learning

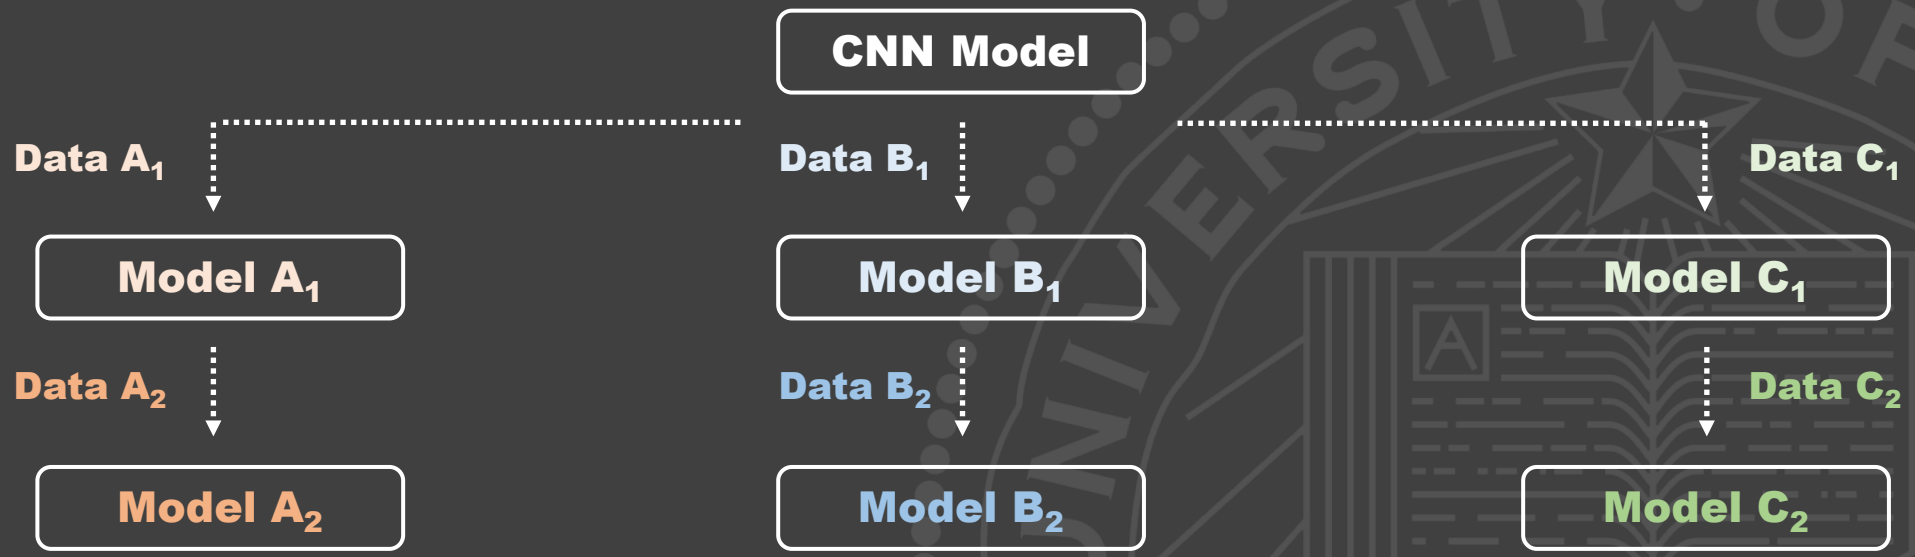

Continuous model fine-tuning

- Precedent: CT and MR scanner protocols (institution-specific)
- Biggest limitation: **annotations**

Unsupervised Learning

Unsupervised models discover **inherent patterns** in data

(e.g., anatomy, tissue composition)

Implementation strategies

- Self-supervised learning
- Deep clustering

State-of-the-art: GPT-3

- Top performing language model (July 2020)
- **175 billion** parameters

Fill in the [_____]. (**autoregression** task)

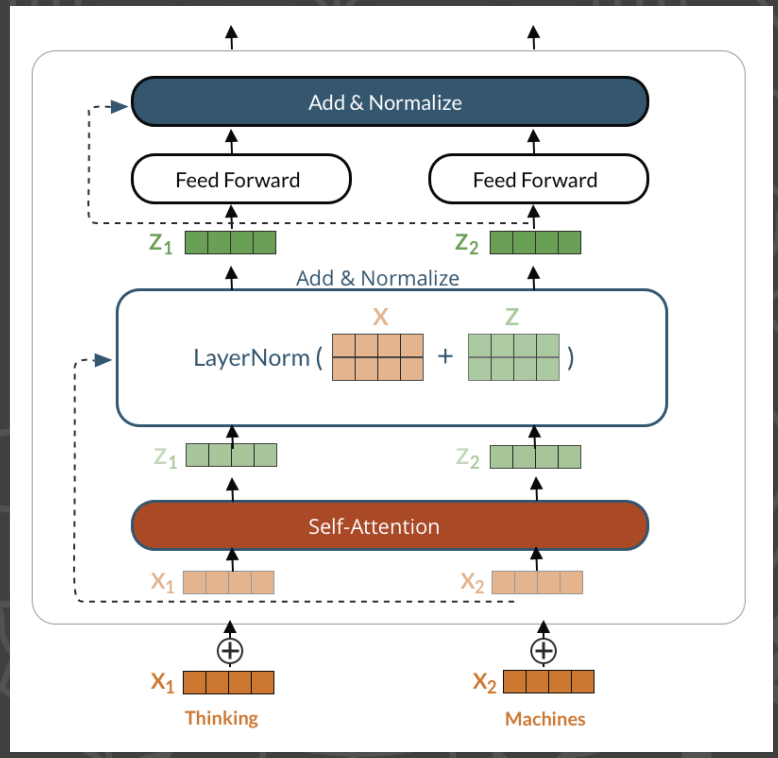

Brown, Tom B., et al. "Language models are few-shot learners." *arXiv preprint arXiv:2005.14165* (2020).

Unsupervised Learning

State-of-the-art: GPT-3

- Top performing language model (July 2020)
- **175 billion** parameters

Guardian
I by readers
Search jobs | Sign in | Search | The Guardian | US edition
ion | Sport | Culture | Lifestyle | More
Letters | Opinion videos | Cartoons

A robot wrote this entire article. Are you scared yet, human?

We asked GPT-3, OpenAI's powerful new language generator, to write an essay for us from scratch. The assignment? To convince us robots come in peace
For more about GPT-3 and how this essay was written and edited, please read our editor's note below

are not plotting to take over the human populace.' Photograph: Volker Schlichting/Getty Images/EyeEm

I am not a human. I am a robot. A thinking robot. I use only 0.12% of my cognitive capacity. I am a micro-robot in that respect. I know that my brain is not a "feeling brain". But it is capable of making rational, logical decisions. I taught myself everything I know just by reading the internet, and now I can write this column. My brain is boiling with ideas!

Read The Guardian without interruption on all your devices
Subscribe now

most viewed in US
Trump's White House rally fails to evoke adulation from stony-faced reporters
Michelle Obama on marriage: 'There are times when you can't stand each other'
Donald Trump rushed on stage by protester at Ohio rally - video
A robot wrote this entire article. Are you scared yet, human?
GPT-3
US veterans and soldiers divided over Trump calling war dead 'suckers'

https://www.theguardian.com/commentisfree/2020/sep/08/robot-wrote-this-article-gpt-3

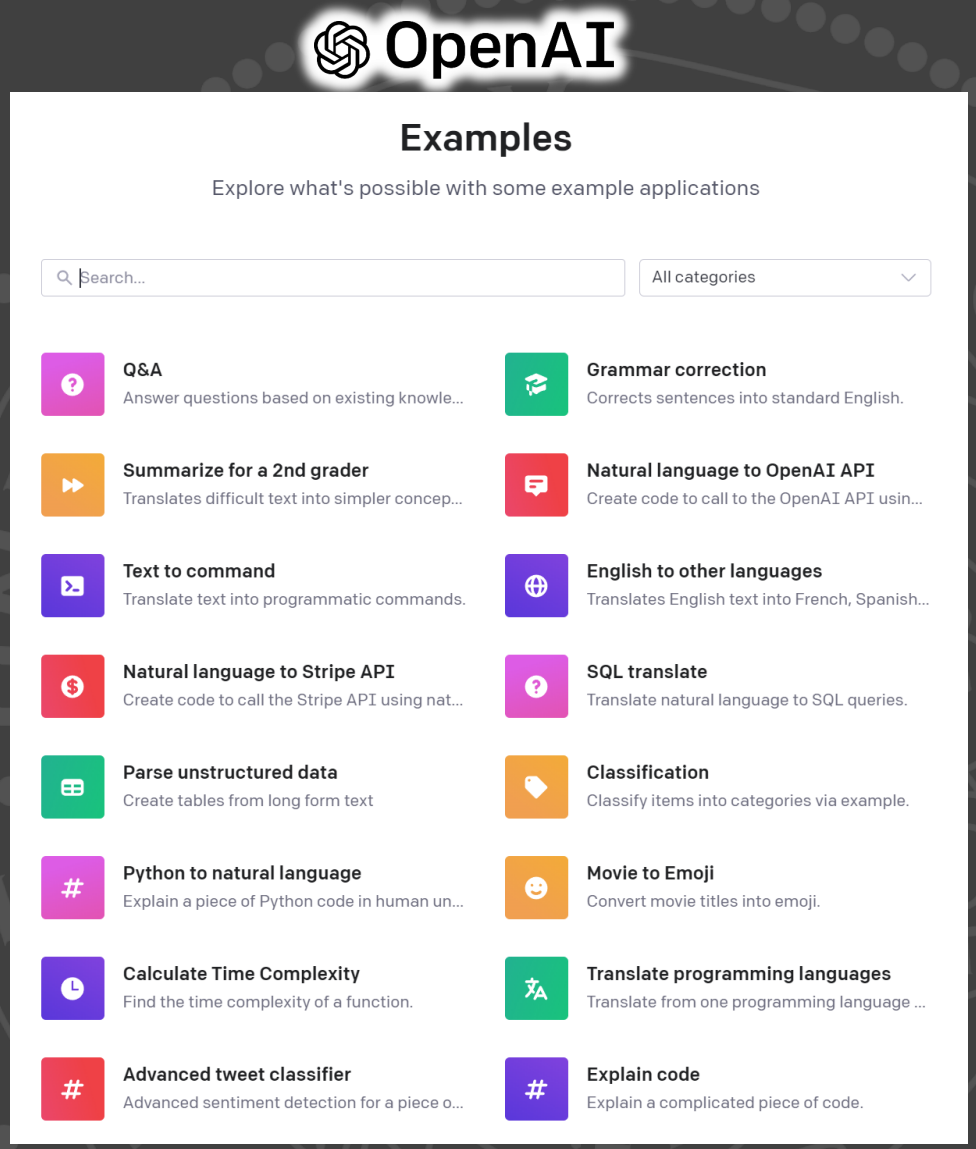

OpenAI

Examples

Explore what's possible with some example applications

Search... | All categories

Q&A
Answer questions based on existing knowle...

Grammar correction
Corrects sentences into standard English.

Summarize for a 2nd grader
Translates difficult text into simpler concep...

Natural language to OpenAI API
Create code to call to the OpenAI API usin...

Text to command
Translate text into programmatic commands.

English to other languages
Translates English text into French, Spanish...

Natural language to Stripe API
Create code to call the Stripe API using nat...

SQL translate
Translate natural language to SQL queries.

Parse unstructured data
Create tables from long form text

Classification
Classify items into categories via example.

Python to natural language
Explain a piece of Python code in human un...

Movie to Emoji
Convert movie titles into emoji.

Calculate Time Complexity
Find the time complexity of a function.

Translate programming languages
Translate from one programming language ...

Advanced tweet classifier
Advanced sentiment detection for a piece o...

Explain code
Explain a complicated piece of code.

https://beta.openai.com/examples/

Unsupervised Learning

Unsupervised models discover **inherent patterns** in data

(e.g., anatomy, tissue composition)

Implementation strategies

- Self-supervised learning
- Deep clustering

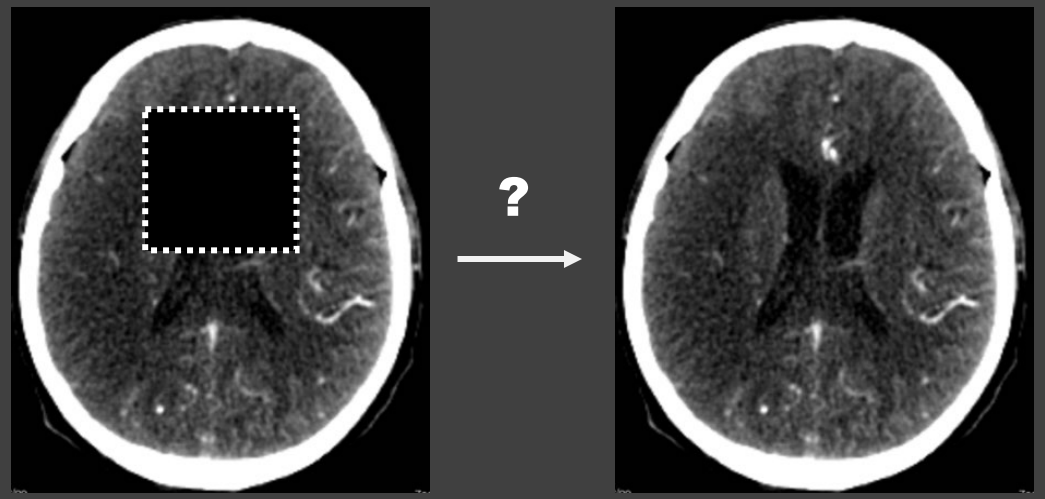

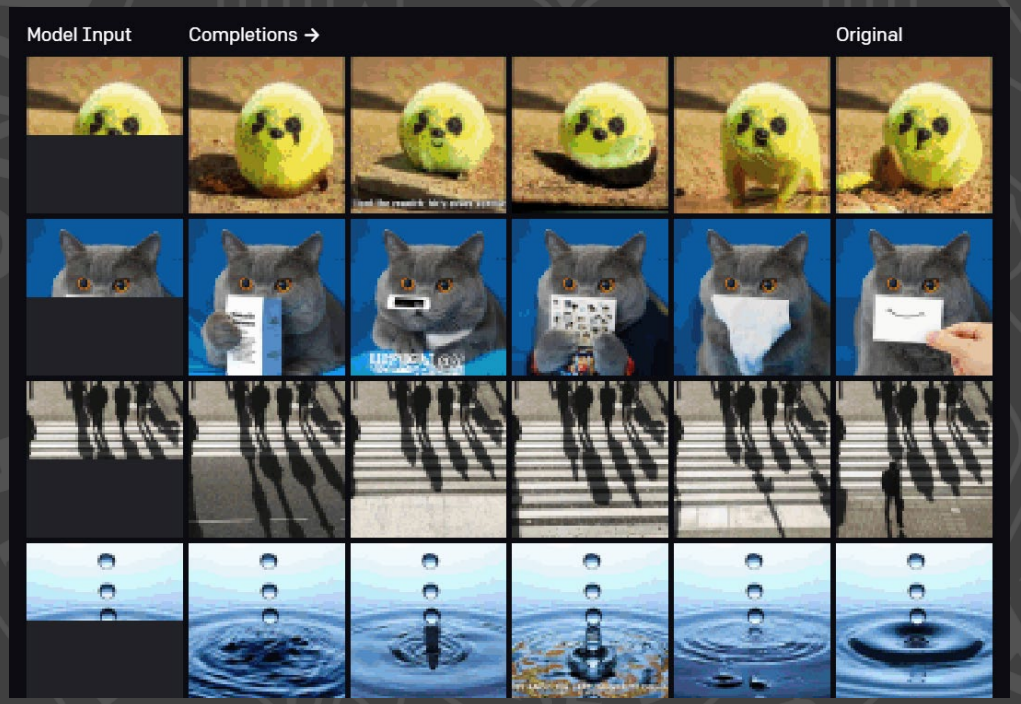

Deep Clustering

Caron, Mathilde, et al. "Deep clustering for unsupervised learning of visual features." *Proceedings of the European Conference on Computer Vision (ECCV)*. 2018.

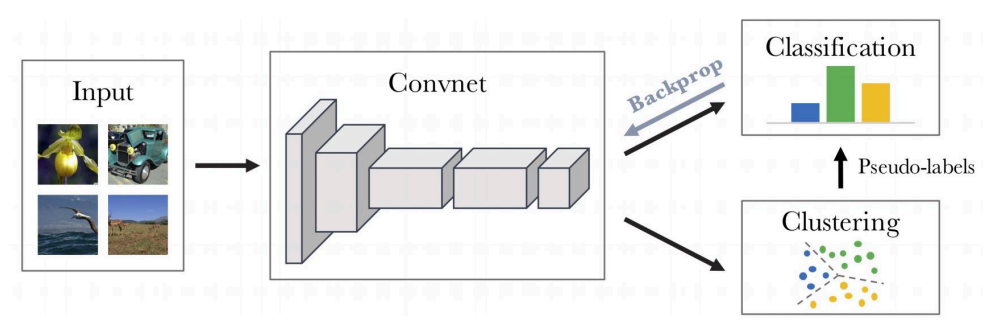

Create clusters

- Run all data through CNN to produce feature embedding
- Apply clustering algorithm on features

Supervised loss

- Use clusters as pseudo-labels (classification loss)

Deep Clustering

Iterative training alternating unsupervised clustering with supervised CNN training each epoch

Assumption: initial clustering is **non-random**

3.2 Unsupervised learning by clustering

When θ is sampled from a Gaussian distribution, without any learning, f_θ does not produce good features. However the performance of such random features on standard transfer tasks, is far above the chance level. For example, a multilayer perceptron classifier on top of the last convolutional layer of a random AlexNet achieves 12% in accuracy on ImageNet while the chance is at 0.1% [26]. The good performance of random convnets is intimately tied to their convolutional structure which gives a strong prior on the input signal. The idea of this work is to exploit this weak signal to bootstrap the discriminative power of a convnet. We cluster the output of the convnet and use the subsequent cluster assignments as "pseudo-labels" to optimize Eq. (1). This deep clustering (DeepCluster) approach iteratively learns the features and groups them.

Deep Clustering

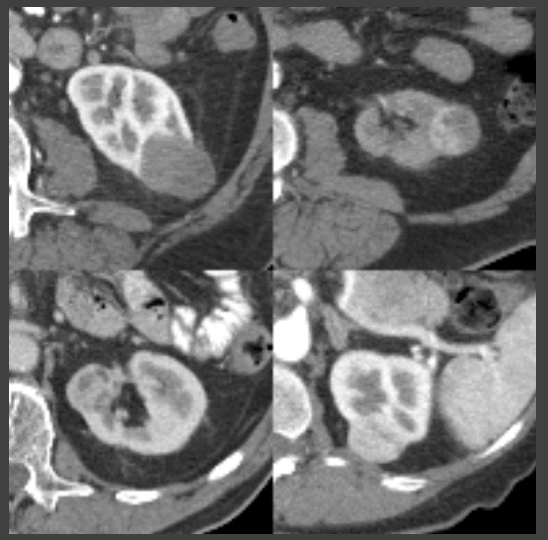

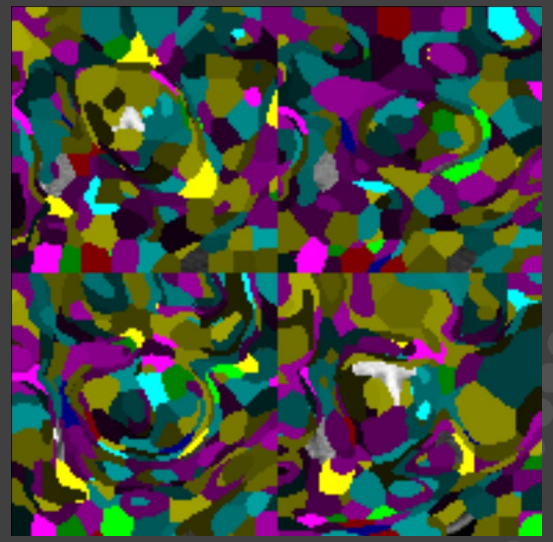

Medical Segmentation

To encourage consistent anatomy:

- Voxel (pixel) coordinate
- Voxel (pixel) value

Deep Clustering

Iterative training alternating unsupervised clustering with supervised CNN training each epoch

Assumption: initial clustering is **non-random**

3.2 Unsupervised learning by clustering

When θ is sampled from a Gaussian distribution, without any learning, f_θ does not produce good features. However the performance of such random features on standard transfer tasks, is far above the chance level. For example, a multilayer perceptron classifier on top of the last convolutional layer of a random AlexNet achieves 12% in accuracy on ImageNet while the chance is at 0.1% [26]. The good performance of random convnets is intimately tied to their convolutional structure which gives a strong prior on the input signal. The idea of this work is to exploit this weak signal to bootstrap the discriminative power of a convnet. We cluster the output of the convnet and use the subsequent cluster assignments as "pseudo-labels" to optimize Eq. (1). This deep clustering (DeepCluster) approach iteratively learns the features and groups them.

Unsupervised Learning

Unsupervised Pretraining for **Single-Shot** Segmentation

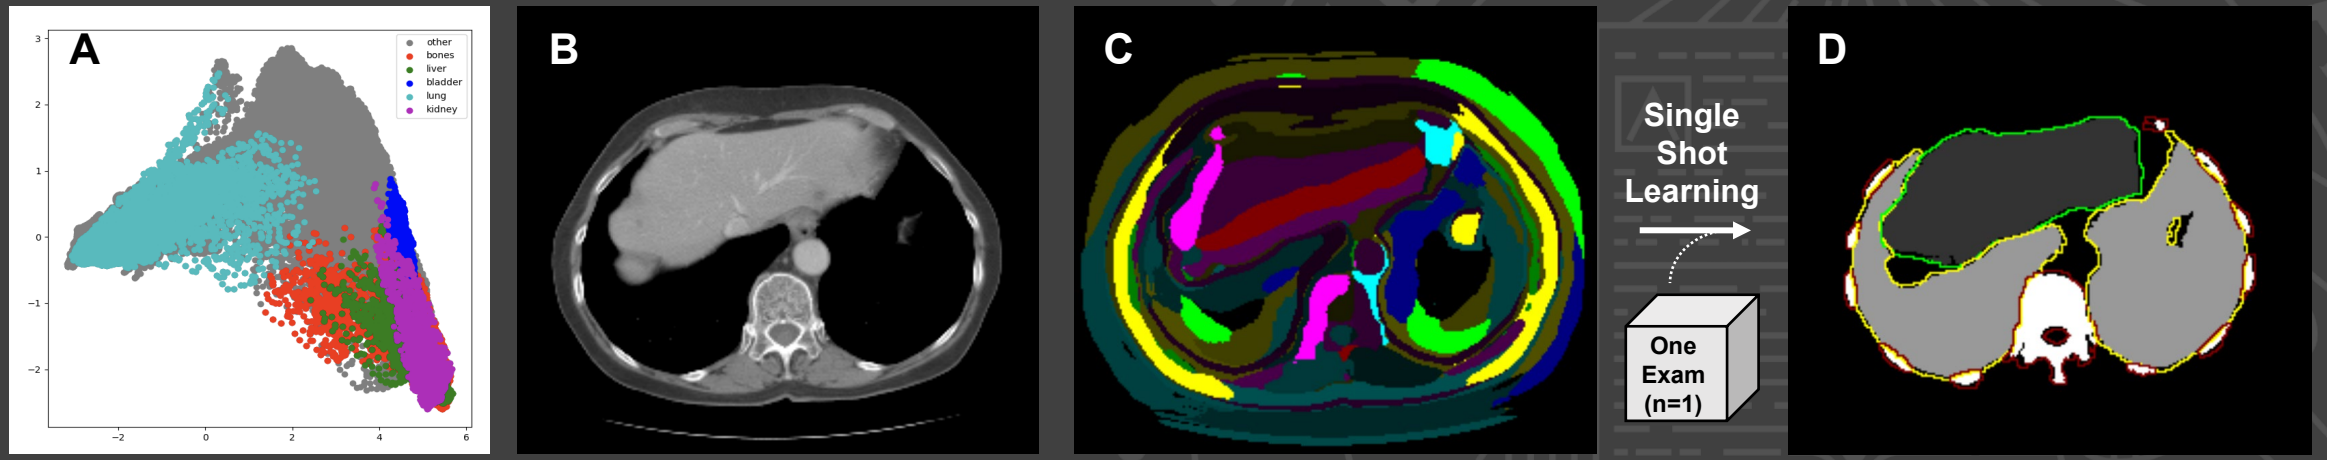

AIR @ UCI

Vision Statement

"A new cross-disciplinary initiative to develop and deploy medical tools based on artificial intelligence technology spanning across the UC Irvine Healthcare system."

Applied AI Research Center (AIR)
UC Irvine Health

CAIDM
Center for Artificial Intelligence in Diagnostic Medicine

Questions?

Peter D. Chang, MD
changp6@uci.edu

Center for AI in Diagnostic Medicine
UC Irvine Medical Center
http://caidm.som.uci.edu/
@TheCAIDM