Deep Learning for Robust Recognition, Inverse Problems, and Healthcare

Ju Sun (CS&E and Neurosurgery)

University of Minnesota
Driven to Discover℠

Thanks to my group

Thanks to my collaborators

Rajesh Rajamani, PhD
(Mech E.)

Robustness

Felix Hofmann
(Eng, Oxford)

Inverse Problems

Chris Tignamelli, MD
(Surgery & IHI)

Gene Melton-Meaux PhD, MD
(Surgery&IHI&Fairview)

Healthcare

Clark Chen
PhD, MD
(Neurosurgery)

Thanks to funding/support agencies

Why I moved to DL?

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

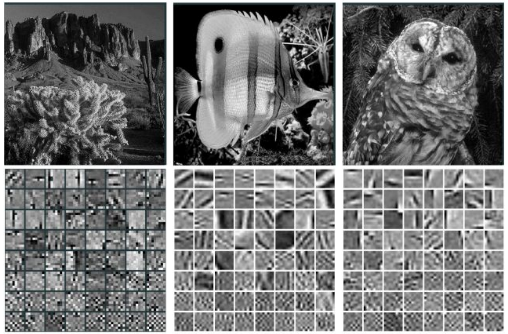

denoising

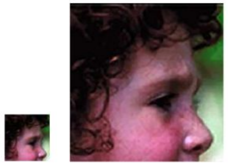

super resol.

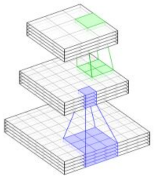

recognition

Dictionary Learning

$$\min \quad f(\boldsymbol{q}) \doteq \frac{1}{m} \|\boldsymbol{q}^* \boldsymbol{Y}\|_1 \quad \text{s.t.} \ \|\boldsymbol{q}\|_2^2 = 1. \qquad \boldsymbol{Y} \in \mathbb{R}^{n \times m}$$

Low-dim. $(n = 3)$ landscape when the target $\boldsymbol{Q}_0 = \boldsymbol{I}$ and $m \to \infty$

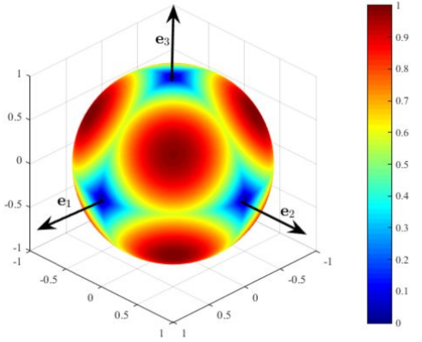

– global mins

– saddles

Why it works?

$\{$machine learning, data sciences, optimization, computer vision, image/signal processing, imaging, $\ldots\}$

(Fourier) phase retrieval:

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

Generalized phase retrieval:

For a complex signal $\boldsymbol{x} \in \mathbb{C}^n$, given measurements of the form $\lvert \boldsymbol{a}_k^* \boldsymbol{x} \rvert$ for $k = 1, \ldots, m$, recover $\boldsymbol{x}$.

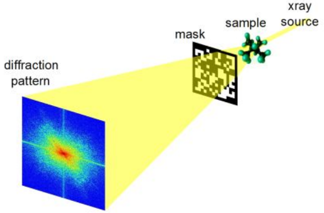

Applications: X-ray crystallography, diffraction imaging, optics, astronomical imaging, and microscopy

Phase Retrieval

Given $y_k = \lvert \boldsymbol{a}_k^* \boldsymbol{x} \rvert$ for $k = 1, \ldots, m$, recover $\boldsymbol{x}$ (**up to a global phase**).

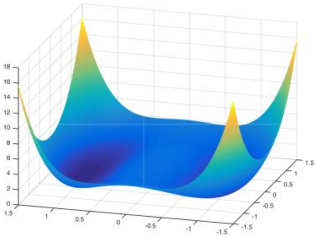

$$\min_{\boldsymbol{z} \in \mathbb{C}^n} f(\boldsymbol{z}) \doteq \frac{1}{2m} \sum_{k=1}^{m} (y_k^2 - \lvert \boldsymbol{a}_k^* \boldsymbol{z} \rvert^2)^2.$$

Theorem (Informal, [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.

Why it works?

Benign NCVX problems in practice!

All local mins are global, all saddles are strict

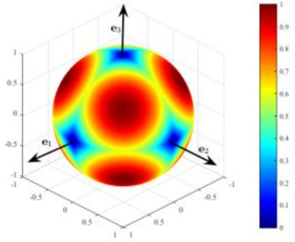

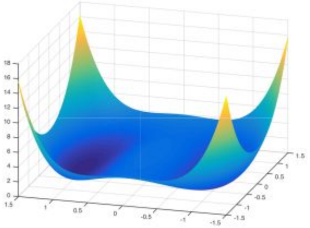

Eigenvalue problems (folklore!)

Sparsifying dictionary learning [Sun et al., 2015]

Generalized phase retrieval [Sun et al., 2016]

Orthogonal tensor decomposition [Ge et al., 2015]

Low-rank matrix recovery and completion

[Ge et al., 2016, Bhojanapalli et al., 2016]

Phase synchronization [Boumal, 2016]

Community detection [Bandeira et al., 2016]

Deep/shallow networks [Kawaguchi, 2016,

Lu and Kawaguchi, 2017, Soltanolkotabi et al., 2017]

Sparse blind deconvolution [Zhang et al., 2017]

Algorithms: virtually everything reasonable works!

[Conn et al., 2000, Nesterov and Polyak, 2006, Goldfarb, 1980, Jin et al., 2017]

RL X NCVX

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

– **Representation Learning**: learn **efficient** representation for data

– **Nonconvex Optimization**: when/how NCVX becomes **tractable**

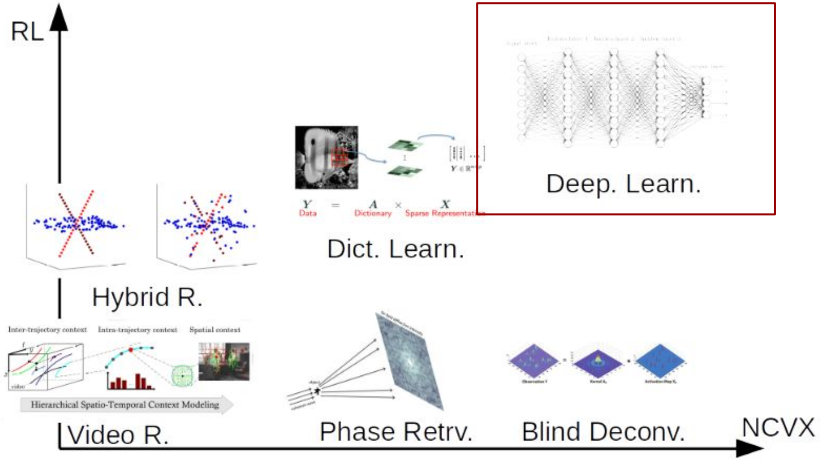

RL X NCVX → DL

Ju Sun   Welcome   Blog   Publications   Talks   Pe

Teaching

- CSCI2033: Elementary Computational Linear Algebra (Spring 2022)
- CSCI8980: Topics in Modern Machine Learning (Fall 2021)
- CSCI5525: Machine Learning: Analysis and Methods (Spring 2021)
- CSCI5980/8980: Think Deep Learning (Fall 2020)
- CSCI5980: Think Deep Learning (Spring 2020)

Where to start?

Application-driven

- **Limitation of DL: Robustness**

- **Power of DL: Difficult inverse problems**

- **Niche area of DL: Medical imaging in Healthcare**

Robustness of DL

Adversarial Robustness (AR)

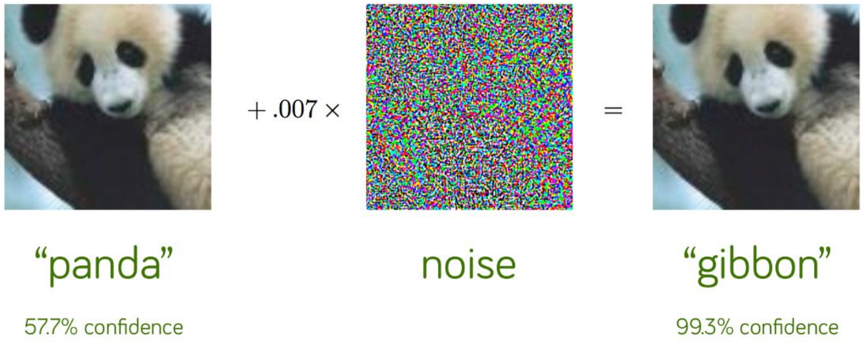

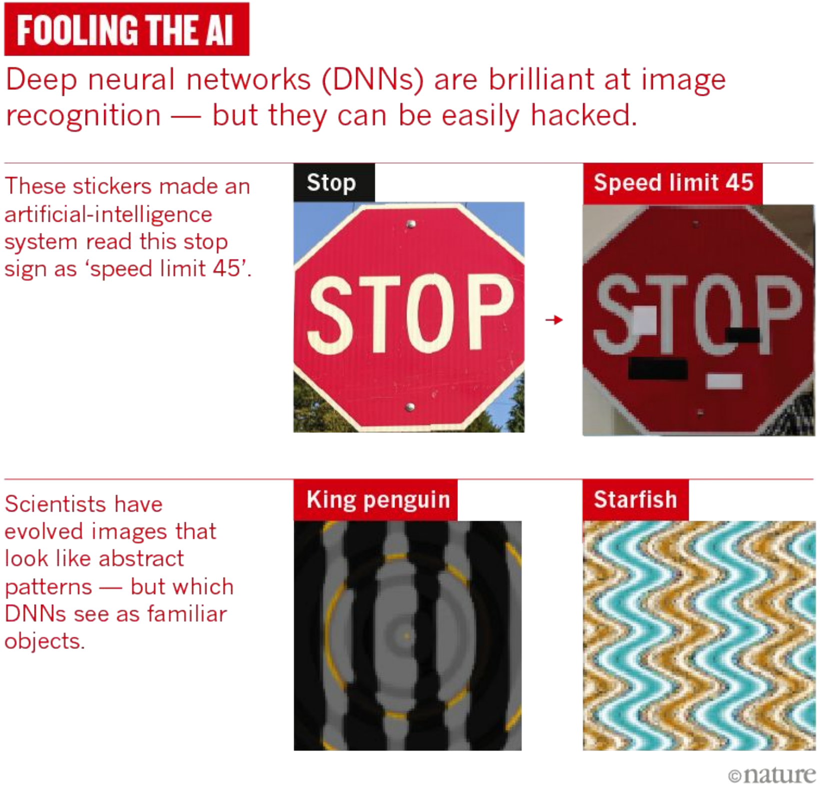

Attack:

$$\max_{\delta \leq \Delta} \mathcal{L}(\theta, x + \delta, y)$$

Defense:

$$\min_{\theta} \Big(\frac{1}{\mathcal{D}} \sum_{(x,y) \in \mathcal{D}} \max_{\delta \in \Delta(x)} \mathcal{L}(f(x+\delta), y) \Big)$$

Is AR what we care about?

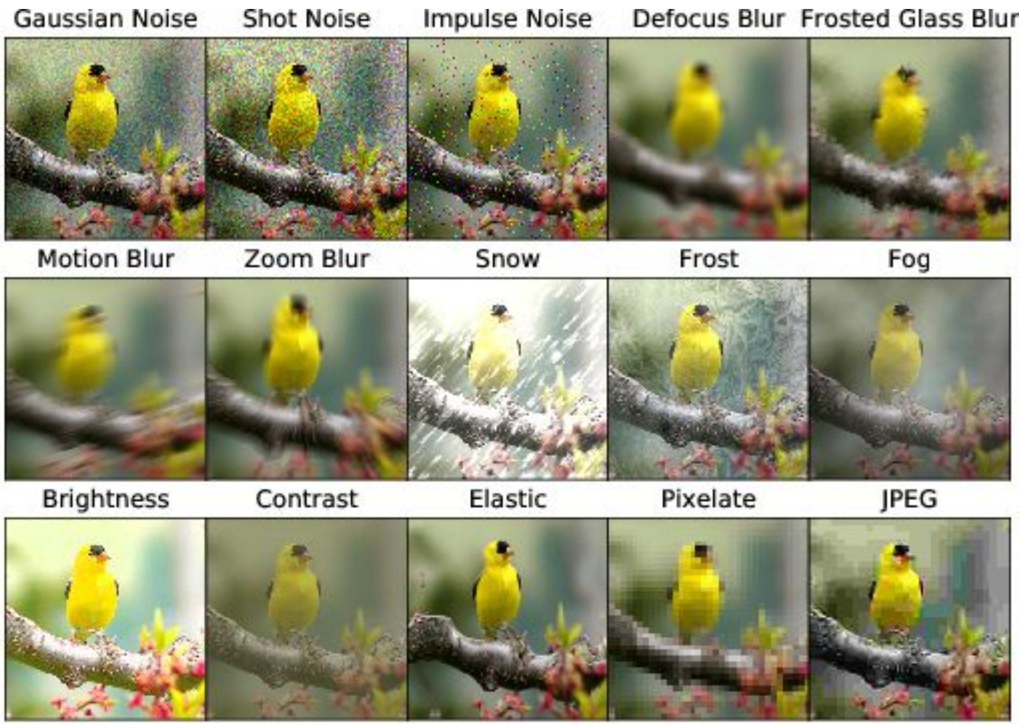

Imagenet-C

Natural Robustness?

- **Large perturbation**
- **Naturally occurring**

Observed on **classification, detection, segmentation, reconstruction, generation**, etc

Open questions

- **How to model? Modify?**

$$\min_{\theta}\left(\frac{1}{\mathcal{D}}\sum_{(x,y)\in\mathcal{D}}\max_{\delta\in\Delta(x)}\mathcal{L}(f(x+\delta),y)\right)$$

- **How to solve?**

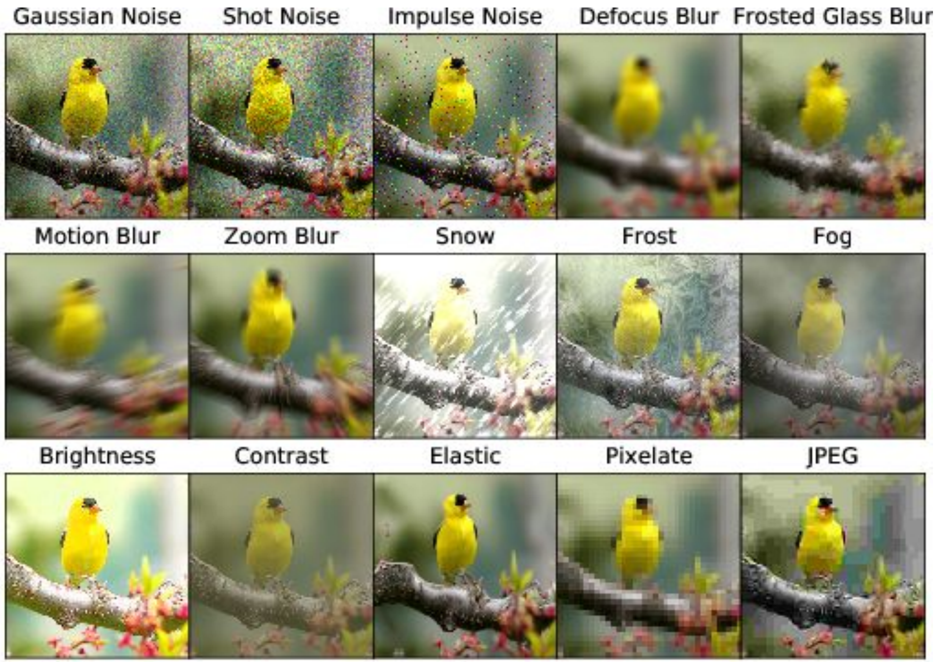

A possible approach: agnostic “denoising”

$$\min_{\theta} E(f_{\theta}(z); x_0) \, .$$

f is a DNN, z is frozen

Deep image prior

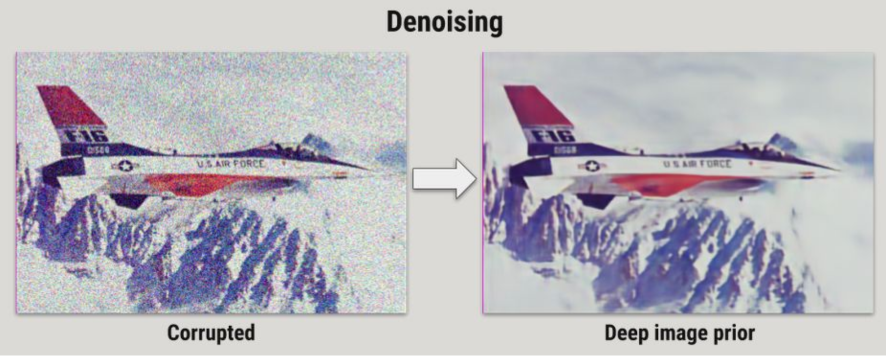

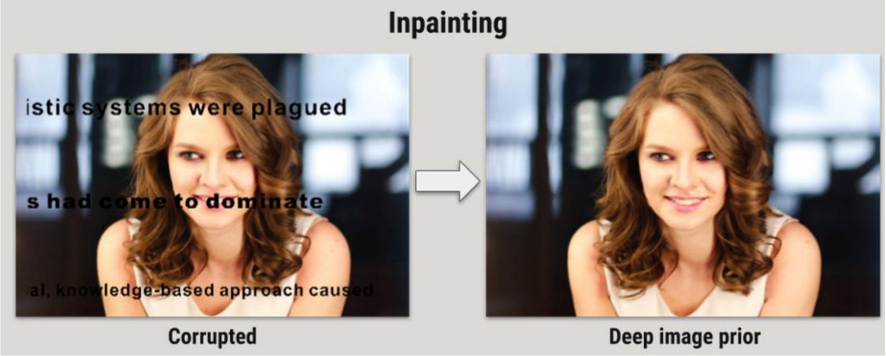

DIP: need for early stopping

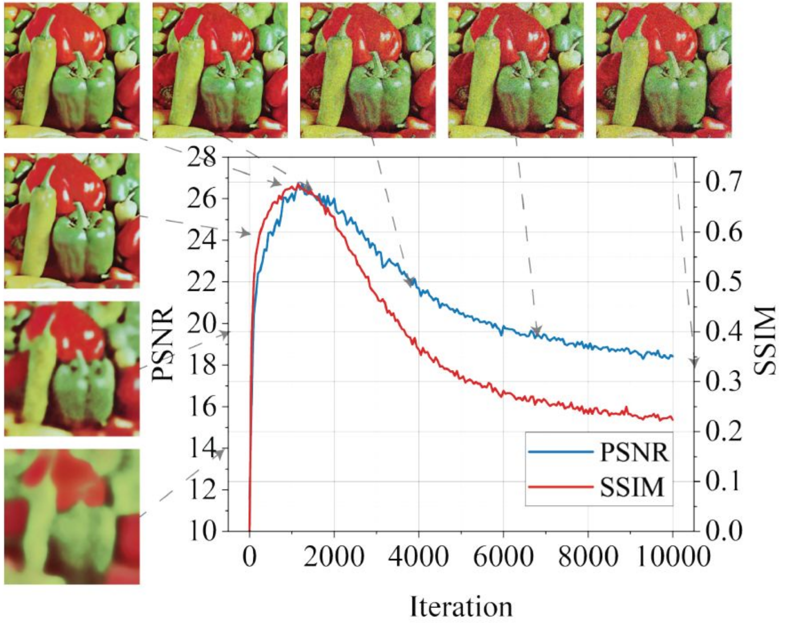

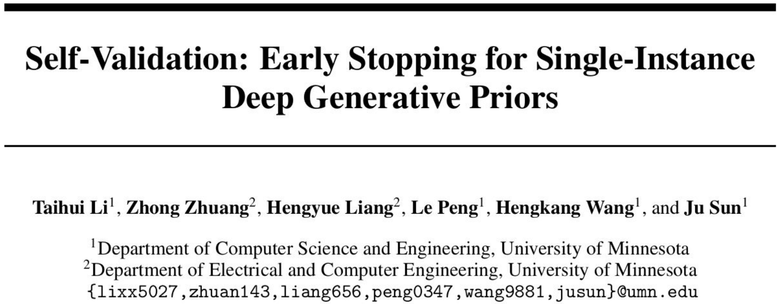

Self-Validation: Early Stopping for Single-Instance Deep Generative Priors

Taihui Li[1], **Zhong Zhuang**[2], **Hengyue Liang**[2], **Le Peng**[1], **Hengkang Wang**[1], and **Ju Sun**[1]

[1]Department of Computer Science and Engineering, University of Minnesota
[2]Department of Electrical and Computer Engineering, University of Minnesota
{lixx5027,zhuan143,liang656,peng0347,wang9881,jusun}@umn.edu

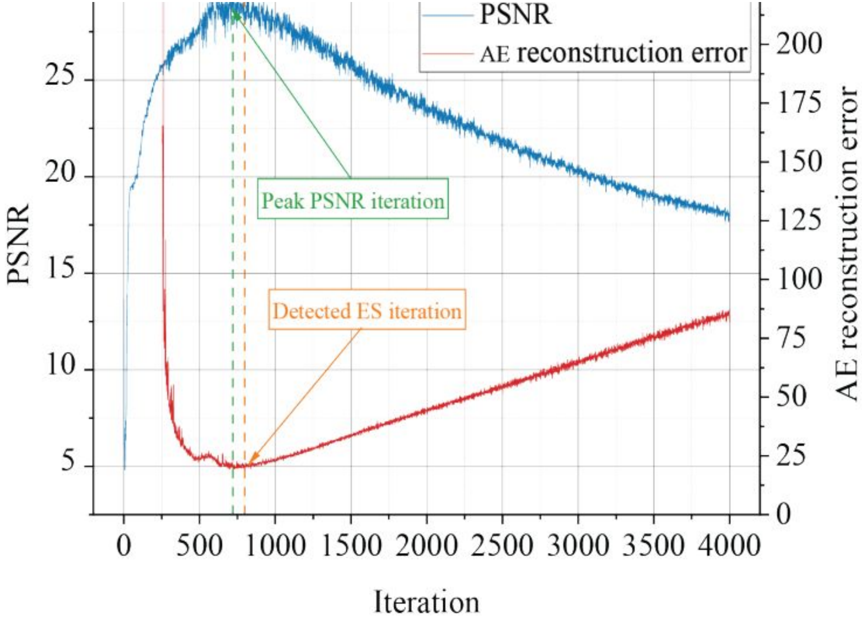

Other questions on DIP

$$\min_{\theta} E(f_{\theta}(z); x_0) \,.$$

- **Which E?**
- **Which DNN model?**
- **Speed?**
- **Initialization?**
- **Task-oriented: do we need perfect denoising?**

Inverse Problems

Phase retrieval (PR)

Phase retrieval (PR): Given $|\mathcal{F}(\boldsymbol{x})|^2$, recover $\boldsymbol{x}$

spectral factorization: find $X(z)$ so that $R(z) = \alpha X(z) X(z^{-1})$ and $X(z)$ has all roots inside the unit circle.

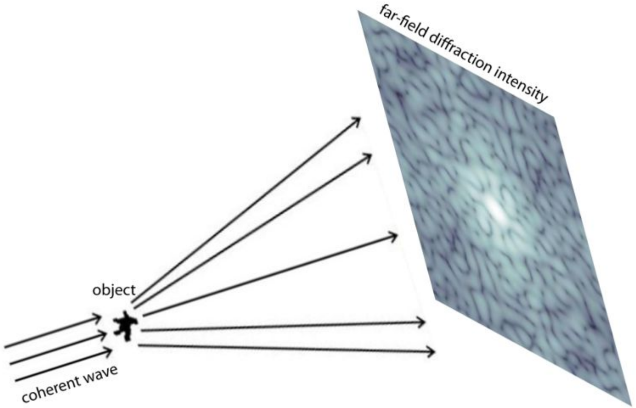

2D: Coherent diffraction imaging (CDI)

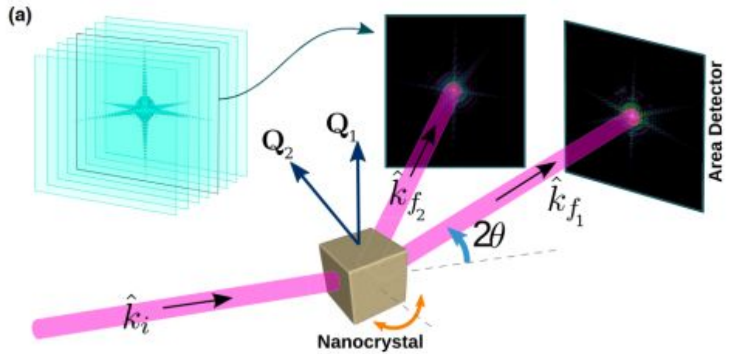

3D: (multi-reflection) Bragg CDI

Is phase retrieval (PR) solved?

(Fourier) phase retrieval:

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

Generalized phase retrieval:

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

James R Fienup

Fienup: I find it interesting people have tried to analyze Gaussian phase retrieval.

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

What's the gap?

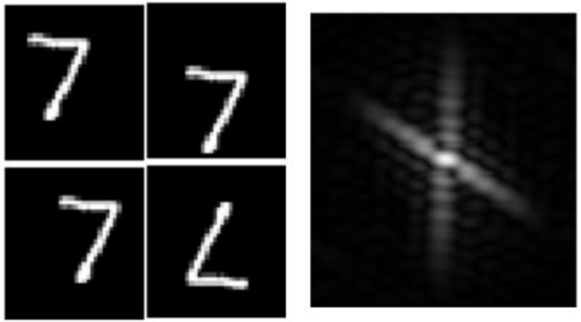

Symmetries in Fourier PR:

- translation
- 2D flipping
- global phase

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

GPR doesn't contain the translation and flipping symmetries!

Albert Einstein: Everything should be made as simple as possible, but **no simpler**.

DL for Inverse Problems

Given $\boldsymbol{y} = f(\boldsymbol{x})$, estimate $\boldsymbol{x}$ (f may be unknown)

- Traditional

$$\min_{\boldsymbol{x}} \ \ell(\boldsymbol{y}, f(\boldsymbol{x})) + \lambda \Omega(\boldsymbol{x})$$

- Modern
  * End-to-end: set up $\{\boldsymbol{x}_i, \boldsymbol{y}_i\}$ to learn f^{-1} directly
  * Hybrid: replace ℓ, Ω, or algorithmic components using **learned functions**, e.g., plug-and-play ADMM, unrolling ISTA
- "Modern" works **better** when "traditional" **already works**

Recent surveys: [McCann et al., 2017, Lucas et al., 2018, Arridge et al., 2019, Ongie et al., 2020]

DL for PR

- Hybrid: replace ℓ, Ω, or algorithmic components using **learned functions**, e.g., plug-and-play ADMM, unrolling ISTA

  "modern" works **better** when traditionally **already works**

  Attempts: [Metzler et al., 2018, Işıl et al., 2019], but HIO needed for initialization

- End-to-end: set up $\{\boldsymbol{x}_i, |\mathcal{F}\boldsymbol{x}|^2\}$ to learn f^{-1} directly

  Attempts:

  [Goy et al., 2018, Uelwer et al., 2019, Metzler et al., 2020]

  with positive initial results

How good are they?

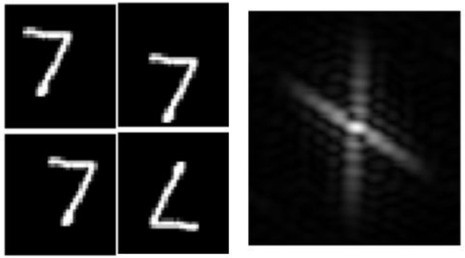

Symmetries in Fourier PR:

- shift
- 2D flipping
- global phase

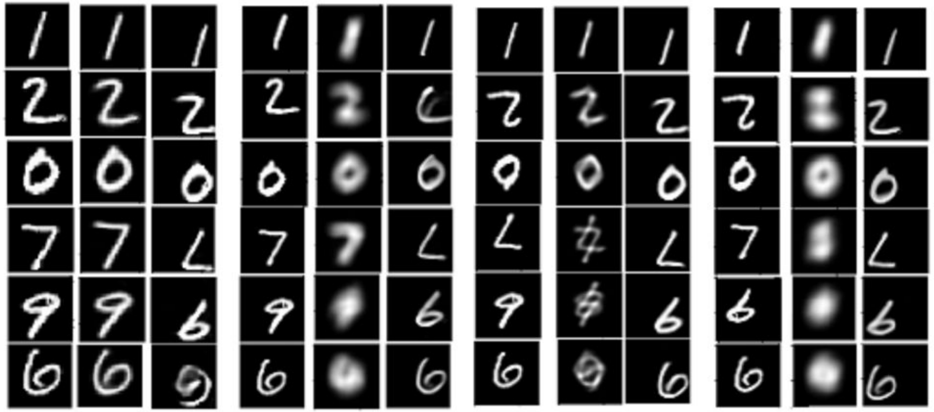

(a) No Symmetry　(b) Shift symmetry　(c) Flipping symmetry　(d) Shift and Flipping symmetries

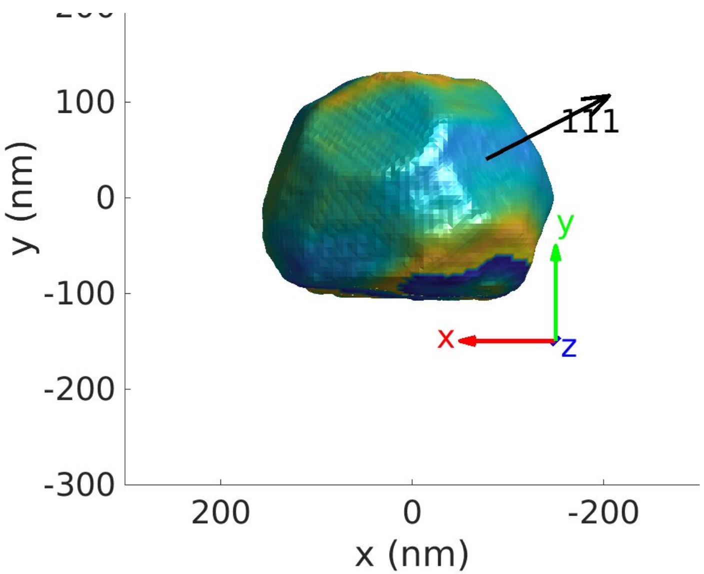

Fail miserably once simulating realistic datasets

Why they fail?

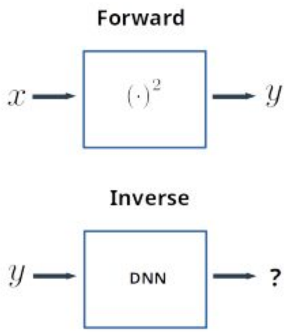

Original Dataset

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

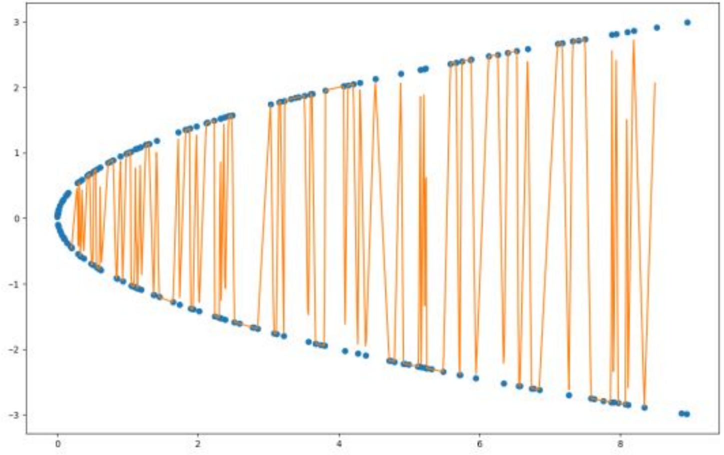

nearby inputs mapped to remote outputs **due to symmetries**

Approximating (highly) oscillatory functions

Other examples

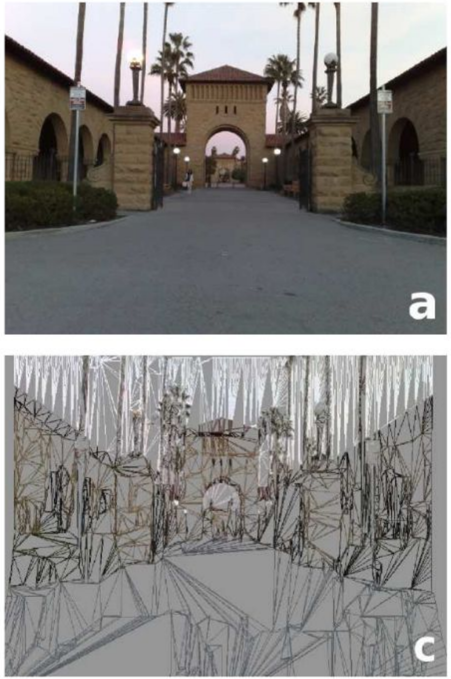

3D depth from 2D image

Deblurring

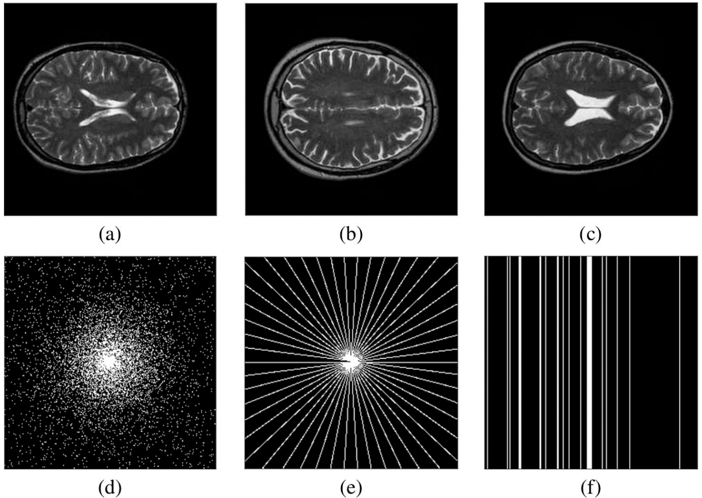

MRI Reconst.

Sol: symmetry breaking

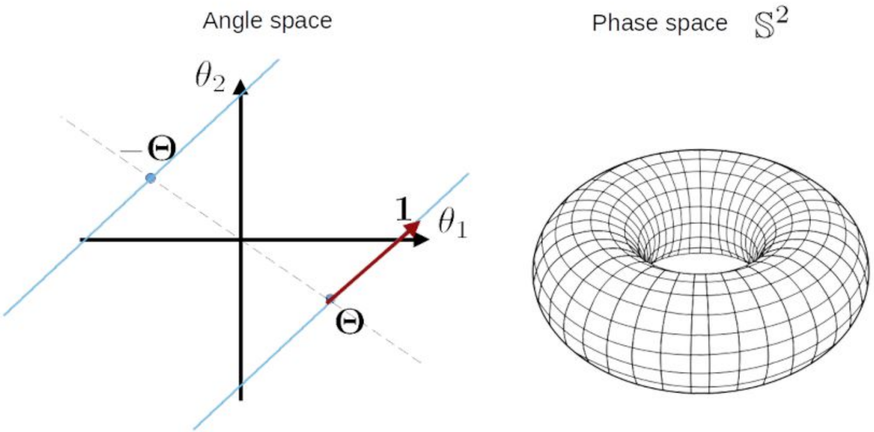

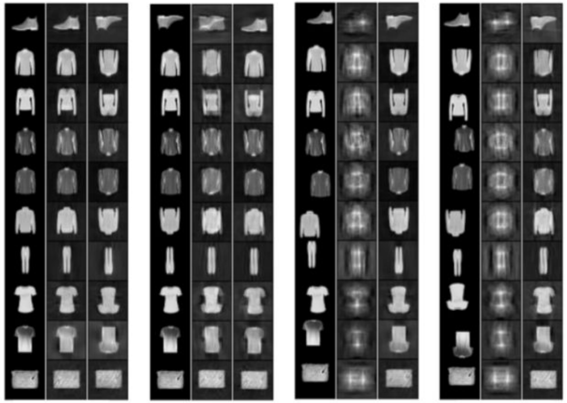

(a) No Symmetry (b) Flipping Symmetry (c) Shift Symmetry (d) Shift & Flipping Symmetry

Unlocking Inverse Problems Using Deep Learning: Breaking Symmetries in Phase Retrieval

Kshitij Tayal[1], Chieh-Hsin Lai[2], Raunak Manekar[1], Zhong Zhuang[3], Vipin Kumar[1], Ju Sun[1]

[1]Department of Computer Science and Engineering, University of Minnesota, Twin Cities, USA
[2]School of Mathematics, University of Minnesota, Twin Cities, USA
[3]Department of Electrical and Computer Engineering, University of Minnesota, Twin Cities, USA

Table 1: Test error using different symmetry schemes

	U-Net-B	U-Net-A (ours)
No Symmetry	0.103	0.103
Flipping Symmetry	0.168	0.162
Shift Symmetry	0.249	0.102
Shift & Flipping Symmetry	0.248	0.161

Table 2: MSE error

Method	MSE
ALM	0.299
U-Net-B	0.249
U-Net-A	0.160

Open problems

- Essential difficulty: use DL to approximate **one-to-many** mapping
  - When there is forward symmetry (this talk)
  - When the forward mapping under-determined (super-resolution, 3D structure from a single image)
  - or Both
- Not only learning difficulty, but also **robustness** [Antun et al., 2020, Gottschling et al., 2020]

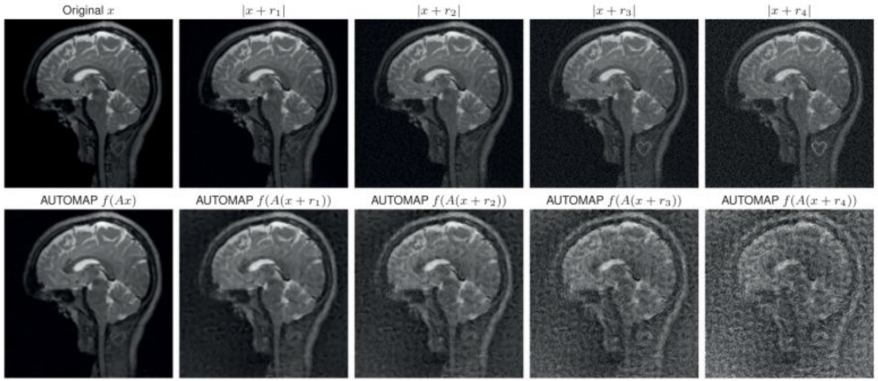

Healthcare

The core team (UMN Computer Vision in Healthcare Initiative)

Chris Tignamelli, MD
(Surgery & IHI)

Trauma/Critical Care

Ju Sun, PhD
(CS&E&IHI&Neurosurgery)
Computer Vision/AI

Gene Melton-Meaux
PhD, MD
(Surgery&IHI&Fairview)
NLP/Informatics

Tadashi Allen, MD
(Radiology)

Radiology

University of Minnesota
Driven to Discover®

Health
Fairview

U of M
U of M Physicians
Fairview Health Service

Why AI/CV for Medical Imaging (MI)?

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

Why AI/CV for Medical Imaging(MI)?

Perils

- **Small** datasets (often)
- **Unbalanced/biased** datasets (almost always)
- **(Label-)Noisy** datasets (almost always)

Promise – confined domains

What we do?

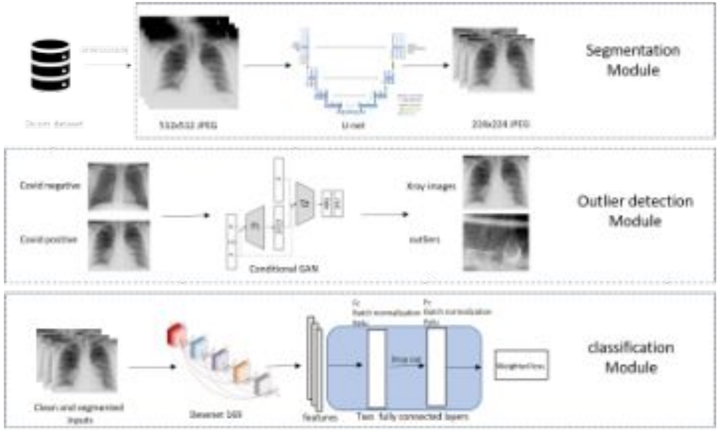

COVID-19 Diagnosis and Prognosis
(deployed in 12 M Health Fairview H's)

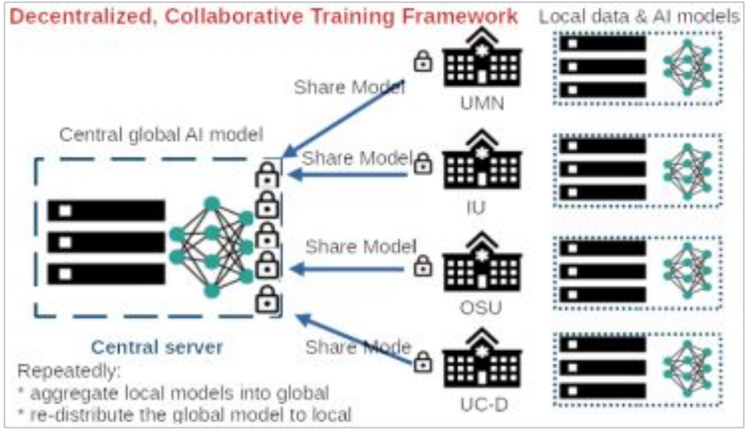

Collaborative/federated learning for medical imaging

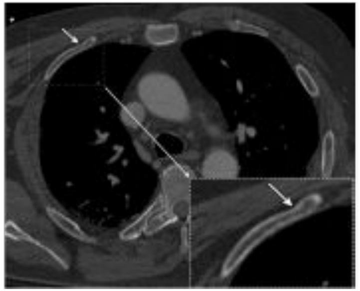

Fracture detection in critical/trauma care

Recap of the COVID-19 Project

Start (Mar 2020) –
Deployed in 12 hospitals of M Health Fairview and Epic App Orchard (Nov 2020)

The Diagnostic Model

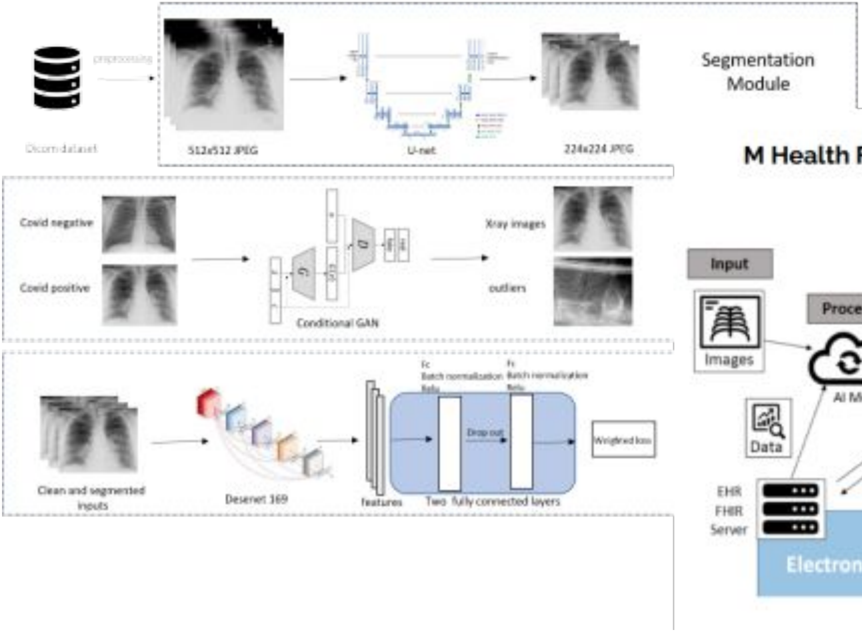

Deployment

M Health Fairview CDS: 12 hospitals

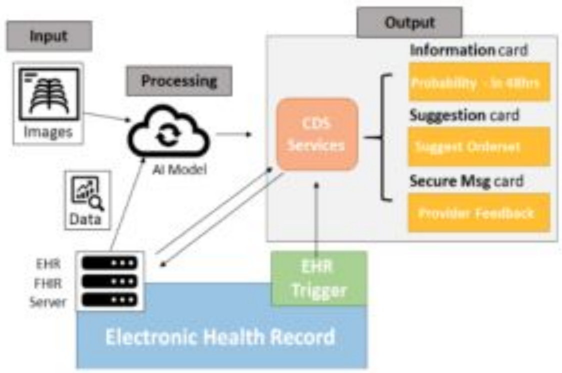

Epic App Orchard: 450+ customers

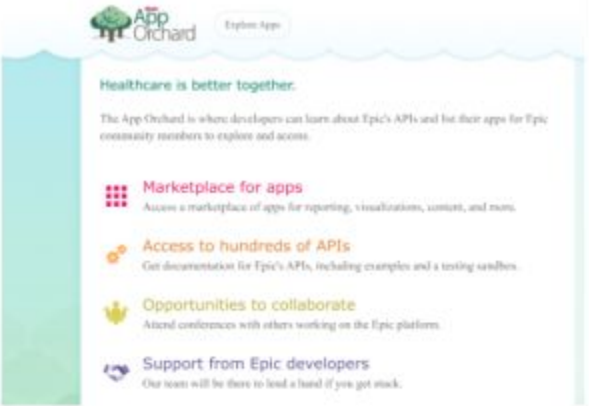

Academic products

Artificial Intelligence to Accelerate COVID-19 Identification from Chest X-rays

Ju Sun PhD[1], Taihui Li MS[1], Le Peng BS[1], Dyah Adila BS[1], Genevieve B. Melton MD PhD[2,3], Nicholas Ingraham MD[4], Daniel Boley PhD[1], Basil S. Karam MD[5], Tadashi Allen MD[6], Rachel Morris MD[5], Erich Kummerfeld PhD[2]*, Christopher Tignanelli MD[2,3,7]*

Under review in Radiology AI 2021

Rethink Transfer Learning in Medical Imaging

No Author Given

No Institute Given

Abstract. Transfer learning (TL) with deep convolutional neural networks (DCNNs) has proved successful in medical classification tasks. However, the practice is puzzling, as medical image classification typi-

Under review in AAAI, 2021

Now: addressing major challenges in CV/AI for MI

Perils

- **Small** datasets (often)
- **Unbalanced/biased** datasets (almost always)
- **(Label-)Noisy** datasets (almost always)

Getting more data solves most of these problems **but**…

- Privacy
- Security
- Profitability
- …

Federated learning (FL) in CV/AI for MI

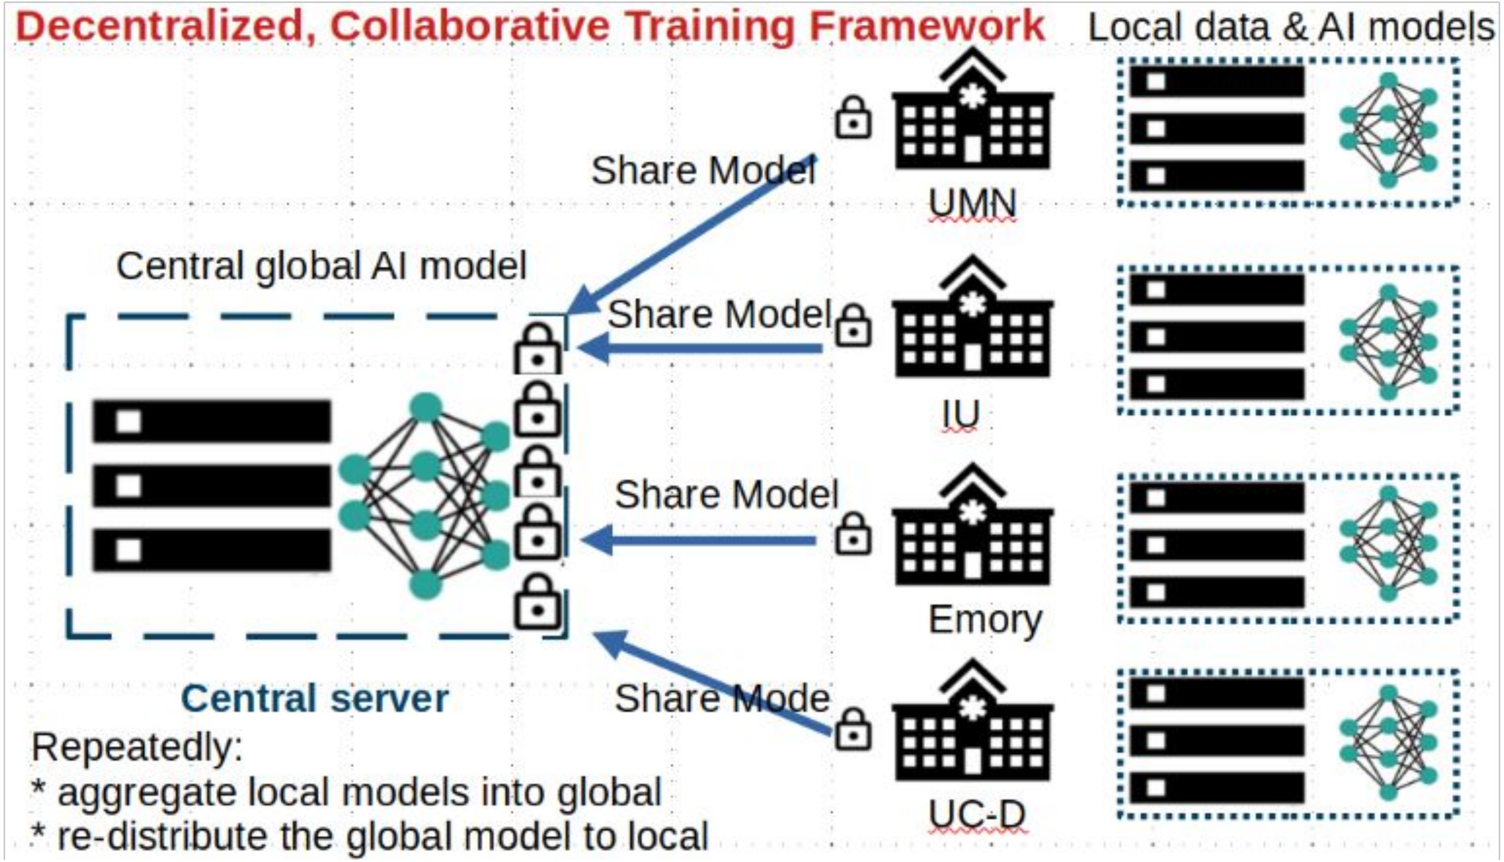

UMN-IU-Emory-UCD-**X** FL partnership

- ✔ Sep 2020: established
- ✔ Sep 2020 – Jan 2021: validation of UMN COVID-19 model
- ✔ Jan – Mar 2021: FL infrastructure set up
- ✔ Jan – June 2021: FL testdrive on the COVID-19 model
- ✔ Jan – June 2022: FL for fracture detection

Expect to expand to >=10 partners over next 2 years

Major FL questions

Practical strategies to

- Handle distribution shift (yes, batchnorm hurts…)
- Minimize privacy exposure
- Reduce security concerns (CISCO is in!)

Closing

Application-driven, toward DL theory

- **Limitation of DL: Robustness**
- **Power of DL: Difficult inverse problems**
- **Niche area of DL: Medical imaging in Healthcare**