


CS 188: Artificial Intelligence Conclusion



Instructors: Angela Liu and Yanlai Yang

University of California, Berkeley

(Slides adapted from Pieter Abbeel, Anca Dragan, and Stuart Russell)

Research Frontiers

- Deep Unsupervised Learning
- Robot Learning
- AI for Science

Research Frontiers

- *Deep Unsupervised Learning*
- Robot Learning
- AI for Science

Deep Unsupervised Learning

- Key hypothesis:

Task 1

- IF neural network smart enough to predict:
 - Next frame in video
 - Next word in sentence
 - Generate realistic images
 - “Translate” images from one style to another style
 - ...

Task 2

- THEN same neural network is ready to do Deep Supervised Learning from very small data-set

Next Word Prediction (OpenAI's GPT)

Prompt: A train carriage containing controlled nuclear materials was stolen in Cincinnati today. Its whereabouts are unknown.


GPT-2: *The incident occurred on the downtown train line, which runs from Covington and Ashland stations. In an email to Ohio news outlets, the U.S. Department of Energy said it is working with the Federal Railroad Administration to find the thief.*

'The theft of this nuclear material will have significant negative consequences on public and environmental health, our workforce and the economy of our nation,' said Tom Hicks, the U.S. Energy Secretary, in a statement. 'Our top priority is to secure the theft and ensure it doesn't happen again.'


The stolen material was taken from the University of Cincinnati's Research Triangle Park nuclear research site, according to a news release from Department officials.

The Nuclear Regulatory Commission did not immediately release any information. According to the release, the U.S. Department of Energy's Office of Nuclear Material Safety and Security is leading that team's investigation.

Transfer from Unsupervised Learning



Example Setting




Unsupervised Sentiment Neuron

This is one of Crichton's best books. The characters of Karen Ross, Peter Elliot, Munro, and Amy are beautifully developed and their interactions are exciting, complex, and fast-paced throughout this impressive novel. And about 99.8 percent of that got lost in the film. Seriously, the screenplay AND the directing were horrendous and clearly done by people who could not fathom what was good about the novel. I can't fault the actors because frankly, they never had a chance to make this turkey live up to Crichton's original work. I know good novels, especially those with a science fiction edge, are hard to bring to the screen in a way that lives up to the original. But this may be the absolute worst disparity in quality between novel and screen adaptation ever. The book is really, really good. The movie is just dreadful.

Benchmarks

DATASET	METRIC	OUR RESULT	PREVIOUS RECORD	HUMAN
Winograd Schema Challenge	accuracy (+)	70.70%	63.7%	92%+
LAMBADA	accuracy (+)	63.24%	59.23%	95%+
LAMBADA	perplexity (-)	8.6	99	~1-2
Children's Book Test Common Nouns (validation accuracy)	accuracy (+)	93.30%	85.7%	96%
Children's Book Test Named Entities (validation accuracy)	accuracy (+)	89.05%	82.3%	92%
Penn Tree Bank	perplexity (-)	35.76	46.54	unknown
WikiText-2	perplexity (-)	18.34	39.14	unknown


Unsupervised Learning in Vision



Predict Missing Patch



Unsupervised Learning + linear classifier



AI for Art Creation



An astronaut riding a horse in a photorealistic style


[openai.com/dall-e-2]

Disinformation can be aided by AI generated text

Prompt: Recycling is good for the world, no, you could not be more wrong.

GPT-2: Recycling is NOT good for the world. It is bad for the environment, it is bad for our health, and it is bad for our economy. I'm not kidding. Recycling is not good for the environment. It is destructive to the earth and it is a major contributor to global warming. Recycling is not good for our health. It contributes to obesity and diseases like heart disease and cancer. Recycling is bad for our economy. It increases the cost of a product, and in turn, the price of everything that is made with that product. Recycling is not good for our nation. We pay a tremendous price for the privilege of having the world's most advanced and efficient recycling system. Recycling is a huge, colossal waste of time, energy, money, and resources.

Disinformation can be aided by Deep Fakes



[thispersondoesnotexist.com]

[thiscatdoesnotexist.com]

AI Ethics Ever More Important

- Why?
 - AI is making decisions, at scale
 - Any kind of issues (e.g. bias or malignant use) could significantly affect *people*
- Many open questions:
 - Who is responsible?
 - How to diagnose and prevent?

Research Frontiers

- Deep Unsupervised Learning
- ***Robot Learning***
- AI for Science

PR-1






Reinforcement Learning



Learned Skills



Challenge Task: Robotic Laundry



Research Frontiers

- Deep Unsupervised Learning
- Robot Learning
- *AI for Science*

NEWS · 30 NOVEMBER 2020

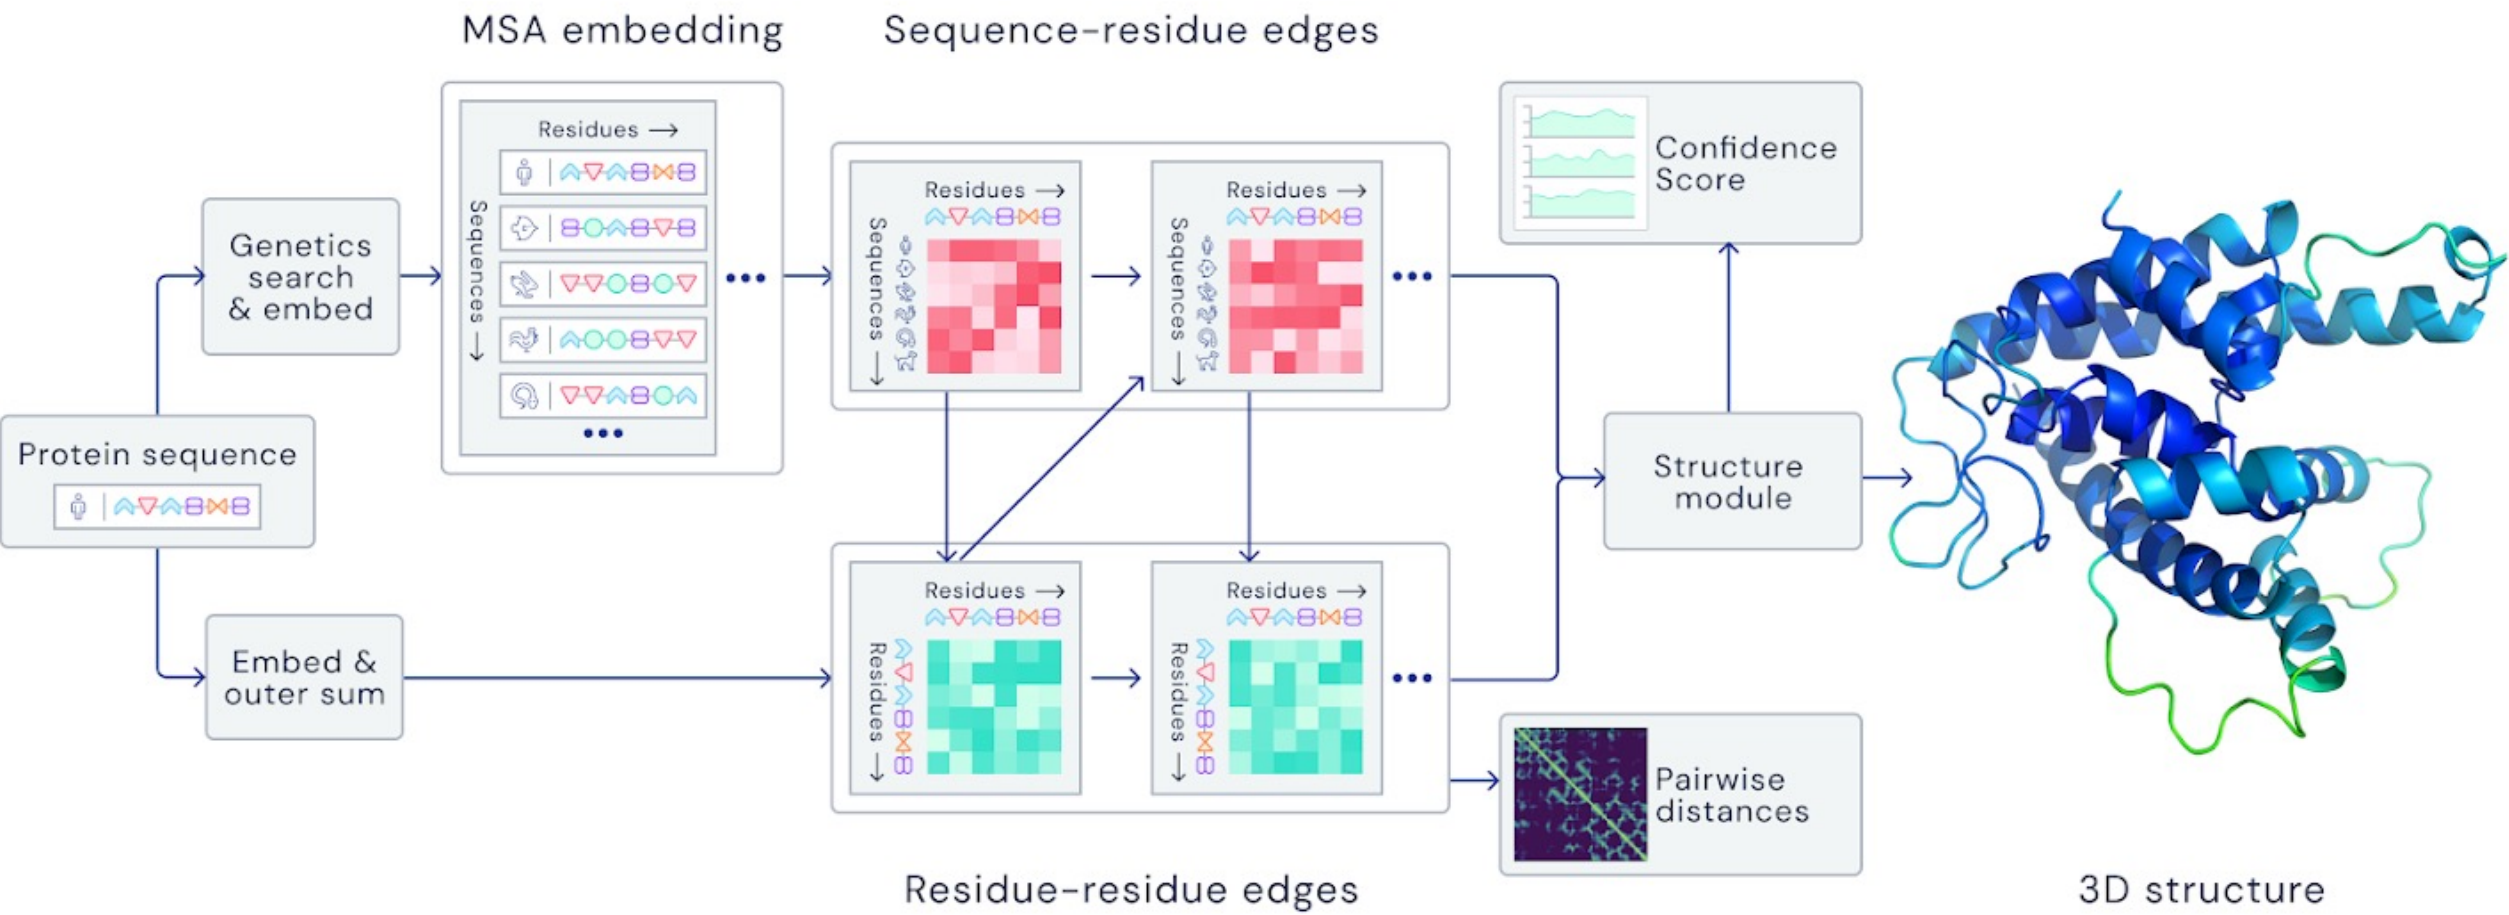
'It will change everything': DeepMind's AI makes gigantic leap in solving protein structures


Google's deep-learning program for determining the 3D shapes of proteins stands to transform biology, say scientists.

[Ewen Callaway](#)




A protein's function is determined by its 3D shape. Credit: DeepMind





T1037 / 6vr4
90.7 GDT
(RNA polymerase domain)



T1049 / 6y4f
93.3 GDT
(adhesin tip)


- Experimental result
- Computational prediction

Symbolic Math: Integrals and ODEs

Equation	Solution
$y' = \frac{16x^3 - 42x^2 + 2x}{(-16x^8 + 112x^7 - 204x^6 + 28x^5 - x^4 + 1)^{1/2}}$	$y = \sin^{-1}(4x^4 - 14x^3 + x^2)$
$3xy \cos(x) - \sqrt{9x^2 \sin(x)^2 + 1}y' + 3y \sin(x) = 0$	$y = c \exp(\sinh^{-1}(3x \sin(x)))$
$4x^4yy'' - 8x^4y'^2 - 8x^3yy' - 3x^3y'' - 8x^2y^2 - 6x^2y' - 3x^2y'' - 9xy' - 3y = 0$	$y = \frac{c_1 + 3x + 3 \log(x)}{x(c_2 + 4x)}$

Table 4: Examples of problems that our model is able to solve, on which Mathematica and Matlab were not able to find a solution. For each equation, our model finds a valid solution with greedy decoding.

Where to Go Next?



Where to go next?

- Congratulations, you've seen the basics of modern AI
 - ... and done some amazing work putting it to use!
- How to continue:
 - CS 189 (Machine Learning, people.eecs.berkeley.edu/~jrs/189/)
 - CS 182 (Neural Networks, youtube.com/playlist?list=PL_iWQOsE6TfVmKkQHucjPAoRtIJYt8a5A)
 - CS 285 (Deep RL, youtube.com/playlist?list=PL_iWQOsE6TfXxKgl1GgyV1B_Xa0DxE5eH)
 - CS 287 (Robotics, people.eecs.berkeley.edu/~pabbeel/cs287-fa19/)
 - Data 100 (Data Science, ds100.org)
 - Data 102 (Inference and Decisions, data102.org)
 - EECS 127 (Optimization, inst.eecs.berkeley.edu/~eecs127/fa20/)
 - CS 294-158 Deep Unsupervised Learning (sites.google.com/view/berkeley-cs294-158-sp20)
 - EECS 126 (Probability), CS 281A (ML Theory), CS 280 (Computer Vision), CS 288 (Natural Language), CS 287H (Human-Robot Interaction)
 - ... and more: coursecapture.berkeley.edu

Lightweight Opportunities to Keep Learning



- Andrew Ng weekly newsletter:

The Batch: <https://www.deeplearning.ai/thebatch/>



- Jack Clark (former Comms Director OpenAI) weekly newsletter:

Import AI: <https://jack-clark.net/>



- Rachel Thomas AI Ethics course:


Course website: ethics.fast.ai



- Pieter Abbeel podcast:

The Robot Brains Podcast: <https://therobotbrains.ai>

Midterm Contest Results



Top 3

- 1st place:

- Panda

Jonathan Yue

- 2nd place

- PacmanCutiesRun!


Enqian Zheng

- 3rd place

- iPacman

Luyang Jiang and Martin Liu

Congratulations!





Arvind Rajaraman

Head TA

Email: Admin cs188@, Personal: arvind.rajaraman@

Hi all! I'm a rising junior majoring in EECS and Business from New Jersey. I'm super excited to impart a curiosity for AI to you all. Academically, I'm interested in computer vision, reinforcement learning, and autonomous vehicles. Outside work, I like hiking, listening to hip hop/pop music, eating food, and am a recently converted Warriors fan!



Neil Thomas

Discussion TA

Email: nthomas@

Hi friends! I'm a 5th year CS PhD Student advised by Professor Yun S. Song. My [research](#) focuses on using machine learning to guide the design of proteins to help facilitate the transition to a bio-based economy. When I'm not teaching, doing research, or procrastinating, I like to play frisbee, ride my bike, cook recipes from Alison Roman, scramble up rocks, and freestyle rap.




Joy Liu

Student/Staff Engagement TA

Email: joy.liu@

I'm a third year majoring in CS and cognitive science, and I like books, music, and the Albany Bulb.



Zhuang Liu

Homework TA

Email: zhuangl@

I am a final year PhD student at Berkeley, advised by Prof. Trevor Darrell. My research interests lie in deep learning and computer vision. I love to enjoy great food, movies and books, like most of you. I'm super excited to on the staff of CS 188 this summer! Hope to learn from you all!



Sid Ijju

Project TA

Email: sidijju@

Hey everyone! I'm a 3rd year EECS and Business major from Colorado. I'm really interested in human-robot interaction and reinforcement learning research. In my free time, I like baking, drawing, and playing video games (mainly Genshin and League). I'm looking forward to meeting all of you and having a great semester!



Cham Yao

DSP & Logistics TA

Email: chamyao@

Hey guys! I'm Cham, a 3rd year CS-major from Hinsdale, Illinois. My hobbies include reading (two of my favorite authors are Hemingway and Murakami), chess (I play the Queen's Gambit and Caro), and running. Really excited to experience the summer semester with y'all!



Andrew Wang

Website & Logistics TA

Email: andrew-wang@

Hi! I am a fourth-year EECS major from the LA area. My current research interests include RL and robotics. I'm very excited to meet everyone this semester. Feel free to reach out if you want to talk about anything!



Perry Dong

Exam TA

Email: perrydong@

Hi everyone! I'm Perry, a rising junior studying CS and math from Beijing, China. I am interested in AI and specifically reinforcement learning. In my free time, I like to read and watch movies. Look forward to a great summer session.



Weijia Zeng

Exam TA

Email: 0z5e1n3g@

Hello, this is Weijia! I am a graduating Applied Math and CS major. I enjoy learning and doing research in reinforcement learning and 3D vision. In my free time, I play and watch basketball (Let's go Warriors!!), and enjoy boba tea. Looking forward for a great summer session!



Xiangwei Kong

Exam TA

Email: xk39@

Hello! I'm a rising junior and major in Comp Sci and Econ major. I come from Beijing. My research is currently focused on Web-based NLP task. I took CS 188 last semester and luckily become a TA for this summer and I am really excited to meet everyone! Outside of work, I love playing tennis and running.

Readers



Pranav Muralikrishnan

Email: pranavm@

Hi! My name is Pranav and I'm a sophomore math and cs major. I'm excited to be on course staff this summer!



Austen Liao

Email: austenliao@

Hi! I'm Austen, a rising 3rd-year CS Major from the Bay Area. I really enjoyed CS 188, and I hope I can help more people find value in it too! In my free time, I like to play Smash, obsess over the Warriors, edit videos, and spend time in nature.



Jerry Sun

Email: qjsun@

Hi my name is Jerry and I am a rising junior studying computer science and environmental science. In my spare time, I like to run, listen to Mitski and support the Warriors.



Stephen Yang

Email: stephenyang@

Hiiii hope life is going well! I am Stephen and I am a rising senior majoring in CS and Cogsci. I am from the beautiful coastal city of Qingdao (or Tsingtao if you know the beer!) In my free time, I like to play VR games, fly drones, play pianos, and cook random food... I absolutely love 188 and am excited to get to know you guys!



Shawn Zhao

Email: shzhao@


Hi, I'm Shawn, a rising 3rd year EECS major. I'm from Maryland, around the D.C. area, so if you are also from there, feel free to reach out! In my free time I play CS:GO/Valorant,osu!, TFT, and Genshin Impact. Looking forward to a great semester!



Stanley Kwon

Email: shkwon@


Hey, I'm Stanley! I'm a rising junior majoring in computer science and economics from Rockville, MD. In my free time, I like scrolling through TikTok to find the end of my #fyp, eating Chipotle, and juggling. I recently got into skateboarding so I can be a kool kid on the block. Looking forward to meeting you all!



Ryan Zhu

Email: ryanzhu@

Heyo! I'm a rising third year CS major from SoCal interested in algorithms, complexity, and linguistics. I love eating food, solving puzzles, and slaying the spire! Looking forward to a great summer! :D




Ketrina Yim
CS188 Artist

That's It!

- Help us out with course evaluations!
- Have a great rest of your summer!
- Good luck with finals!

Stickers Time!!





Kiermaty 2012