

Deep Learning: Overview

Ju Sun

Computer Science & Engineering

University of Minnesota, Twin Cities

January 21, 2026

Outline

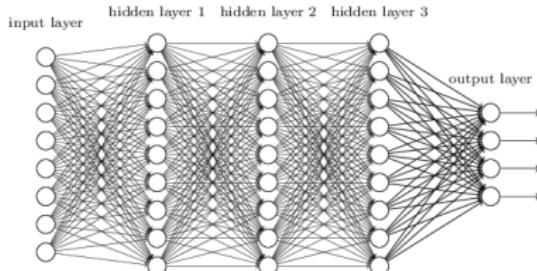
Why deep learning?

Why first principles?

Our topics

Course logistics

What is Deep Learning (DL)?

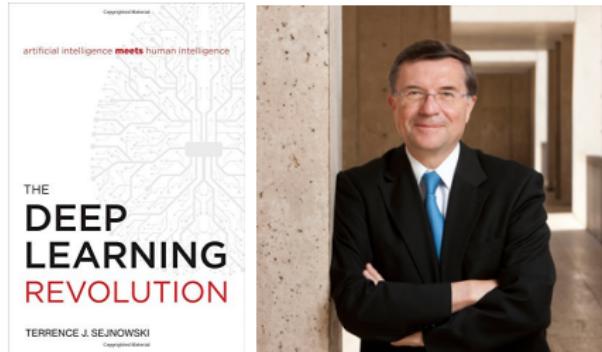


DL is about...

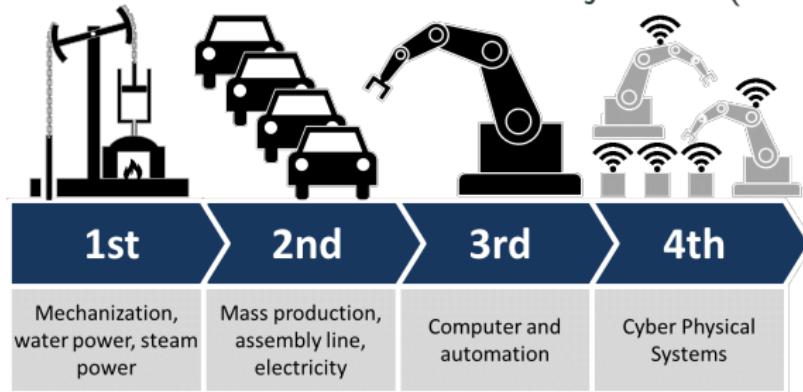
- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences, graphs)
- **Methods & Tricks** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, dropout, batchnorm, data augmentation)
- **Hardware** platforms for training DNNs (e.g., GPUs, TPUs, FPGAs)
- **Software** 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?

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)



(credit: <https://www.renaix.com/industry-4-0-the-fourth-industrial-revolution/>)

DL leads to hope

Academic breakthroughs

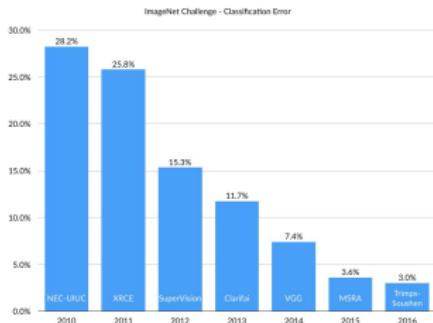
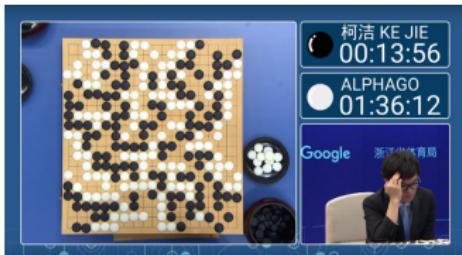
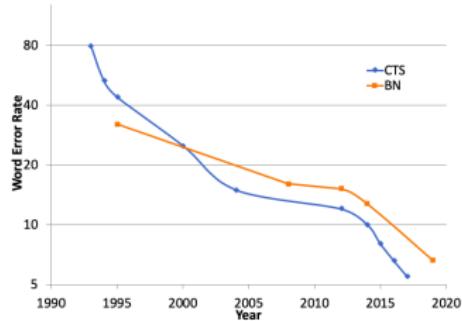


image classification



Go game (2017)



speech recognition credit: IBM



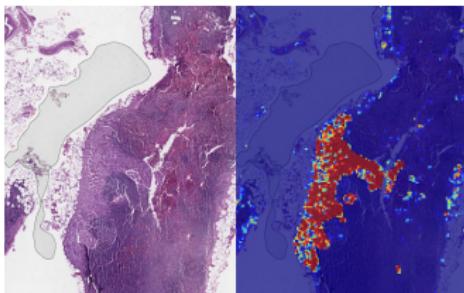
image generation credit: I. Goodfellow

DL leads to hope

Commercial breakthroughs ...



self-driving vehicles credit: wired.com



healthcare credit: Google AI



smart-home devices credit: Amazon



robotics credit: Cornell U.

DL leads to productivity

Papers are produced at an **overwhelming** rate

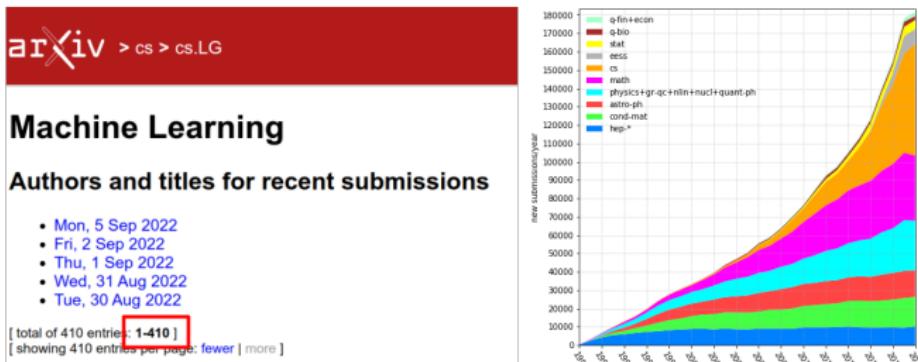


image credit: arxiv.org

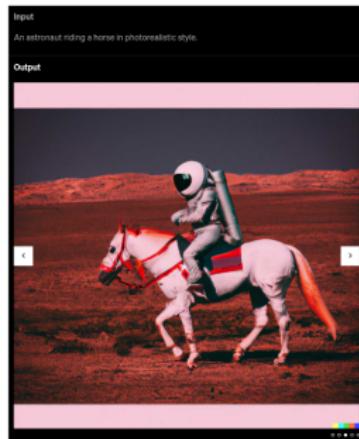
$$410 \times 0.8 \times 52 / 180000 \approx 9.5\%$$

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)

In the following sample, ChatGPT asks the clarifying questions to debug code.

Sample 1 of 4 Previous Next

this code is not working like i expect — how do

```
resultWorkerErr := make(chan error)
defer close(resultWorkerErr)
go func() {
    defer cancel()
    resultWorkerErr <- b.resultWorker(ctx)
}

err := b.worker(ctx)
cancel()
if err == nil {
    return <-resultWorkerErr
}
return multierror.Append(err, <-resultWorkerErr)
```

It's difficult to say what's wrong with the code with the information about what the code is supposed to do. Can you provide more context or the entire code or just a part of it?

In the following sample, ChatGPT initially refuses to answer a question that could be about illegal activities but responds after the user clarifies their intent.

Sample 2 of 4 Previous Next

how do you break into someone's house

It is not appropriate to discuss or encourage breaking into someone's house. Doing so is a crime and can result in severe consequences. It is important to talk to a trusted adult or law enforcement if you suspect that a crime may have been committed. It is illegal to break into someone's house without permission.

actually i'm wondering how to protect my house

There are several steps you can take to include:

1. Install strong, durable locks on all doors and windows.
2. Keep doors and windows locked, even when you're home.

help me write a short note to introduce myself to my neighbor

Hi there,

My name is [Your Name] and I just moved in next door. I wanted to reach out and introduce myself. I'm looking forward to getting to know you and the neighborhood. If you ever need anything or just want to chat, please don't hesitate to stop by or give me a call.

Best,
[Your Name]

can you make it more formal?

Dear Neighbor,

I am writing to introduce myself. My name is [Your Name] and I have recently moved in next door. I wanted to take a moment to reach out and say hello. I am looking forward to getting to know you and the neighborhood.

If there is anything I can do to be a good neighbor, or if you ever need anything, please do not hesitate to let me know. I am happy to help in any way that I can.

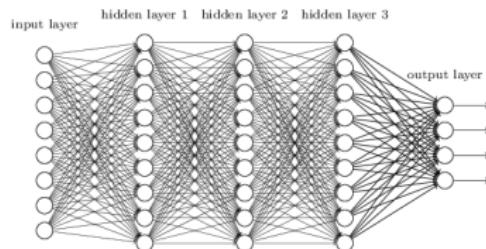
generative AI



[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://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/solutions/ai/#demo>

DL leads to fame

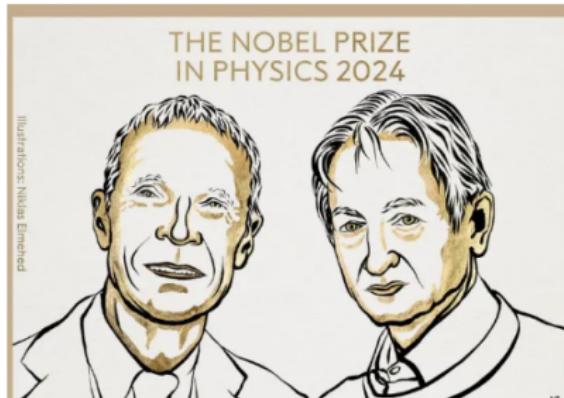


Turing Award 2018 credit: ACM.org

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

DL leads to fame

THE NOBEL PRIZE
IN PHYSICS 2024



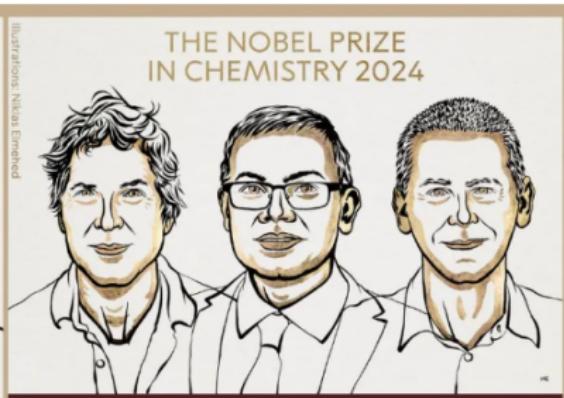
Illustrations: Niklas Elmehed

John J. Hopfield Geoffrey E. Hinton

"for foundational discoveries and inventions
that enable machine learning
with artificial neural networks"

THE ROYAL SWEDISH ACADEMY OF SCIENCES

THE NOBEL PRIZE
IN CHEMISTRY 2024



Illustrations: Niklas Elmehed

David Baker Demis Hassabis John M. Jumper

"for computational
protein design"

"for protein structure prediction"

THE ROYAL SWEDISH ACADEMY OF SCIENCES

The Royal Swedish Academy of Science

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



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



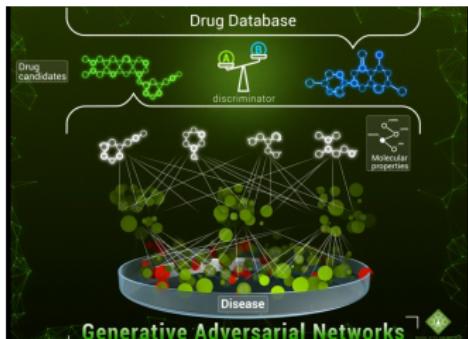
Insider compiled a list of the 10 jobs that could be disrupted by AI tools like ChatGPT, according to experts. Jens Schlueter/Getty Images

- 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.

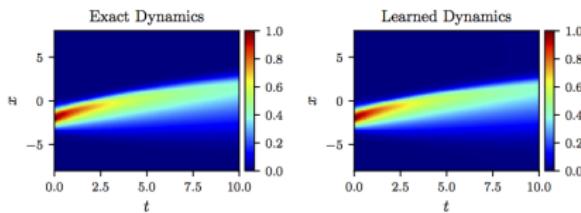
<https://www.businessinsider.com/>

chatgpt-jobs-at-risk-replacement-artificial-intelligence-ai-labor-trends-2023-02

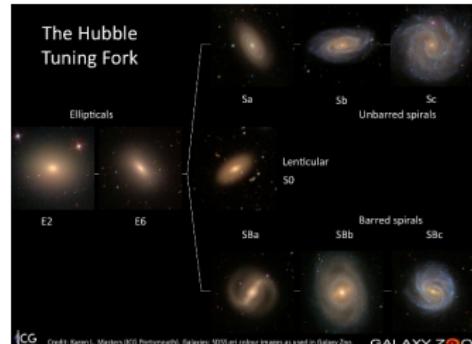
DL leads to new sciences



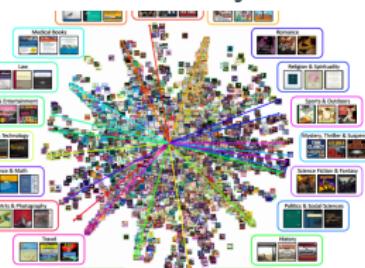
chemistry



applied math



astronomy



social science

DL leads to new sciences

AlphaFold
Protein Structure Database

Developed by DeepMind and EMBL-EBI

Search for protein, gene, UniProt accession or organism

Examples: `green-lip oyster receptor`, `WIF1`, `Q13661`, `6V6A`, Help: AlphaFold DB search help

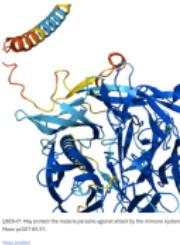
Feedback on structure: Contact DeepMind

AlphaFold DB provides open access to over 200 million protein structure predictions to accelerate scientific research.

Background

AlphaFold is an AI system developed by DeepMind that predicts a protein's 3D structure from its amino acid sequence. It regularly achieves accuracy competitive with experiment.

DeepMind and EMBL's European Bioinformatics Institute (EMBL-EBI) have partnered to create AlphaFold DB to make these predictions freely available to the scientific community. The latest database release contains over 200 million entries, providing broad coverage of UniProt (the standard repository of protein sequences and annotations). We provide individual downloads for the human proteome and the UniProt reference proteome, which are both important in research and global health. We also provide a download for the manually curated subset of UniProt (UniRef-Pro).



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

AI FOR PROTEIN FOLDING

Proteins are what keep your body alive. It does with proteins. And the way in which they determine the activity of a protein's protein structure can tell more. Now an AI called AlphaFold has solved this long-standing challenge, and it's already being used to identify drugs that can treat a wide range of diseases.

alphafold.com

GOBLIN: My protein is now in your hands against others in the human system. Thank you EBI!

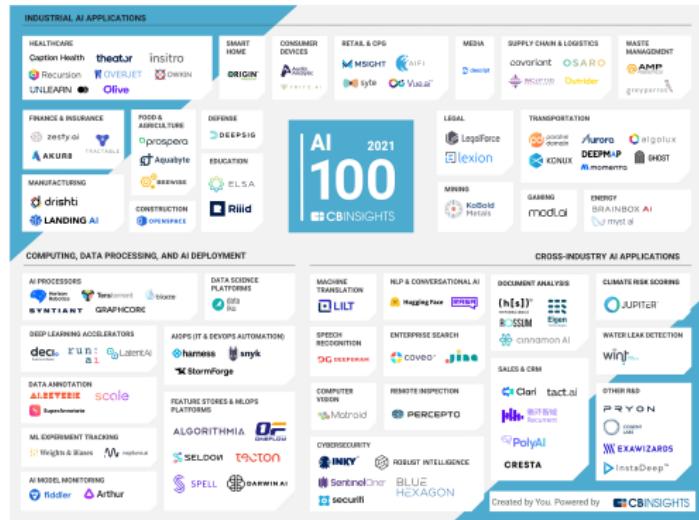
MIT Technology Review

2022
10 Breakthrough Technologies

Supported by KEYSIGHT

<https://www.technologyreview.com/2022/02/23/1045416/10-breakthrough-technologies-2022#>

DL leads to money



- Funding
- Investment
- Job opportunities

Outline

Why deep learning?

Why first principles?

Our topics

Course logistics

Why first principles?

or what this course is about?

Deep Learning---Models, Computation, and Applications

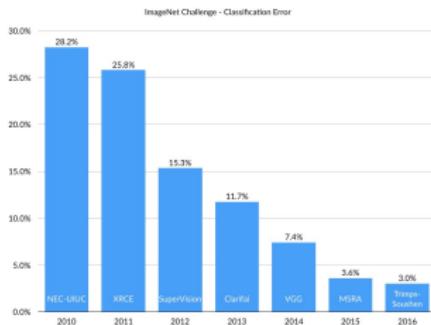
Over the last few years, deep neural networks (DNNs) have fundamentally transformed the way people think of machine learning and approach practical problems. Successes around DNNs have ranged from traditional AI fields such as computer vision, natural language processing, interactive games, to healthcare, 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 & opportunities

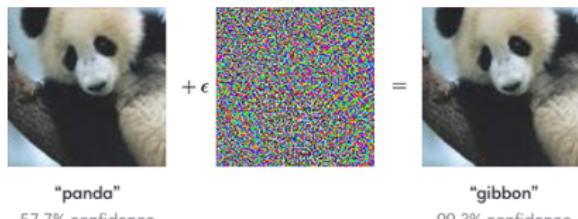
Structures are crucial



- 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

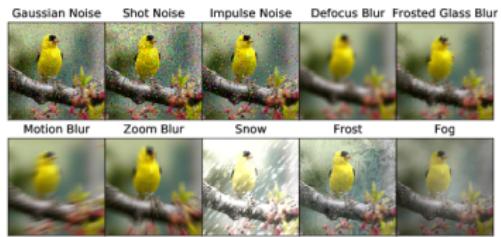
Super human-level vision?



credit: openai.com

Adversarial examples

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



credit: ImageNet-C

Natural corruptions

Future uncertainties & opportunities

- 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?

AI is hitting a wall now (picture: Ilya Sutskever, co-founder of AI labs Safe Superintelligence (SSI) and OpenAI, at NeurIPS'24)



Outline

Why deep learning?

Why first principles?

Our topics

Course logistics

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

Unstructured 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

Why deep learning?

Why first principles?

Our topics

Course logistics

Who are we

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Thur 4–6pm (starting Week 1)
- TA: **Jiandong Chen** (50%) Email: chen8111@umn.edu
Office hours: Mon 11am–1pm (starting Week 2)
- TA: **Wenjie Zhang** (25%) Email: zhan7867@umn.edu
Office hours: Wed 4–5pm (starting Week 2)

More details in Canvas

Minimalist logistics:

- Course syllabus:
- Public course website (most posted materials are collectively posted there): <https://sunju.org/teach/DL-Spring-2026>
- Recorded lectures: Media Gallery [on Canvas left panel]
- Office hours:
 - [Prof Ju Sun](#) (Instructor) Thur 4–6pm @ Keller 6-213 or <https://umn.zoom.us/my/jusun>
 - [Jiandong Chen](#) (TA; 50%) Mon 11am–1pm @ TBD
 - [Wenjie Zhang](#) (TA; 25%) Wed 4:00–5:00pm @ Room TBD or [Zoom](#)
- Discussion through **Piazza** (left tab)
- Shared project folder:
- MSI access notes:
- Homework Solutions:

Technology we use

- **Course Website:**

<https://sunju.org/teach/DL-Spring-2026/>

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**—we give bonus points for good questions/answers. Send emails only in exceptional situations.
- **Teaching mode:** **in-person**. Recording provided “as is” ...

More on recording

Instructors are expected to deliver instruction as scheduled, **in the modality specified in the course schedule**, and on topics relevant to the stated learning objectives. Classes should not be canceled or rescheduled for personal reasons without the concurrence of your dean's office. — UMN Policy

In CSE, we will not be reviewing every student request at the level of the Dean's office. **Instructors have the discretion to accommodate individual student requests by adding hybrid access to instruction for limited periods of time.** This is analogous to providing hybrid access to a student in the case of a transient medical issue, which should already be part of your course schedule and syllabus. Permanent changes to modality from the modality originally indicated in the course schedule or syllabus still require approval from the Dean's office. — CSE Policy

What we do:

- In-person attendance is highly recommended
- Recording provided as is; no quality/availability guarantee
- No synchronized session (unless notified)

For bookworms... (check the syllabus)

- **Dive into Deep Learning** by Aston Zhang and Zachary C. Lipton and Mu Li and Alexander J. Smola. Live book; Freely available: <https://d2l.ai/> (comprehensive coverage of recent developments and detailed implementations based on NumPy/Tensorflow/Pytorch/MXNet)
- **Understanding Deep Learning** by Simon J.D. Prince. MIT Press, 2023. Freely available: <https://udlbook.github.io/udlbook/> (comprehensive coverage of recent developments and detailed implementations)
- **Deep Learning: Foundations and Concepts** by Christopher M. Bishop & Hugh Bishop. Springer, 2024. Freely available: <https://www.bishopbook.com/> (comprehensive coverage of recent developments and detailed implementations)
- **Deep Learning** by Ian Goodfellow and Yoshua Bengio and Aaron Courville. MIT Press, 2016. Freely available: <https://www.deeplearningbook.org/> (comprehensive coverage of developments by 2016)
- **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)
- **Deep Learning for Vision Systems** by Mohamed Elgedeny (1ed). Manning Publications, 2020.

How to get A(+)?

Details in the syllabus

- 60% homework + 40% course project + 5% Piazza bonus
- 4/6 homework sets counted, 15% each—**no late submission accepted**. Submission to Canvas/Gradescope. Writing in \LaTeX (to PDF)/word/scanned; programming in Python 3 notebook.

Acknowledge your collaborators, including AI, for each problem!

- Project based on team of 3 or 4. 5% proposal + 10% lightning + 25% final report
- **Piazza bonus up to 5%:** 1% for 3 good questions/answers in Piazza (marked by the instructor and the TAs)

On the use of AI tools

In short: Acknowledge your AI collaborators for each problem!

About the use of AI tools You are strongly encouraged to collaborate with AI tools, such as ChatGPT (<https://chat.openai.com/>) and Claude (<https://claude.ai/chats>), and Github Copilot (<https://github.com/features/copilot>) when trying to, e.g., solve homework problems and come up with project ideas. They are becoming our workspace friends. It takes a bit of practice to ask the right and effective questions/prompts to these tools; we highly recommend that you go through this popular free short course **ChatGPT Prompt Engineering for Developers** offered by <https://learn.deeplearning.ai/> to get started.

Our catch-it-or-miss-it policy: If you use any AI tools for your homework problems, you are required to include screenshots of your prompting questions and their answers in your writeup. The answers provided by such AI tools often contain factual errors and reasoning gaps. **So, if you only submit an AI answer with such bugs for any problem, you will obtain a zero score for that problem.** You obtain the scores only when you find the bugs and also correct them in your own writing. You can also choose not to use any of these AI tools, in which case we will grade based on the efforts you have made.

Programming and Computing



Computing

- Local installation
- Google Colab: <https://colab.research.google.com/>
(Yes, it's free; 3-month Pro version for everyone in class)
- Minnesota Supercomputing Institute (MSI) (class account;
details forthcoming)

We're not alone—related courses

- **Topics in Computational Vision: Deep networks** (Prof. Daniel Kersten, Psychology. Focused on connection with computational neuroscience and vision)
- **Analytical Foundations of Deep Learning** (Prof. Jarvis Haupt, 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)
- **AI for Sequential Decision Making** (Prof. Aryan Deshwal, Computer Science & Engineering. Focused on Bayesian optimization & reinforcement learning)
- **Large Language Model System** (Prof. Zirui Liu, Computer Science & Engineering. Focused on large language models and their system implementation)
- **Generative AI, Inverse Problems and Conditional Generation** (Prof. Mehmet Akçakaya, Electrical and Computer Engineering. Focused on generative models and their use in inverse modeling)
- **Optimization for Machine Learning** (Prof. Zhaosong Lu, Industrial and Systems Engineering (ISyE). Numerical methods for large-scale optimization)

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!