

Think Deep Learning: Overview

Ju Sun

Computer Science & Engineering

University of Minnesota, Twin Cities

January 21, 2020


Why deep learning?

Why first principles?


Our topics

Course logistics

What is Deep Learning (DL)?




What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)


What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences)


What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences)
- **Methods** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, Dropout)


What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences)
- **Methods** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, Dropout)
- **Hardware** platforms for training DNNs (e.g., GPUs, TPUs, FPGAs)


What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences)
- **Methods** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, Dropout)
- **Hardware** platforms for training DNNs (e.g., GPUs, TPUs, FPGAs)
- **Software** platforms for training DNNs (e.g., Tensorflow, PyTorch, MXNet)

What is Deep Learning (DL)?



DL is about...

- Deep neural networks (DNNs)
- **Data** for training DNNs (e.g., images, videos, text sequences)
- **Methods** for training DNNs (e.g., AdaGrad, ADAM, RMSProp, Dropout)
- **Hardware** platforms for training DNNs (e.g., GPUs, TPUs, FPGAs)
- **Software** platforms for training DNNs (e.g., Tensorflow, PyTorch, MXNet)
- **Applications!** (e.g., vision, speech, NLP, imaging, physics, mathematics, finance)

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*

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)

Academic breakthroughs




image classification

DL leads to hope

Academic breakthroughs





image classification



speech recognition credit: IBM

DL leads to hope

Academic breakthroughs





image classification



chess game (2017)



speech recognition credit: IBM

DL leads to hope

Academic breakthroughs





image classification



chess game (2017)



speech recognition credit: IBM




image generation credit: I. Goodfellow

DL leads to hope

Commercial breakthroughs ...




self-driving vehicles credit: wired.com

DL leads to hope

Commercial breakthroughs ...



self-driving vehicles credit: wired.com




smart-home devices credit: Amazon

DL leads to hope


Commercial breakthroughs ...



self-driving vehicles credit: wired.com



smart-home devices credit: Amazon




healthcare credit: Google AI

DL leads to hope


Commercial breakthroughs ...



self-driving vehicles credit: wired.com



smart-home devices credit: Amazon



healthcare credit: Google AI



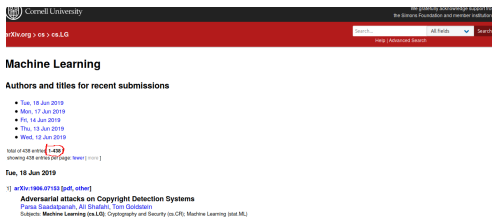
robotics credit: Cornell U.

DL leads to productivity

Papers are produced at an **overwhelming** rate

DL leads to productivity

Papers are produced at an **overwhelming** rate



The screenshot shows the Cornell University arXiv.org search results for 'Machine Learning'. The search bar contains 'cs.LG'. The results list authors and titles for recent submissions from June 12, 2019, to June 18, 2019. A red circle highlights the number '1430' in the text 'Total of 428 entries (1430)'. Below this, a specific entry is shown for a paper titled 'Adversarial attacks on Copyright Detection Systems' by Parisa Saadqapanah, Ali Shafahi, and Tom Goldstein, with the subject 'Machine Learning (cs.LG): Cryptography and Security (cs.CR); Machine Learning (stat.ML)'.




image credit: arxiv.org

DL leads to productivity

Papers are produced at an **overwhelming** rate






image credit: arxiv.org

$$400 \times 0.8 \times 52/140000 \approx 11.9\%$$


DL Supremacy!?

DL leads to fame



Turing Award 2018 credit: ACM.org

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 am keeping 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

DL leads to new sciences




chemistry




astronomy


DL leads to new sciences



chemistry




astronomy




applied math


DL leads to new sciences




chemistry



astronomy



applied math



social science

DL leads to money

Market summary >

NVIDIA Corporation

NASDAQ: NVDA

Overview News Compare Financials

248.24 USD -1.04 (0.42%) ↓

Jan 21, 11:07 AM EST · Disclaimer

1 day 5 days 1 month 6 months YTD 1 year 5 years Max



Open	247.80	Div yield	0.26%
High	249.00	Prev close	249.28

- Funding
- Investment
- Job opportunities

Why deep learning?

Why first principles?

Our topics

Course logistics

Why first principles?



PYTORCH

Why first principles?




- Tuning and optimizing for a task require basic intuitions

Why first principles?



- Tuning and optimizing for a task require basic intuitions
- **Historical lesson:** model structures in data
- **Current challenge:** move toward trustworthiness
- **Future world:** navigate uncertainties

Structures are crucial




Cat



Cat

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 improve and understand



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)

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)

Structured data: images and sequences

Work with images: convolutional neural networks (2)

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

Work with sequences: recurrent neural networks (2)

Structured data: images and sequences

Work with images: convolutional neural networks (2)

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

Work with sequences: recurrent neural networks (2)

Deterministic DNN

To train or not? scattering transforms (2)

Structured data: images and sequences

Work with images: convolutional neural networks (2)

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

Work with sequences: recurrent neural networks (2)

Deterministic DNN

To train or not? scattering transforms (2)

Other settings: generative/unsupervised/reinforcement learning

Learning probability distributions: generative adversarial networks (2)

Learning representation without labels: dictionary learning and autoencoders (1)

Gaming time: deep reinforcement learning (2)

Outline of tutorial/discussion sessions

Python, Numpy, and Google Cloud/Colab

Project ideas

Tensorflow 2.0 and Pytorch

Backpropagation and computational tricks

Research ideas

Why deep learning?

Why first principles?

Our topics

Course logistics

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Th 4–6pm 5-225E Keller H

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Th 4–6pm 5-225E Keller H
- TA: **Yuan Yao** Email: yaoxx340@umn.edu
Office hours: Wed 12:15–2:15pm at Shepherd Lab 234

Who are we

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Th 4–6pm 5-225E Keller H
- TA: **Yuan Yao** Email: yaoux340@umn.edu
Office hours: Wed 12:15–2:15pm at Shepherd Lab 234
- Courtesy TA: **Taihui Li** Email: lixx5027@umn.edu
who is responsible for setting up hard homework problems!

Who are we

- Instructor: **Professor Ju Sun** Email: jusun@umn.edu
Office hours: Th 4–6pm 5-225E Keller H
- TA: **Yuan Yao** Email: yaoux340@umn.edu
Office hours: Wed 12:15–2:15pm at Shepherd Lab 234
- Courtesy TA: **Taihui Li** Email: lixx5027@umn.edu
who is responsible for setting up hard homework problems!
- Guest lecturers (TBA)

- **Course Website:**

<https://sunju.org/teach/DL-Spring-2020/>

All course materials will be posted on the course website.

- **Course Website:**

<https://sunju.org/teach/DL-Spring-2020/>

All course materials will be posted on the course website.

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

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)

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)

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 (available soon). (hands-on machine learning, including deep learning, using Scikit-Learn and Keras)

How to get A(+)?

- 60 % homework + 40 % course project

How to get A(+)?

- 60 % homework + 40 % course project
- 5/7 homework counts. Submission to Canvas. Writing in \LaTeX (to PDF) and programming in Python 3 notebook.

Acknowledge your collaborators for each problem!

How to get A(+)?

- 60 % homework + 40 % course project
- 5/7 homework counts. Submission to Canvas. Writing in \LaTeX (to PDF) and programming in Python 3 notebook.

Acknowledge your collaborators for each problem!

- Project based on team of 2 or 3. 5% proposal + 10% mid-term presentation + 25% final report

How to get A(+)?

- 60 % homework + 40 % course project
- 5/7 homework counts. Submission to Canvas. Writing in \LaTeX (to PDF) and programming in Python 3 notebook.

Acknowledge your collaborators for each problem!

- Project based on team of 2 or 3. 5% proposal + 10% mid-term presentation + 25% final report
- Publish a paper \implies A!



≥ 3



≥ 2.0

PYTORCH

≥ 1.0



≥ 3



≥ 2.0



≥ 1.0

Computing

- Local installation
- Google Colab: <https://colab.research.google.com/>
(Yes, it's free)
- Google Cloud (\$50 credits per student) (similarly AWS and Azure)
- Minnesota Supercomputing Institute (MSI)

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)

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)

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 today!

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

Homework 0 today!

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 week
- **OR** take the course in later iterations

Thank you!