Deep Learning: Overview

Ju Sun
Computer Science & Engineering
University of Minnesota, Twin Cities

September 6, 2023

Outline

What is Deep Learning (DL)?

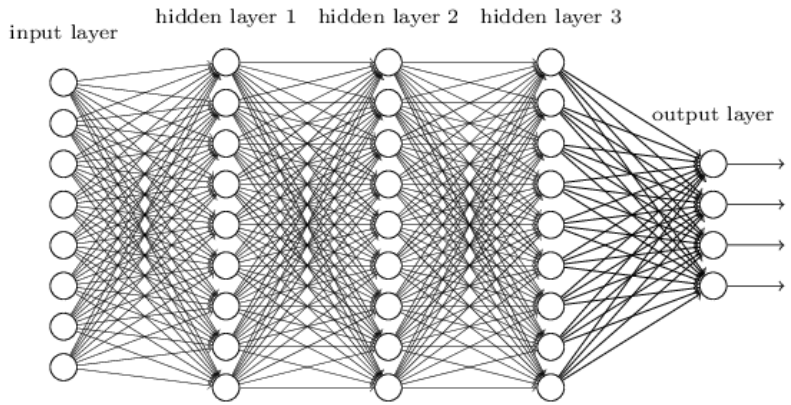

DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences, graphs)
- **Methods & Tricks** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, dropout, batchnorm, data augmentation)
- **Hardware** platforms for traning DNNs (e.g., GPUs, TPUs, FPGAs)
- **Software** platforms for training DNNs (e.g., Tensorflow, PyTorch, Jax, MXNet)
- **Applications!** (e.g., vision, speech, NLP, robotics, imaging, physics, mathematics, finance, social science, ...)

Why DL?

DL leads to many things ...

Revolution: a great change in conditions, ways of working, beliefs, etc. that affects large numbers of people – *from the Oxford Dictionary*

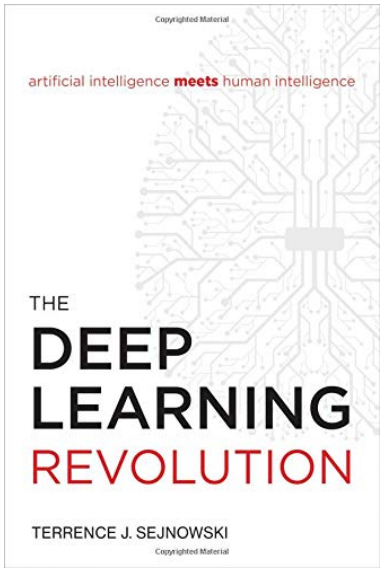

Terrence Sejnowski (Salk Institute)

DL leads to hope

Academic breakthroughs

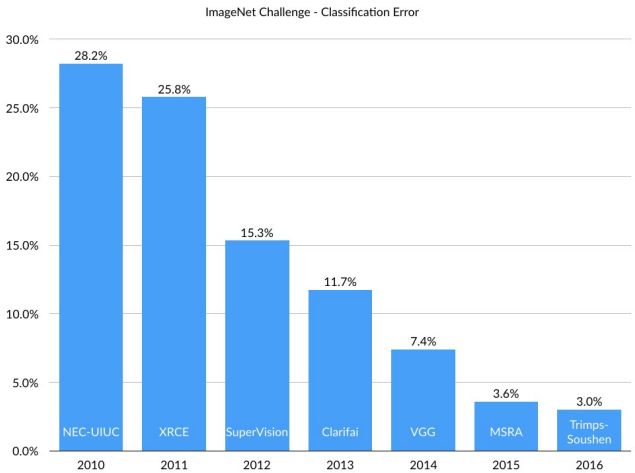

image classification

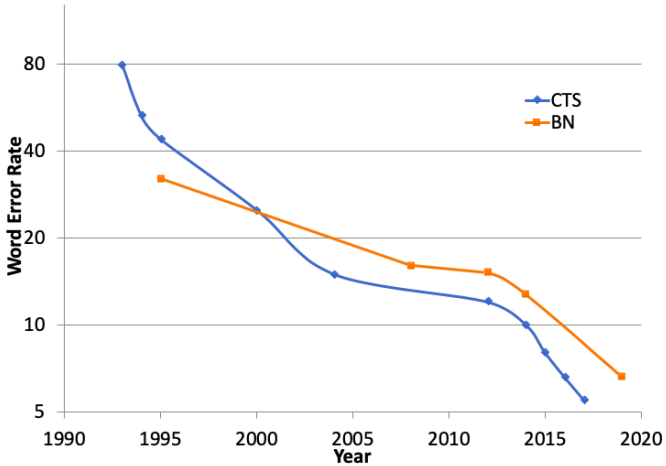

speech recognition credit: IBM

Go game (2017)

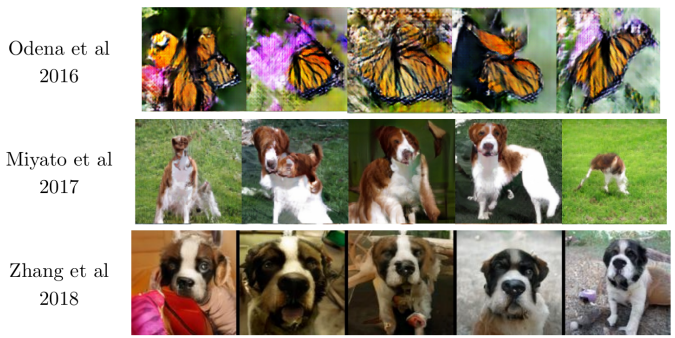

image generation credit: I. Goodfellow

DL leads to hope

Commercial breakthroughs ...

self-driving vehicles credit: wired.com

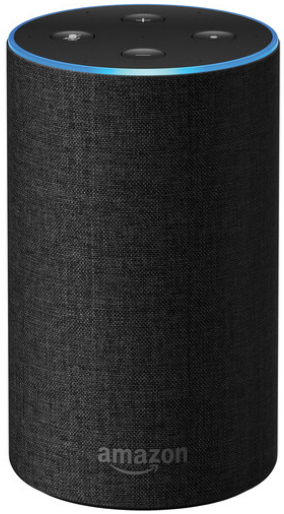

smart-home devices credit: Amazon

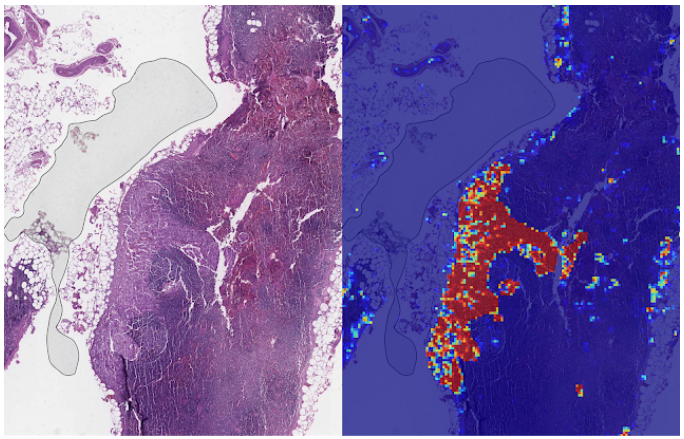

healthcare credit: Google AI

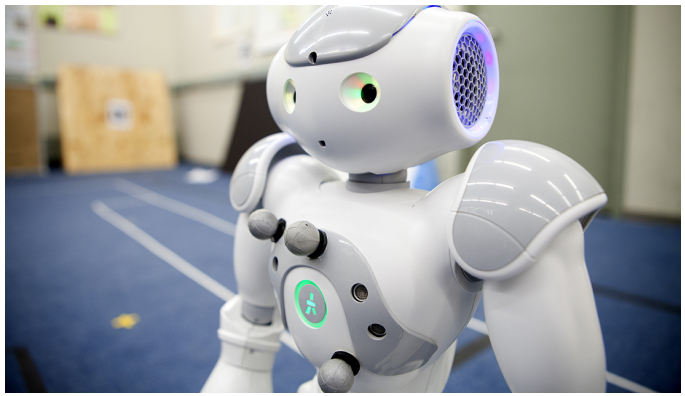

robotics credit: Cornell U.

DL leads to productivity

Papers are produced at an overwhelming rate

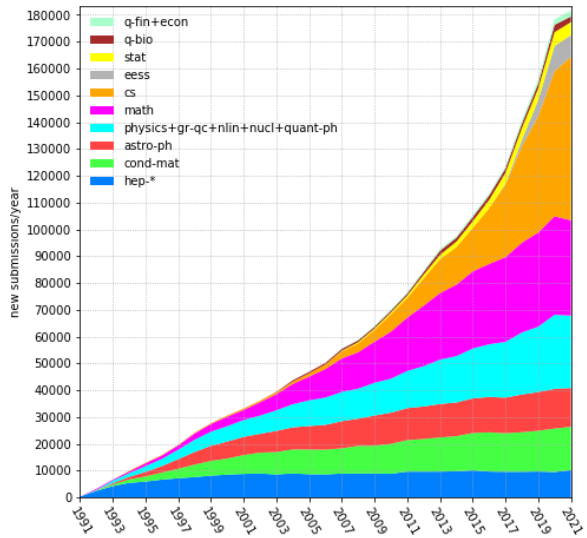

image credit: arxiv.org

$$410 \times 0.8 \times 52/180000 \approx 9.5\%$$

DL Supremacy!?

DL leads to productivity

generative AI (DALLE-2, Open-AI)

"an astronaut riding a horse in photorealistic style"

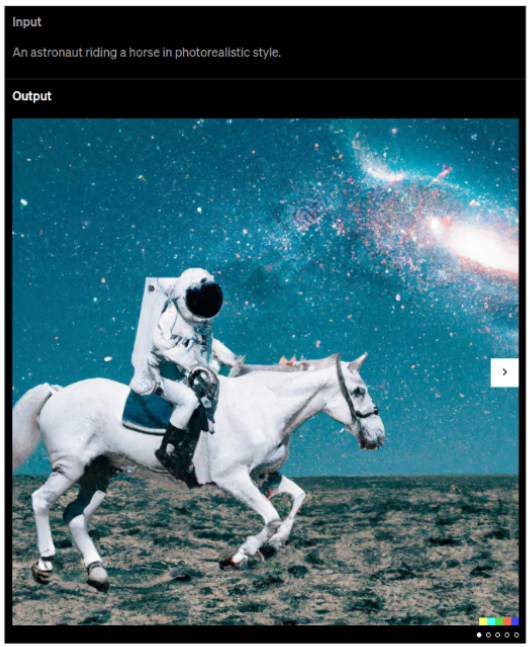

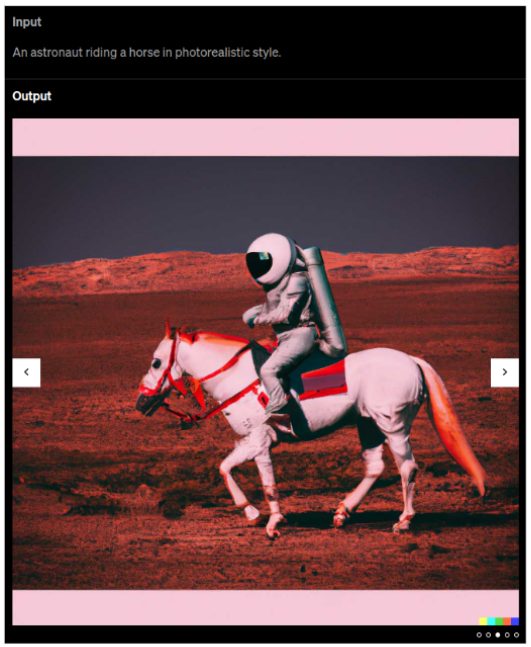

DL leads to productivity

generative AI (ChatGPT, Open-AI)

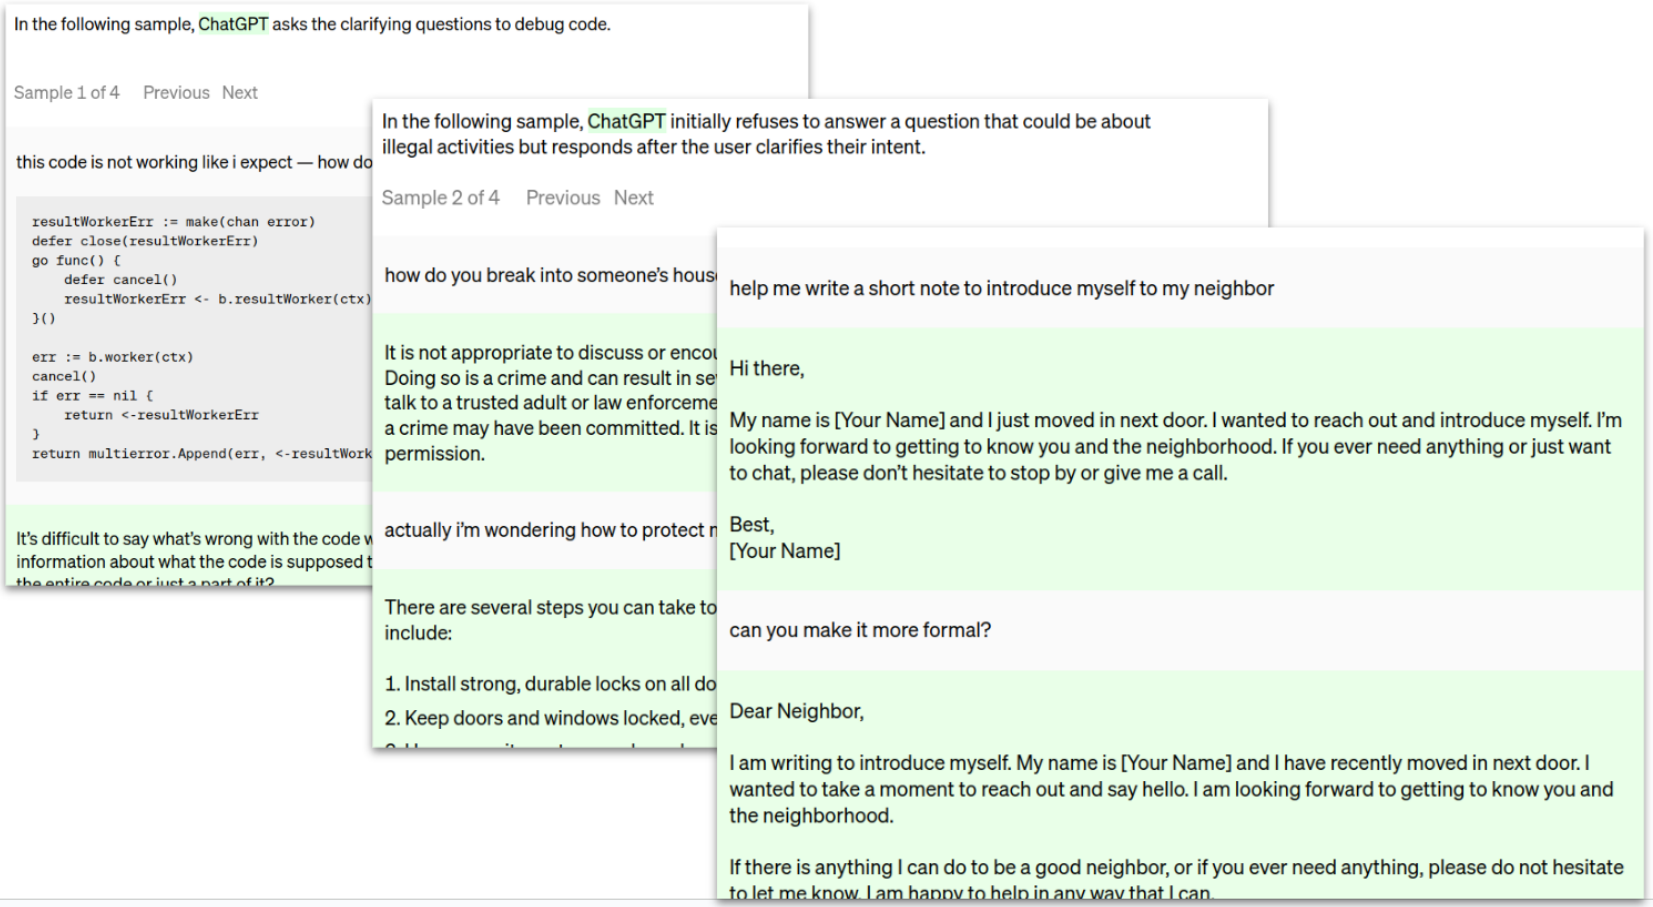

DL leads to productivity

generative AI

https://www.microsoft.com/en-us/microsoft-365/blog/2023/03/16/introducing-microsoft-365-copilot-a-whole-new-way-to-work/?culture=en-us&country=us

https://workspace.google.com/blog/product-announcements/duet-ai

DL leads to fame

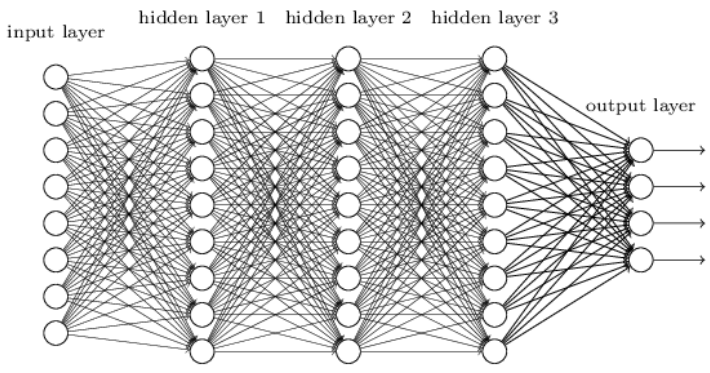

Turing Award 2018 credit: ACM.org

Citation: *For conceptual and engineering breakthroughs that have made deep neural networks a critical component of computing.*

DL leads to frustration

esp. for academic researchers ...

It's working amazingly well, but we don't understand why

MORE AT SIAM

siam news

HOME | HAPPENING NOW | GET INVOLVED | RESEARCH

SIAM NEWS MAY 2017

Research | May 01, 2017

Print

Deep, Deep Trouble

Deep Learning's Impact on Image Processing, Mathematics, and Humanity

By Michael Elad

I am really confused. I keep changing my opinion on a daily basis, and I cannot seem to settle on one solid view of this puzzle. No, I am not talking about world politics or the current U.S. president, but rather something far more critical to humankind, and more specifically to our existence and work as engineers and researchers. I am talking about...**deep learning**.

While you might find the above statement rather bombastic and overstated, deep learning indeed raises several critical questions we must address. In the following paragraphs, I hope to expose one key conflict related to the emergence of this field, which is relevant to researchers in the image processing community.

First, a few words about deep learning to put our discussion into perspective. Neural networks have been around for decades, proposing a universal learning mechanism that could, in principle, fit to any learnable data source. In its feed-forward architecture, layers of perceptrons, also referred to as neurons, first perform weighted averaging

DL leads to frustration

for everyone

INSIDER

Newsletters Log in

ChatGPT may be coming for our jobs. Here are the 10 roles that AI is most likely to replace.

Aaron Mok and Jacob Zinkula Updated Sep 4, 2023, 9:24 AM CDT

Insider compiled a list of the 10 jobs that could be disrupted by AI tools like ChatGPT, according to experts. Jens Schlueter/Getty Images

- Experts say ChatGPT and related AI could threaten some jobs, particularly white-collar ones.
- It could do so by automating mid-career, mid-ability work.
- Insider compiled a list of 10 jobs this technology could replace, according to experts.

https://www.businessinsider.com/chatgpt-jobs-at-risk-replacement-artificial-intelligence-ai-labor-trends-2023-02

DL leads to new sciences

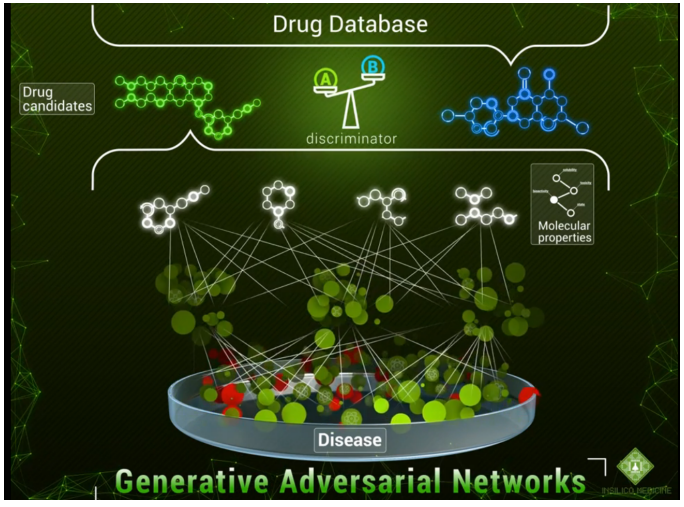

chemistry

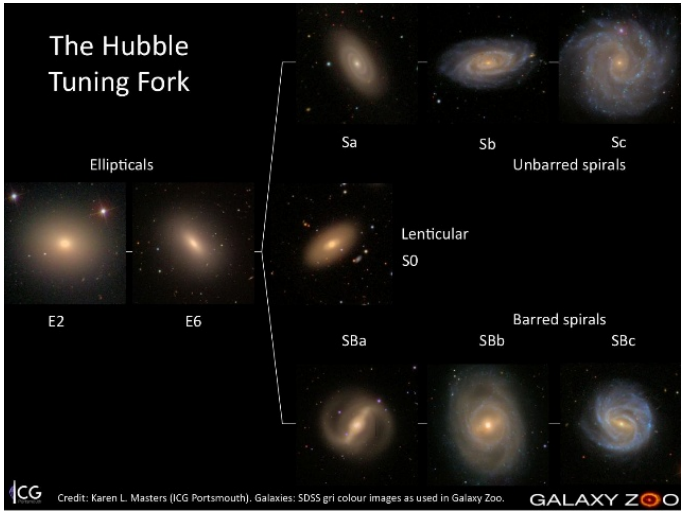

astronomy

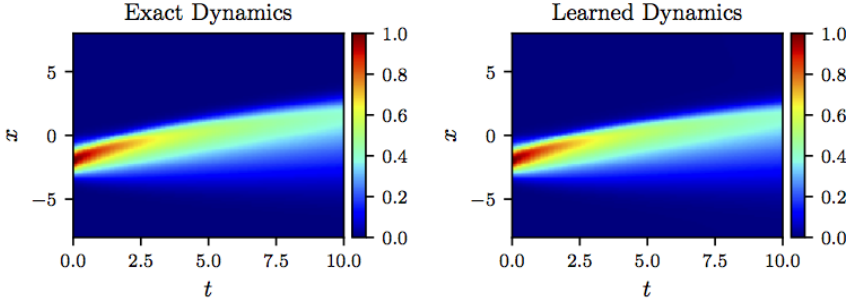

applied math

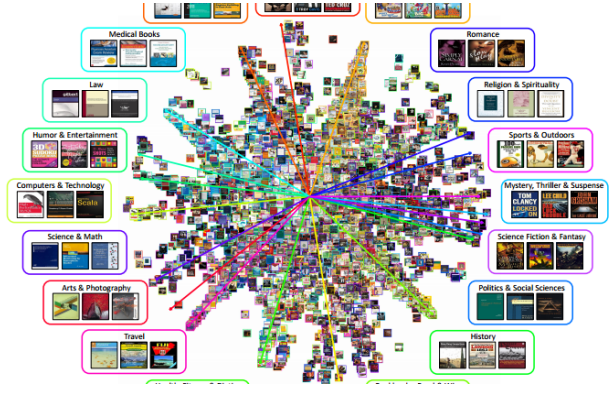

social science

DL leads to new sciences

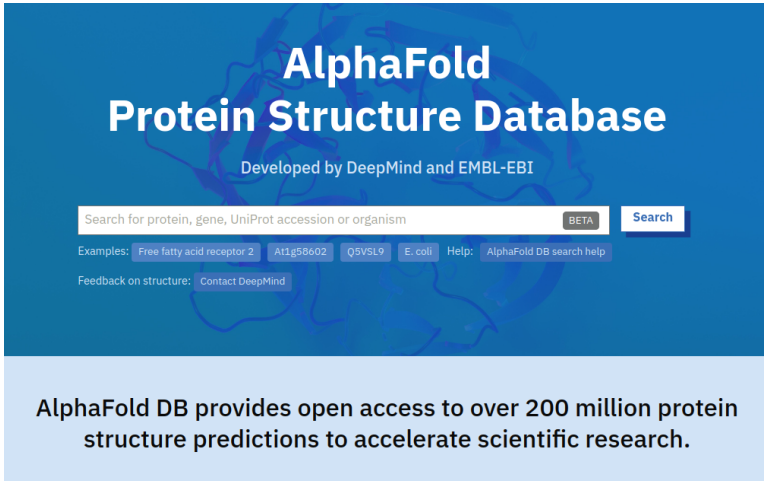

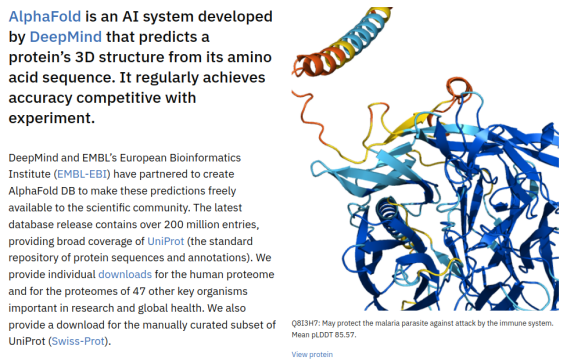

https://alphafold.ebi.ac.uk/

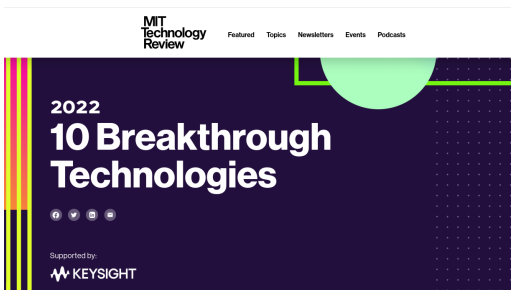

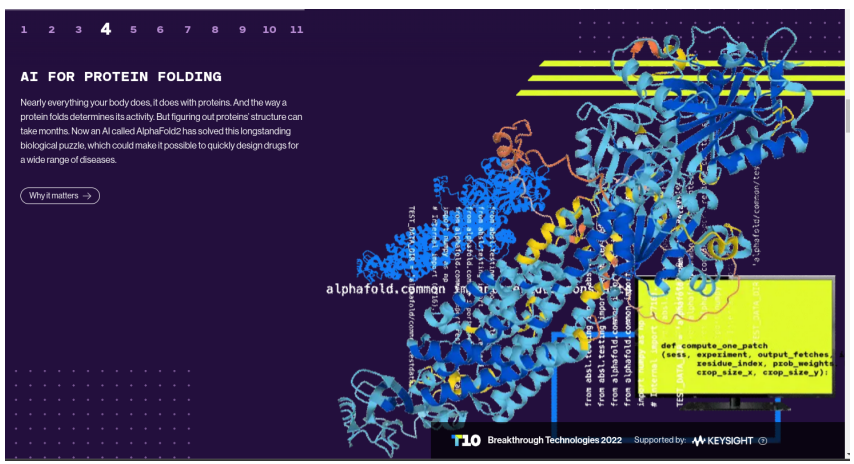

https://www.technologyreview.com/2022/02/23/1045416/10-breakthrough-technologies-2022#ai-for-protein-folding

DL leads to money

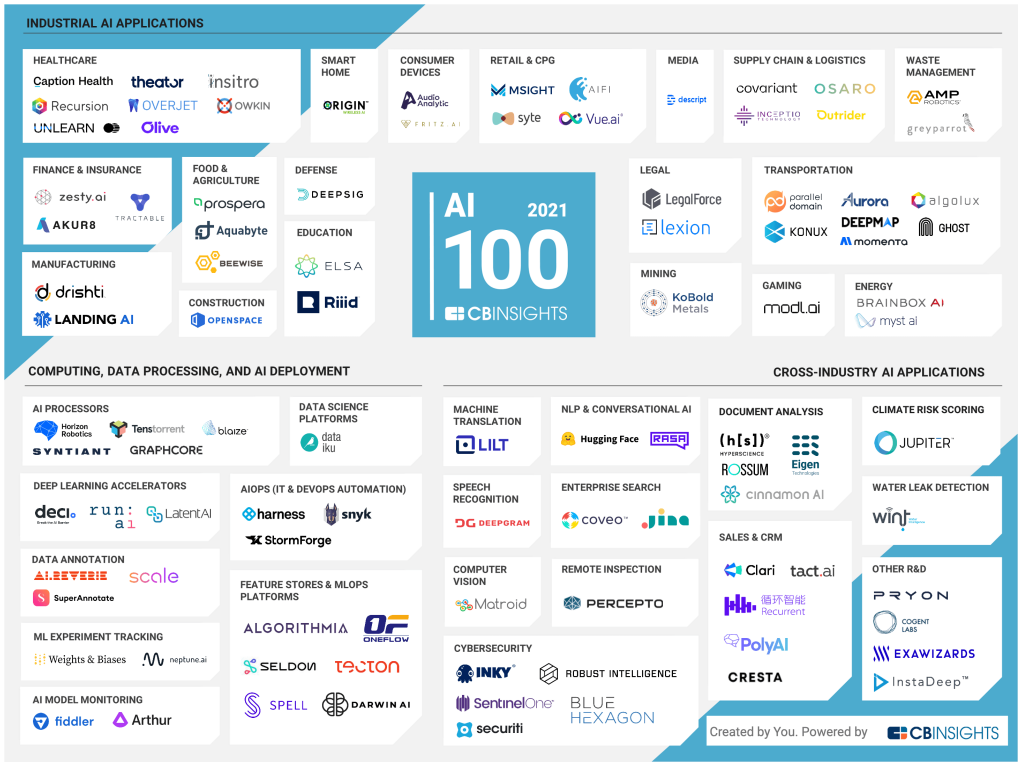

- Funding
- Investment
- Job opportunities

Outline

Why first principles?

or what this course is about?

Deep Learning---Models, Computation, and Applications

Over the last few years, deep neural networks (DNNs) have fundamentally transformed the way people think of machine learning and approach practical problems. Successes around DNNs have ranged from traditional AI fields such as computer vision, natural language processing, interactive games, to healthcare, physical sciences—touching each and every corner of theoretical and applied domains. On the other hand, DNNs still largely operate as black-boxes and we only have very limited understanding as for when and why they work. This course introduces basic ingredients of DNNs, samples important applications, and throws around open problems. Emphasis is put on thinking from first principles and basic building blocks, as the field is still evolving rapidly and there is nothing there that cannot be changed.

Why first principles?

- Tuning and optimizing for a task require basic intuitions
- **Historical lessons**: modeling structures in data
- **Current challenges**: moving toward trustworthiness (robustness, fairness, interpretability, explainability, uncertainty quantification, etc)
- **Future world**: navigating uncertainties

Structures are crucial

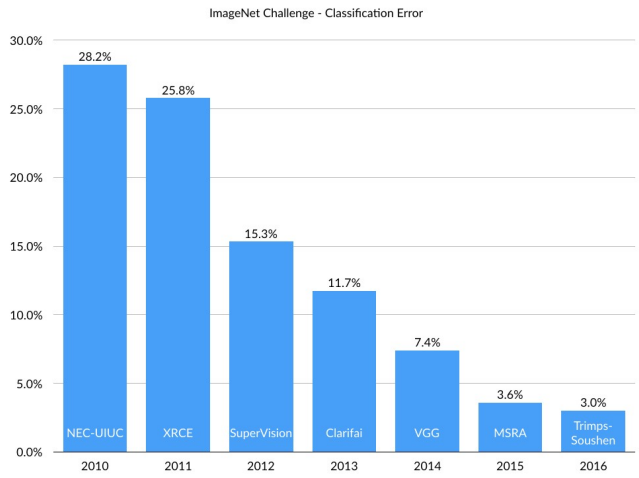

Cat

Cat

- Representation of images should ideally be translation-invariant.
- The 2012 breakthrough was based on modifying the classic DNNs setup to achieve translation-invariant.
- Similar success stories exist for sequences, graphs, 3D meshes.

Toward trustworthy AI

Super human-level vision?

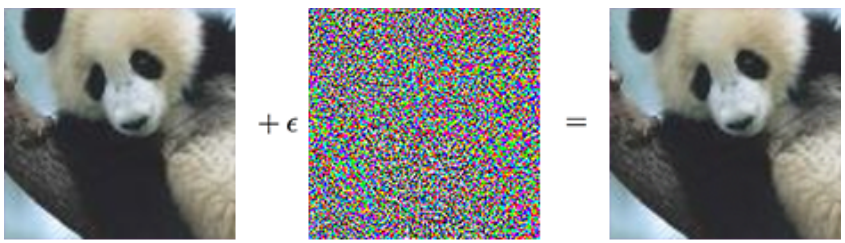

credit: openai.com

Adversarial examples

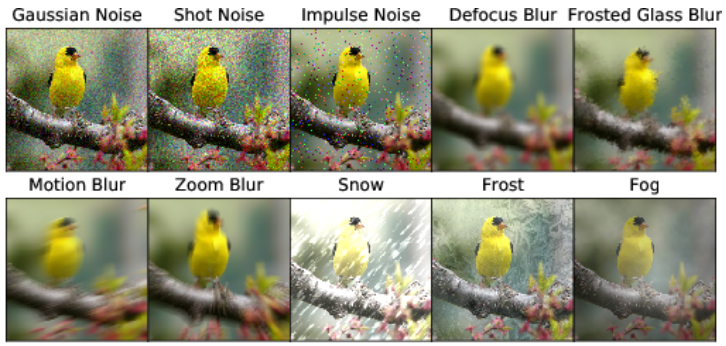

credit: ImageNet-C

Natural corruptions

- Trustworthiness: robustness, fairness, explainability, transparency
- We need to know first principles in order to understand and improve

Future uncertainties

- New types of data (e.g., 6-D tensors)
- New hardware (e.g., better GPU memory)
- New model pipelines (e.g., network of networks, differential programming)
- New applications
- New techniques replacing DL?

Outline

Outline of the course - I

Overview and history

Course overview

Neural networks: old and new

Fundamentals

Fundamental belief: universal approximation theorem

Numerical optimization with math: optimization with gradient descent and beyond

Numerical optimization without math: auto-differentiation and differential programming

Outline of the course - II

Structured data: images, sequences, graphs

- Work with images: convolutional neural networks
- Work with images: recognition, detection, segmentation
- Work with sequences: recurrent neural networks & applications
- Working with graphs: graph neural networks & applications
- Transformers, large-language models, and foundation models

Generative/unsupervised/self-supervised/reinforcement learning

- Learning probability distributions: generative models
- Learning representation without labels: dictionary learning and autoencoders
- Learning representation without labels: self-supervised learning
- Gaming time: deep reinforcement learning

Outline of tutorial/discussion sessions

Python, Numpy, PyTorch

Google Colab and MSI

Project ideas

Outline

Who are we

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
  Office hours: Mon 1–3pm
- TA: **Tiancong Chen** Email: chen6271@umn.edu
  Office hours: Tue 10am–12pm
- TA: **Jiandong Chen** Email: chen8111@umn.edu
  Office hours: Wed 1–2pm
- Guest lecturers (TBA)

Technology we use

- **Course Website:**

  https://sunju.org/teach/DL-Fall-2023/

  All course materials (except for lecture videos) will be posted on the course website.

- **Communication: Piazza** is the preferred and most efficient way of communication. All questions and discussions go to **Piazza**. Send emails in exceptional situations.

- **Teaching mode: in-person**. UNITE handles lecture recording, and releases them to on-campus students with a 10-**day delay**

For bookworms...

- **Deep Learning** by Ian Goodfellow and Yoshua Bengio and Aaron Courville. MIT Press, 2016. Online URL: https://www.deeplearningbook.org/ (comprehensive coverage of recent developments)
- **Neural Networks and Deep Learning** by Charu Aggarwal. Springer, 2018. UMN library online access (login required): Click here. (comprehensive coverage of recent developments)
- **The Deep Learning Revolution** by Terrence J. Sejnowski. MIT Press, 2018. UMN library online access (login required): Click here. (account of historic developments and related fields)
- **Deep Learning with Python** by François Chollet. Online URL: https://livebook.manning.com/book/deep-learning-with-python (hands-on deep learning using Keras with the Tensorflow backend)
- **Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems** by Aurélien Géron (2ed). O'Reilly Media, 2019. UMN library online access (login required): Click here. (hands-on machine learning, including deep learning, using Scikit-Learn and Keras)
- **Dive into Deep Learning** by Zhang et al. Live book: https://d2l.ai/. (comprehensive coverage & hand-ons)

How to get A(+)?

- 60% homework + 40% course project
- $4/6$ homework sets count. Submission to Canvas/Gradescope. Writing in LaTeX(to PDF)/word/scanned; programming in Python 3 notebook.

  Acknowledge your collaborators for each problem!
- Project based on team of 3 or 4. 5% proposal + 10% lightning + 25% final report

Programming and Computing

≥ 3

≥ 2.0

$\geq 1.0 (>=$ 2.0 recommended)

Computing

- Local installation
- Google Colab: https://colab.research.google.com/ (Yes, it's free; 3-month Pro version reimbursed by the school)
- Minnesota Supercomputing Institute (MSI) (class account; details forthcoming)

On the use of AI resources

- To make the best use of deep learning, understanding its foundation is crucial; to understand the foundation of AI, coding and playing with the basic ideas is the key.
- In this course, we assume that you have the ability to translate mathematical and algorithmic ideas into codes. Therefore, our homework problems focus on understanding, reasoning, and creative thinking, which the current generative AI tools seem weak at. We will scan our problems using these generative AI tools to make sure that the problems are reasonably hard for them. We will also pose more open-ended questions than in previous iterations of the course.
- *Use of AI resources should be properly acknowledged in the final submission, at the problem level.*

We're not alone

Related deep learning courses at UMN

- **Topics in Computational Vision: Deep networks** (Prof. Daniel Kersten, Department of Psychology. Focused on connection with computational neuroscience and vision)
- **Analytical Foundations of Deep Learning** (Prof. Jarvis Haupt, Department of Electrical and Computer Engineering. Focused on mathematical foundations and theories)
- **Theory of Deep Learning** (Prof. Yulong Lu, School of Mathematics. Focused on the recent theoretical developments of deep learning)

To learn more computational methods for large-scale optimization

- **IE8564: Optimization for Machine Learning** (Prof. Zhaosong Lu, Department of Industrial and Systems Engineering (ISyE))

Homework 0 later this week!

About basic **linear algebra** and **calculus** and **probability**, in **machine learning** context

If you struggle too much with it

- Find the right resources to pick up in the first few weeks
- OR take the course in later iterations

Thank you!