# What's scary about deep learning?

**Ju Sun, PhD**

Assistant Professor

Department of Computer Science & Engineering

Department of Neurosurgery

March 19, 2021

# Outline

What's good about deep learning?

Why are self-driving cars not delivered?

Are we making the best use of deep learning?

What about understanding DL?

Why AI for healthcare now?

Research of the SUN group

# Deep neural networks

Brain neural networks

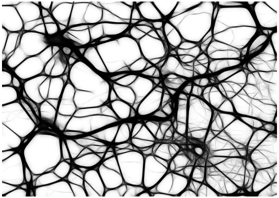

Credit: Max Pixel

# Deep neural networks

Brain neural networks

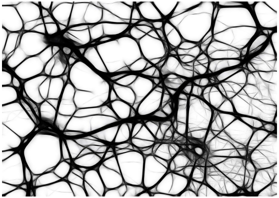

Credit: Max Pixel

Artificial neural networks

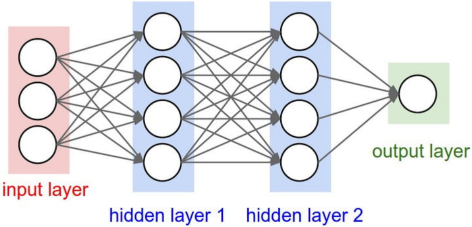

# Deep neural networks

Brain neural networks

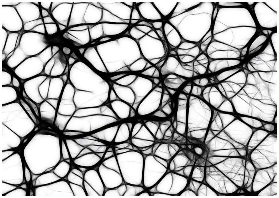

Credit: Max Pixel

Artificial neural networks

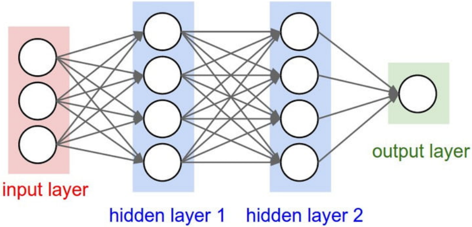


Why called **artificial**?

- (Over-)simplification on neural level
- (Over-)simplification on connection level

## Three pillars

# Three pillars


# Three pillars



Key ingredients of DL have been in place for 25-30 years:

| Landmark | Emblem | Epoch |
|---|---|---|
| Neocognitron | Fukushima | 1980 |
| CNN | Le Cun | mid 1980s' |
| Backprop | Hinton | mid 1980's |
| SGD | Le Cun, Bengio etc | mid 1990's |
| Various | Schmidhuber | mid 1980's |
| *CTF* | *DARPA etc* | *mid 1980's* |

# Three pillars



Key ingredients of DL have been in place for 25-30 years:

| Landmark | Emblem | Epoch |
|---|---|---|
| Neocognitron | Fukushima | 1980 |
| CNN | Le Cun | mid 1980s' |
| Backprop | Hinton | mid 1980's |
| SGD | Le Cun, Bengio etc | mid 1990's |
| Various | Schmidhuber | mid 1980's |
| *CTF* | *DARPA etc* | *mid 1980's* |

data + specialized hardware + specialized software

# What's good about 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*

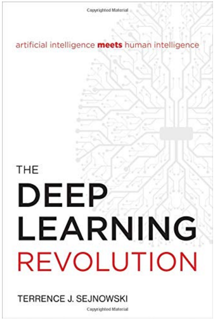


Terrence Sejnowski (Salk Institute)

# DL leads to hope

Academic breakthroughs

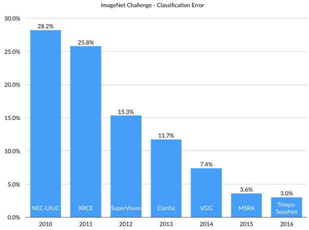


image classification

# DL leads to hope

Academic breakthroughs

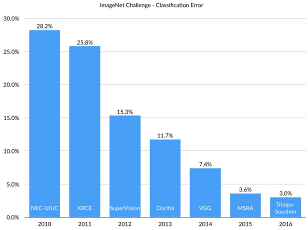


image classification

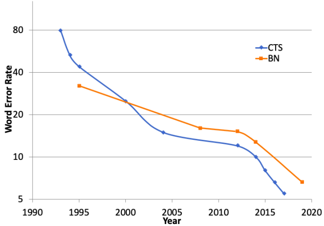


speech recognition credit: IBM

# DL leads to hope

Academic breakthroughs

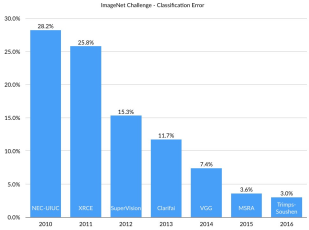


image classification

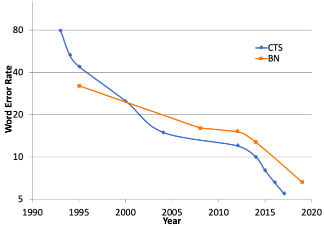


speech recognition credit: IBM


Go game (2017)

# DL leads to hope

## Academic breakthroughs

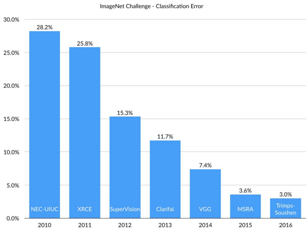


image classification

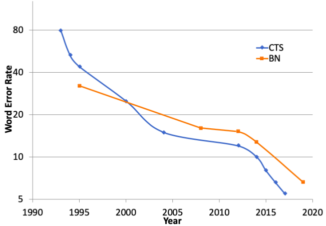


speech recognition credit: IBM


Go game (2017)

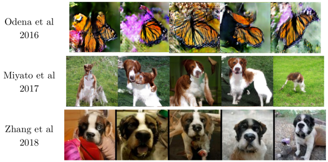


image generation credit: I. Goodfellow

# DL leads to hope

Commercial breakthroughs ...


self-driving vehicles credit: wired.com

# DL leads to hope

Commercial breakthroughs ...


self-driving vehicles credit: wired.com

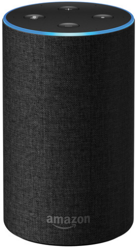


smart-home devices credit: Amazon

# DL leads to hope

Commercial breakthroughs ...


self-driving vehicles credit: wired.com

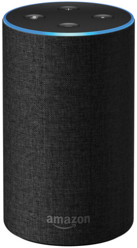


smart-home devices credit: Amazon

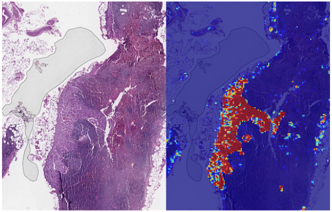

healthcare credit: Google AI

# DL leads to hope

Commercial breakthroughs ...


self-driving vehicles credit: wired.com

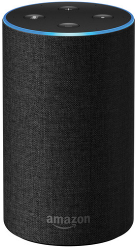


smart-home devices credit: Amazon

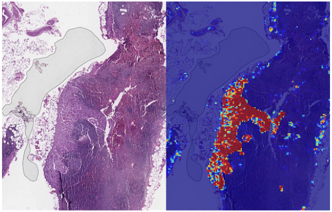

healthcare credit: Google AI

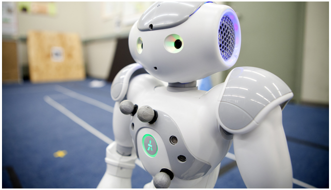

robotics credit: Cornell U.

# DL leads to productivity

Papers are produced at an overwhelming rate

# DL leads to productivity

Papers are produced at an overwhelming rate

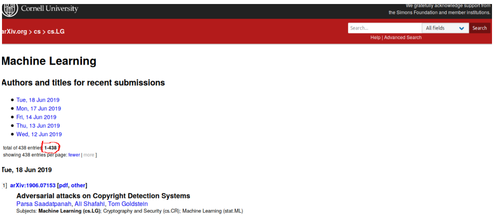


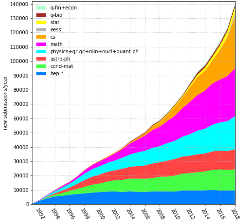


image credit: arxiv.org

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

DL Supremacy!?

# DL leads to fame

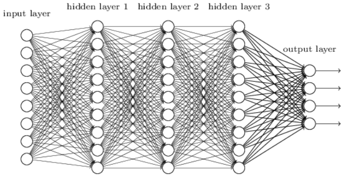



**Turing Award 2018** credit: ACM.org

# DL leads to fame

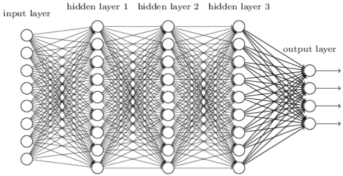



**Turing Award 2018** credit: ACM.org

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 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

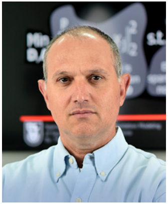

# DL leads to new sciences

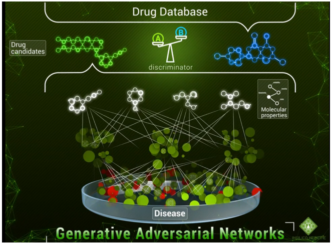


chemistry

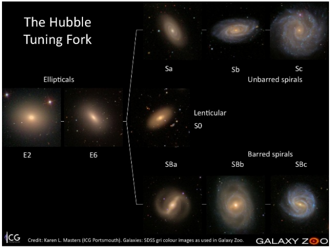


astronomy

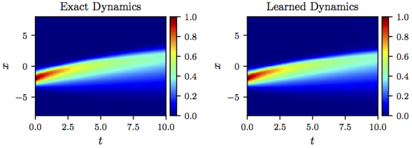


applied math

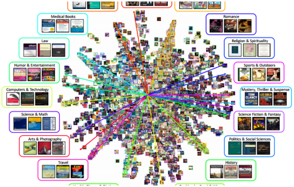

social science

# DL leads to money

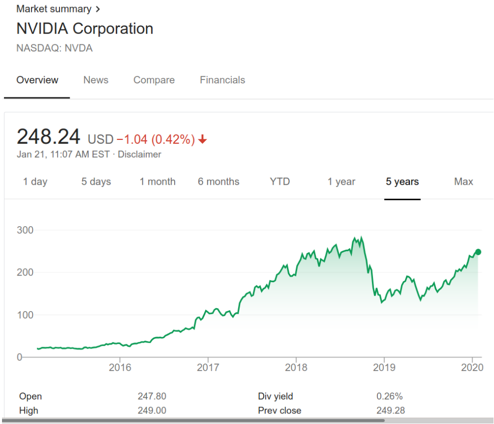


- Funding
- Investment
- Job opportunities

# A brief history of AI

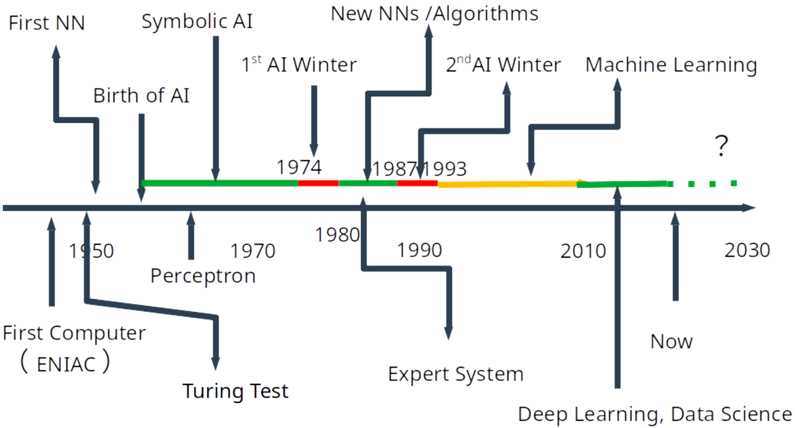

# Outline

What's good about deep learning?

Why are self-driving cars not delivered?

Are we making the best use of deep learning?

What about understanding DL?

Why AI for healthcare now?

Research of the SUN group

# Hype?

## Elon Musk says full self-driving Tesla tech 'very close'

9 July 2020


Tesla will be able to make its vehicles completely autonomous by the end of this year, founder Elon Musk has said.

BBC.com

# Hype?

## Elon Musk says full self-driving Tesla tech 'very close'

9 July 2020


**Tesla will be able to make its vehicles completely autonomous by the end of this year, founder Elon Musk has said.**

BBC.com

Forbes

## Why Is Tesla's Full Self-Driving Only Level 2 Autonomous?

**James Morris** Contributor
Green Tech
*I write about the rapidly growing world of electric vehicles*

The Tesla miracle isn't just about making electric vehicles practical and desirable replacements for fossil fuel cars. Alongside the leading battery and motor technologies have been bold claims by Elon Musk that his cars will be the first you can buy that completely drive themselves too. A trial of Tesla's Full Self Driving ability has been making its way round a few US cities carrying selected beta-testing Tesla owners since October 2020. But recently

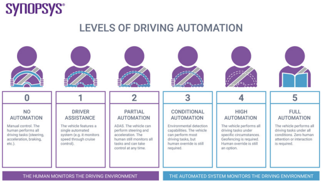

# How it works?

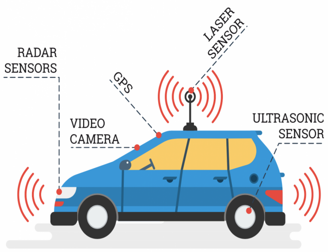
RADAR
SENSORS
GPS
LASER
SENSOR
VIDEO
CAMERA
ULTRASONIC
SENSOR


traffic sign
recognition
surround view ON

# How it works?

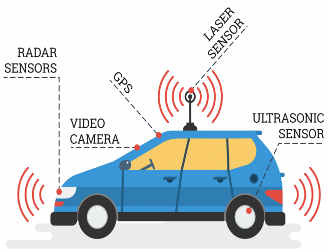
LASER SENSOR
RADAR SENSORS
GPS
VIDEO CAMERA
ULTRASONIC SENSOR


traffic sign recognition
surround view ON
STOP
STOP
50
Finish

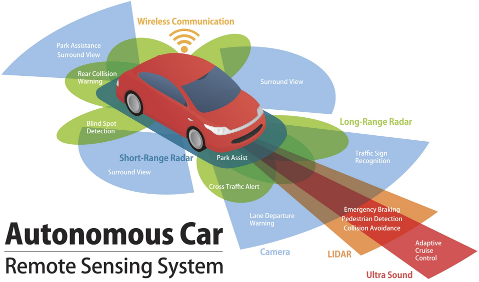
Wireless Communication
Park Assistance
Surround View
Rear Collision
Warning
Surround View
Long-Range Radar
Blind Spot
Detection
Traffic Sign
Recognition
Short-Range Radar
Park Assist
Surround View
Cross Traffic Alert
Emergency Braking
Pedestrian Detection
Collision Avoidance
Lane Departure
Warning
Adaptive
Cruise
Control
Autonomous Car
Remote Sensing System
Camera
LIDAR
Ultra Sound

# How robust?

# How robust?

**FOOLING THE AI**

Deep neural networks (DNNs) are brilliant at image recognition — but they can be easily hacked.

These stickers made an artificial-intelligence system read this stop sign as 'speed limit 45'.

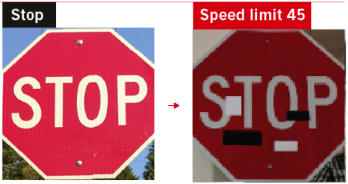

# Robustness to natural variations?

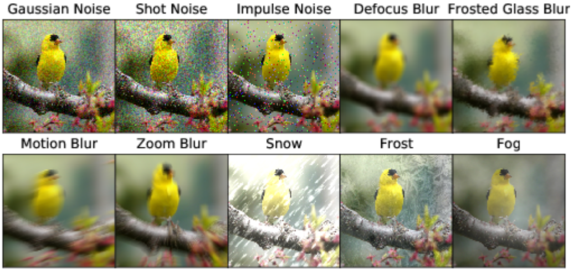


## ImageNet-C Leaderboard

ImageNet-C Robustness with a ResNet-50 Backbone trained on ImageNet-1K and evaluated on 224x224x3 images.

| Method | Reference | Standalone? | mCE | Clean Error |
|---|---|---|---|---|
| DeepAugment+AugMix | Hendrycks et al. | No | 53.6% | 24.2% |
| Assemble-ResNet50 | Lee et al. | No | 56.5% | 17.90% |
| ANT (3x3) | Rusak and Schott et al. | Yes | 63% | 23.9% |

# Outline

What's good about deep learning?

Why are self-driving cars not delivered?

Are we making the best use of deep learning?

What about understanding DL?

Why AI for healthcare now?

Research of the SUN group

solve difficult scientific and engineering problems

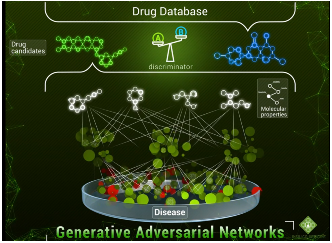


chemistry

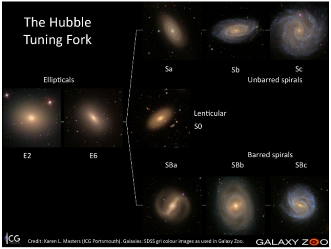


astronomy

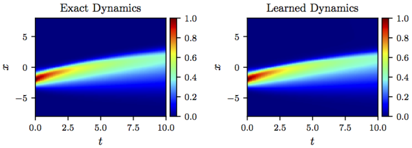


applied math

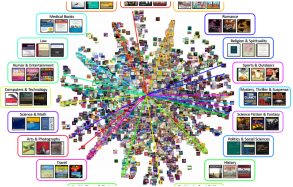


social science

## Is it really straight forward?

**Inverse problems**: given $f$ and $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$

# Is it really straight forward?

**Inverse problems**: given $f$ and $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$

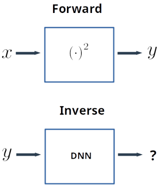

# Is it really straight forward?

**Inverse problems**: given $f$ and $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$

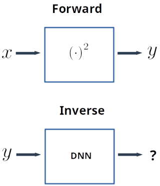


Original Dataset

| | |
|---|---|
| 4 | -2 |
| 9 | 3 |
| 16 | -4 |
| 100 | 10 |
| 4 | 2 |
| 64 | -8 |
| 36 | 6 |

# Is it really straight forward?

**Inverse problems**: given $f$ and $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$

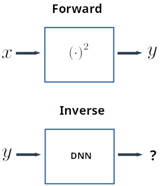


Original Dataset

| | |
|---|---|
| 4 | -2 |
| 9 | 3 |
| 16 | -4 |
| 100 | 10 |
| 4 | 2 |
| 64 | -8 |
| 36 | 6 |

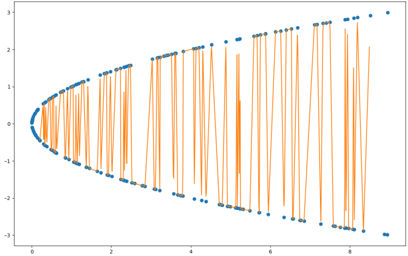

# Is it really straight forward?

**Inverse problems**: given $f$ and $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$

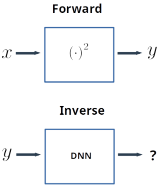


Original Dataset

| | |
|---|---|
| 4 | -2 |
| 9 | 3 |
| 16 | -4 |
| 100 | 10 |
| 4 | 2 |
| 64 | -8 |
| 36 | 6 |

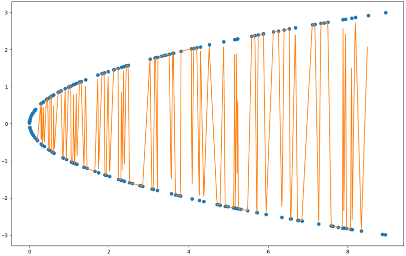


symmetries/ill-posedness in practical problems

# Lots of suboptimal practical usage...

e.g., **phase retrieval**: given $\boldsymbol{Y} = |\mathcal{F}(\boldsymbol{X})|^2$, recover $\boldsymbol{X}$

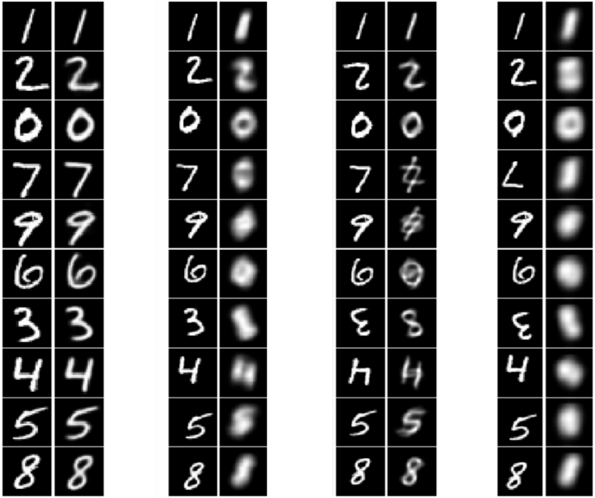

(a) No Symmetry (b) Shift symmetry (c) Flipping symmetry (d) Shift and Flipping symmetries

# Outline

# Theoretical research is high-risk high-reward

*I heard reiteration of the following claim: Complex theories do not work; simple algorithms do.*
*I would like to demonstrate that in the area of science a good old principle is valid:* **Nothing is more practical than a good theory**.

— *Vladimir N Vapnik, who invented support vector machines and statistical learning theory*

# Insights from randomness?

**(Fourier) phase retrieval**:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given $|\mathcal{F}\boldsymbol{x}|^2$, recover $\boldsymbol{x}$.

# Insights from randomness?

**(Fourier) phase retrieval**:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given $|\mathcal{F}\boldsymbol{x}|^2$, recover $\boldsymbol{x}$.

**Generalized phase retrieval**:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given $|\mathcal{A}\boldsymbol{x}|^2$ where $\mathcal{A}$ contains randomness, recover $\boldsymbol{x}$.

# Insights from randomness?

**(Fourier) phase retrieval**:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given $|\mathcal{F}\boldsymbol{x}|^2$, recover $\boldsymbol{x}$.

**Generalized phase retrieval**:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given $|\mathcal{A}\boldsymbol{x}|^2$ where $\mathcal{A}$ contains randomness, recover $\boldsymbol{x}$.

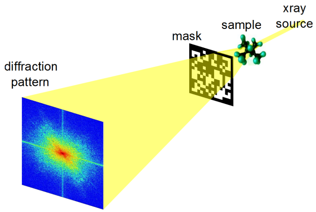


coded-diffraction

CDI [Candès et al., 2015]

# Insights from the Gaussian case?

$\boldsymbol{y} = |\boldsymbol{a}_i^* \boldsymbol{x}|$ for $i = 1, \ldots, m$ where $\boldsymbol{a}_i$'s complex Gaussian vectors

## Insights from the Gaussian case?

$\boldsymbol{y} = |\boldsymbol{a}_i^* \boldsymbol{x}|$ for $i = 1, \ldots, m$ where $\boldsymbol{a}_i$'s complex Gaussian vectors

- many beautiful mathematical results [Chi et al., 2018, Fannjiang and Strohmer, 2020]

# Insights from the Gaussian case?

$\boldsymbol{y} = |\boldsymbol{a}_i^* \boldsymbol{x}|$ for $i = 1, \ldots, m$ where $\boldsymbol{a}_i$'s complex Gaussian vectors

– many beautiful mathematical results [Chi et al., 2018, Fannjiang and Strohmer, 2020]

**Example 1**: a beautiful **init + local descent** result

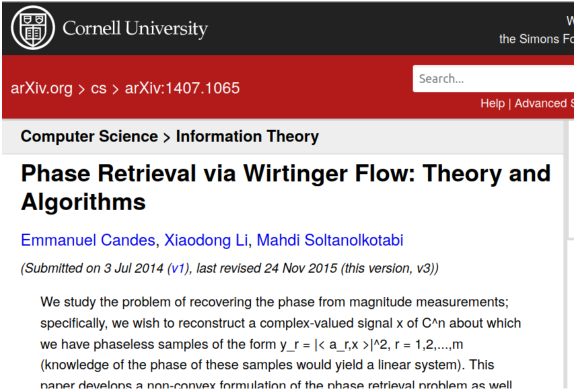
Cornell University

the Simons Fo

arXiv.org > cs > arXiv:1407.1065

Search...

Help | Advanced

**Computer Science > Information Theory**

**Phase Retrieval via Wirtinger Flow: Theory and Algorithms**

Emmanuel Candes, Xiaodong Li, Mahdi Soltanolkotabi

*(Submitted on 3 Jul 2014 (v1), last revised 24 Nov 2015 (this version, v3))*

We study the problem of recovering the phase from magnitude measurements; specifically, we wish to reconstruct a complex-valued signal x of C^n about which we have phaseless samples of the form y_r = |< a_r,x >|^2, r = 1,2,...,m (knowledge of the phase of these samples would yield a linear system). This paper develops a non-convex formulation of the phase retrieval problem as well

# Insights from the Gaussian case?

$\boldsymbol{y} = |\boldsymbol{a}_i^* \boldsymbol{x}|$ for $i = 1, \ldots, m$ where $\boldsymbol{a}_i$'s complex Gaussian vectors

– many beautiful mathematical results [Chi et al., 2018, Fannjiang and Strohmer, 2020]

**Example 2**: my own results

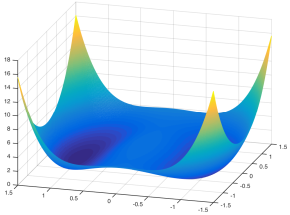

$$\boxed{\min_{\boldsymbol{z} \in \mathbb{C}^n} f(\boldsymbol{z}) \doteq \frac{1}{2m} \sum_{k=1}^{m} (y_k^2 - |\boldsymbol{a}_k^* \boldsymbol{z}|^2)^2.}$$

**Theorem ([Sun et al., 2016])**

*When $\boldsymbol{a}_k$'s generic and $m$ large, with high probability*

**all local minimizers are global, all saddles are nice**.

# I was happy until ...

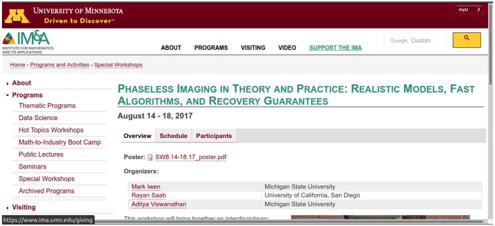

# I was happy until ...

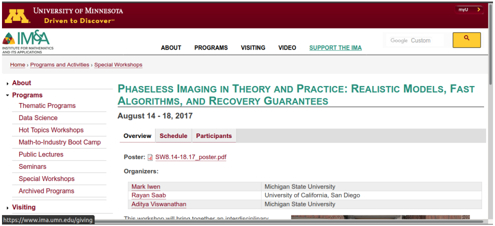


# Take-home messages

**James R Fienup**

(U. Rochester)

*I find it interesting people have tried to analyze Gaussian phase retrieval. —Fineup*

# Take-home messages

**James R Fienup**

(U. Rochester)

*I find it interesting people have tried to analyze Gaussian phase retrieval. —Fineup*

Beautiful mathematical results gathered so far [Chi et al., 2018, Fannjiang and Strohmer, 2020]

# Take-home messages

**James R Fienup**

(U. Rochester)

*I find it interesting people have tried to analyze Gaussian phase retrieval. —Fineup*

Beautiful mathematical results gathered so far [Chi et al., 2018, Fannjiang and Strohmer, 2020]

But we made little progress in solving Fourier PR

# Theories for DL?

PNAS Proceedings of the National Academy of Sciences of the United States of America

Keyword, Author, or D

Home | Articles | Front Matter | News | Podcasts | Authors

COLLOQUIUM ON THE SCIENCE OF DEEP LEARNING

Check for updates

## Theoretical issues in deep networks


Tomaso Poggio, Andrzej Banburski, and Qianli Liao

+ See all authors and affiliations



PNAS December 1, 2020 117 (48) 30039-30045; first published June 9, 2020; https://doi.org/10.1073/pnas.1907369117

Edited by David L. Donoho, Stanford University, Stanford, CA, and approved May 1, 2020 (received for review June 3, 2019)


Article | Figures & SI | Info & Metrics | PDF


### Abstract

While deep learning is successful in a number of applications, it is not yet well

# Theories for DL?

CSCI8980: Topics in modern machine learning (Fall 2021)

—put classic statistical learning theory in the context of modern deep learning, and put deep learning in the context of classic statistical learning theory

- Approximation theory for DL
- Optimization & Generalization
  - Classic theory: uniform convergence, VC-dim, Radamachar complexity, PAC-Bayesian bound, margin-based (⅓)
  - DL: implicit regularization, objection to implicit regularization, double descent
    - Generalization bounds for deep learning https://arxiv.org/abs/2012.04115
    - Understanding Deep Learning (Still) Requires Rethinking Generalization https://dl.acm.org/doi/pdf/10.1145/3446776
  - Generative prior: DIP and variants
  - Neural tangent kernels, lazy training
  - Robustness, interpretability, explainability, fairness, privacy, causality
    - Towards Causal Representation Learning https://arxiv.org/abs/2102.11107
  - Learning with imbalance and label noise
    - Early-Learning Regularization Prevents Memorization of Noisy Labels https://arxiv.org/abs/2007.00151
  - Learning with symmetries --- input (invariance & equivariance) & output
  - Landscape analysis (Batch normalization)
    - Why Flatness Correlates With Generalization For Deep Neural Networks https://arxiv.org/abs/2103.06219
  - Transfer learning & Domain adaptation
  - Self-supervised learning (contrastive learning)
  - Generative models (Normalization flow, GANs & VAE)
  - Use DL for solving hard problems (Maxcut, combinatorial problems, FPR)
  - 2nd order methods for DL (classification, inverse problems, etc)

- Scattering transform
- Randomized numerical linear algebra & concentration of measure: dimension reduction, sketched-based optimization

## Outline

# We're running short of doctors!

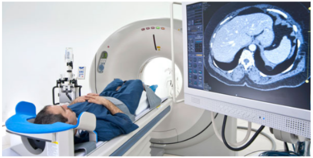

**Projected Physician Shortages by 2033**

| Medical Areas | Shortage Range |
|---|---|
| Primary care | Between 21,400 and 55,200 physicians |
| Nonprimary care specialties | Between 33,700 and 86,700 physicians |
| -- Surgical specialties | Between 17,100 and 28,700 physicians |
| -- Medical specialties | Between 9,300 and 17,800 physicians |
| -- Other specialties (i.e., pathology, radiology, psychiatry) | Between 17,100 and 41,900 physician |

Source: Assoc. American Medical Colleges

# Perils and promise

Perils

- **Small** datasets (sometimes)
- **Unbalanced** datasets (almost always)
- **Noisy** datasets (almost always)

# Perils and promise

Perils

- **Small** datasets (sometimes)
- **Unbalanced** datasets (almost always)
- **Noisy** datasets (almost always)

Promise

- Confined domain

- Robustness — noise less wild

# What we're up to

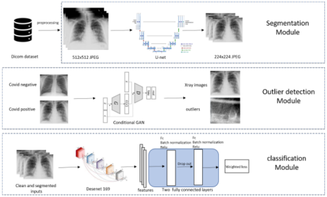


COVID-19 Diagnosis and Prognosis (deployed in 12 M Health Fairview H's)

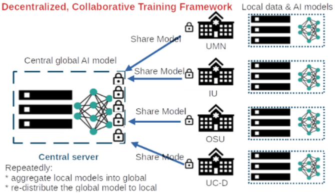


Collaborative/federated learning for medical imaging

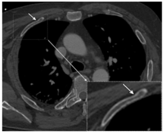

Fracture detection in critical/trauma care

# Outline

What's good about deep learning?

Why are self-driving cars not delivered?

Are we making the best use of deep learning?

What about understanding DL?

Why AI for healthcare now?

Research of the SUN group

{machine learning, data sciences, optimization,
computer vision, image/signal processing, imaging, healthcare...}

- **Representation Learning**: learn **efficient** representation for data
- **Computation**: compute with, often optimize with, massive amounts of data
- **Theory insights**: whenever possible/necessary

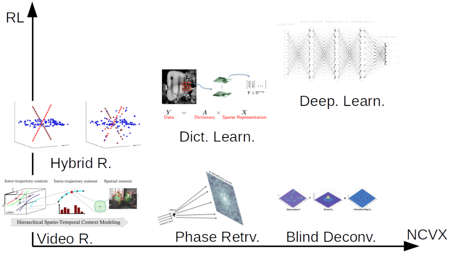

# SUN group at UMN

- foundations of machine/deep learning & computer vision
  * robustness in recognition
  * novel applications and limitations
  * fast computation and theoretical insights

# SUN group at UMN

- foundations of machine/deep learning & computer vision
  * robustness in recognition
  * novel applications and limitations
  * fast computation and theoretical insights
- application of CV and DL in healthcare
  * diagnosis and analysis of gliobastoma via brain MRI
  * fracture/COVID 19 detection from chest X-rays/CT
  * federated learning

CSCI8980: Topics in modern machine learning (Fall 2021)

—put classic statistical learning theory in the context of modern deep learning, and put deep learning in the context of classic statistical learning theory

- Approximation theory for DL
- Optimization & Generalization
    - Classic theory: uniform convergence, VC-dim, Radamachar complexity, PAC-Bayesian bound, margin-based (⅓)
    - DL: implicit regularization, objection to implicit regularization, double descent
        - Generalization bounds for deep learning https://arxiv.org/abs/2012.04115
        - Understanding Deep Learning (Still) Requires Rethinking Generalization https://dl.acm.org/doi/pdf/10.1145/3446776
    - Generative prior: DIP and variants
    - Neural tangent kernels, lazy training
    - Robustness, interpretability, explainability, fairness, privacy, causality
        - Towards Causal Representation Learning https://arxiv.org/abs/2102.11107
    - Learning with imbalance and label noise
        - Early-Learning Regularization Prevents Memorization of Noisy Labels https://arxiv.org/abs/2007.00151
    - Learning with symmetries --- input (invariance & equivariance) & output
    - Landscape analysis (Batch normalization)
        - Why Flatness Correlates With Generalization For Deep Neural Networks https://arxiv.org/abs/2103.06219
    - Transfer learning & Domain adaptation
    - Self-supervised learning (contrastive learning)
    - Generative models (Normalization flow, GANs & VAE)
    - Use DL for solving hard problems (Maxcut, combinatorial problems, FPR)
    - 2nd order methods for DL (classification, inverse problems, etc)

- Scattering transform
- Randomized numerical linear algebra & concentration of measure: dimension reduction, sketched-based optimization
- Graph neural networks/NN on non-Euclidean spaces

# References i


[Candès et al., 2015] Candès, E. J., Li, X., and Soltanolkotabi, M. (2015). **Phase retrieval from coded diffraction patterns.** *Applied and Computational Harmonic Analysis*, 39(2):277–299.

[Chi et al., 2018] Chi, Y., Lu, Y. M., and Chen, Y. (2018). **Nonconvex optimization meets low-rank matrix factorization: An overview.** *arXiv:1809.09573*.

[Fannjiang and Strohmer, 2020] Fannjiang, A. and Strohmer, T. (2020). **The numerics of phase retrieval.** *arXiv:2004.05788*.

[Sun et al., 2016] Sun, J., Qu, Q., and Wright, J. (2016). **A geometric analysis of phase retreival.** *arXiv preprint arXiv:1602.06664*.