

Deep Learning: Overview

Ju Sun

Computer Science & Engineering

University of Minnesota, Twin Cities

September 6, 2022


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?

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



Go game (2017)



speech recognition credit: IBM





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 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 reform weighted outputs




DL leads to new sciences




chemistry



astronomy




applied math



social science

DL leads to new sciences



AlphaFold
Protein Structure Database

Developed by DeepMind and EMBL-EBI

Search for proteins, genes, UniProt accession or organism

Examples: [Free fatty acid receptor 2](#) [ADP/ATPase](#) [Glyceral](#) [Lactate](#) [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 for the proteomes of 47 other key organisms important in research and global health. We also provide a download for the manually curated subset of UniProt (Swiss-Prot).



EMBL-EBI may protect the trade name logo against attack by the immune system. News 2022/04/01 News 2022/04/01


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



MIT Technology Review

2022
10 Breakthrough Technologies

KEYSIGHT



1 2 3 4 5 6 7 8 9 10 11

AI FOR PROTEIN FOLDING

Nearly every living organism has a 3D shape. And the way a protein folds determines its activity. But figuring out protein structure can be difficult. The Nobel Prize-winning method has created the longstanding biological puzzle, which could make it possible to quickly design drugs for a broader range of diseases.



alphaFold.com@nd

T10 Breakthrough Technologies 2022 - Report by KEYSIGHT

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

[ai-for-protein-folding](#)

DL leads to money



- Funding
- Investment
- Job opportunities

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

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

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



credit: ImageNet-C

Natural corruptions

- 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?

Why deep learning?

Why first principles?

Our topics

Course logistics

Overview and history

Course overview (1)

Neural networks: old and new (1)

Fundamentals

Fundamental belief: universal approximation theorem (2)

Numerical optimization with math: optimization with gradient descent and beyond (2)

Numerical optimization without math: auto-differentiation and differential programming (2)

Outline of the course - II

Structured data: images, sequences, graphs

Work with images: convolutional neural networks (2)

Work with images: recognition, detection, segmentation (2)

Work with sequences: recurrent neural networks & applications (3)

Working with graphs: graph neural networks & applications (3)

Generative/unsupervised/self-supervised/reinforcement learning

Learning probability distributions: generative models (3)

Learning representation without labels: dictionary learning and autoencoders (2)

Learning representation without labels: self-supervised learning (2)

Gaming time: deep reinforcement learning (2)

Outline of tutorial/discussion sessions

Python, Numpy, and Google Cloud/Colab

Project ideas

Intro to Pytorch

Research ideas

Why deep learning?

Why first principles?

Our topics

Course logistics

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Mon 2–4pm
- TA: **Hengkang Wang** Email: wang9881@umn.edu
Office hours: Thur 1–3pm
- TA: **Yash Travadi** Email: trava029@umn.edu
Office hours: Wed 2–4pm
- Guest lecturers (TBA)

- **Course Website:**

<https://sunju.org/teach/DL-Fall-2022/>

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. (Unofficial Discord group available, but the instruction team won't be there)
- **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 L^AT_EX(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% mid-term presentation + 25% final report



≥ 3



≥ 2.0



≥ 1.0

Computing

- Local installation
- Google Colab: <https://colab.research.google.com/>
(Yes, it's free)
- Google Cloud (\$100 credits per student) (similarly AWS and Azure)
- Minnesota Supercomputing Institute (MSI) (class account; details forthcoming)

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)

To learn more computational methods for large-scale optimization

- **IE5080: Optimization Models and Methods 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!