## How can Al enhance

## measurement?

University of Chicago and Fermilab

## Eric Jonas

Assistant Professor, Department of Computer Science
Committee on Computational and Applied Mathematics
Physical Sciences Division, University Of Chicago


## I work on AI + Scientific Measurement

(This is the best place to be!)


## I work on AI + Scientific Measurement

(This is the best place to be!)


## I work on AI + Scientific Measurement

(This is the best place to be!)


## I work on AI + Scientific Measurement

(This is the best place to be!)


Wowza!

## Wowza!

## NeurIPS papers

Number of papers over years

nature

## Wowza!

## NeurlPS papers

Number of papers over years


Explore content $\checkmark$
nature > news > article

## NEWS 22 September 2022

## AlphaFold developers win US\$3million Breakthrough Prize

DeepMind's system for predicting the 3D structure of proteins is among five recipients of science's most lucrative awards.


Demis Hassabis (left) and John Jumper (right) from DeepMind developed AlphaFold, an Al that can predict the structure of proteins. Credit: Breakthrough Prize

Yikes!

## Yikes!

ARTIFICIAL INTELLIGENCE
Hundreds of Al tools have been built to catch covid. None of them helped.

Some have been used in hospitals, despite not being properly tested. But the
pandemic could help make medical Al better.

By Will Douglas Heaven
July 30, 2021

## Yikes!

ARTIFICIAL INTELLIGENCE
Hundreds of Al tools have been built to catch covid. None of them helped.

Some have been used in hospitals, despite not being properly tested. But the
pandemic could help make medical Al better.

By Will Douglas Heaven
July 30, 2021

## Google medical researchers humbled when Al screening tool falls short in real-life testing



What makes
Al for Experiments
special?


## We take our models seriously

Alexander, Kaitlin, and Stephen M. Easterbrook. "The software architecture of climate models: a graphical comparison of CMIP5 and EMICAR5 configurations." Geoscientific Model Development 8.4 (2015): 1221-1232.


## We take our models seriously

Alexander, Kaitlin, and Stephen M. Easterbrook. "The software architecture of climate models: a graphical comparison of CMIP5 and EMICAR5 configurations." Geoscientific Model Development 8.4 (2015): 1221-1232.

## Distribution shift is the point

## Need to think carefully about baselines

> Simple random search provides a competitive approach to reinforcement learning
Horia Mania Aurelia Guy Benjamin Recht

[^0]arXiv: I 803.07055v| March 20I8

An opinionated taxonomy
For AI + Experiments

## An opinionated taxonomy

For AI + Experiments


## Inverse Problems

How can Al help us understand the data that we get from existing systems

## An opinionated taxonomy

For AI + Experiments


Inverse Problems

How can Al help us understand the data that we get from existing systems


Computational Measurement

How can we design new measurement systems to be more interpretable by AI?

## An opinionated taxonomy

For AI + Experiments


Inverse Problems

How can Al help us understand the data that we get from existing systems


Computational Measurement

How can we design new measurement systems to be more interpretable by AI?


## Active Learning

How can Al guide experimentation and measurement?

## An opinionated taxonomy

For AI + Experiments


Inverse Problems

How can Al help us understand the data that we get from existing systems


Computational Measurement

How can we design new measurement systems to be more interpretable by AI?


## Active Learning

How can Al guide experimentation and measurement?

## Al for Inverse Problems

How can Al help us understand the data
that we get from existing systems

## Inverse Problems and Machine Learning

## Inverse Problems and Machine Learning

## (unobserved)

system

## Inverse Problems and Machine Learning

(unobserved)


## Inverse Problems and Machine Learning



## Inverse Problems and Machine Learning



## Inverse Problems and Machine Learning



## Inverse Problems and Machine Learning



## EXAMPLE PROBLEMS

Microscopy

## EXAMPLE PROBLEMS

Microscopy


## EXAMPLE PROBLEMS

Microscopy


## EXAMPLE PROBLEMS

Microscopy


Magnetic Resonance Imaging


## EXAMPLE PROBLEMS

Microscopy


Magnetic Resonance Imaging


## How much do we believe our model?

## How much do we believe our model?

Machine Learning

"What mode!?"

## How much do we believe our model?


"What model?"

## Statistics


capture some aspect of the system

## How much do we believe our model?



## Inverse Problems

 $U(P)=\frac{1}{4 \pi} \int_{S}\left[U \frac{\partial}{\partial n}\left(\frac{e^{i k s}}{s}\right)-\frac{e^{i k s}}{s} \frac{\partial U}{\partial n}\right] d S$,(Kirchhoff's diffraction formula)
"What model?"
capture some aspect of the system

Trust completely

# WHY ARE INVERSE PROBLEMS HARD? 

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem
$y=A x+\epsilon$

# WHY ARE INVERSE PROBLEMS HARD? 

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem
$y=A x+\epsilon$
A

# WHY ARE INVERSE PROBLEMS HARD? 

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem

$$
y=A x+\epsilon
$$

A
identity? easy

# WHY ARE INVERSE PROBLEMS HARD? 

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem

$$
y=A x+\epsilon
$$

A
identity? easy
not full rank? hard

## WHY ARE INVERSE PROBLEMS HARD?

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem

$$
y=A x+\epsilon
$$

$A$
identity? easy
not full rank? hard
poor condition number? hard

$$
\kappa(A)=\frac{\sigma_{\max }(A)}{\sigma_{\min }(A)}
$$

## WHY ARE INVERSE PROBLEMS HARD?

Inverse problems are hard for the same reasons that inverting a matrix is hard.

Linear inverse problem

$$
y=A x+\epsilon
$$

$A$
identity? easy
not full rank? hard
poor condition number? hard

$$
\kappa(A)=\frac{\sigma_{\max }(A)}{\sigma_{\min }(A)}
$$

How to solve? Use prior knowledge!

Learning Fast Approximations of Sparse Coding

Karol Gregor and Yann LeCun \{KGregor,yann\}@cs.nyu.edu Courant Institute, New York University, 715 Broadway, New York, NY 10003, USA

## Accurate Image Super-Resolution Using Very Deep Convolutional Networks

Jiwon Kim, Jung Kwon Lee and Kyoung Mu Lee
Department of ECE, ASRI, Seoul National University, Korea

## Image Super-Resolution Using Deep Convolutional Networks

Chao Dong, Chen Change Loy, Member, IEEE, Kaiming He, Member, IEEE, and Xiaoou Tang, Fellow, IEEE

Deep Convolutional Neural Network for Inverse Problems in Imaging
Kyong Hwan Jin, Michael T. McCann, Member, IEEE, Emmanuel Froustey, and Michael Unser, Fellow, IEEE

ONSAGER-CORRECTED DEEP LEARNING FOR SPARSE LINEAR INVERSE PROBLEMS
Mark Borgerding and Philip Schniter
Dept. of ECE, The Ohio State University, Columbus, OH 43202 Email: borgerding.7@osu.edu, schniter.1@osu.edu

Deeply-Recursive Convolutional Network for Image Super-Resolution

## Lensless computational imaging through deep learning

Ayan Sinha ${ }^{1 *}$, Justin Lee ${ }^{2}$, Shuai Lis ${ }^{1}$, and George Barbastathis ${ }^{1,3}$

LEARNING TO INVERT:
SIGNAL RECOVERY VIA DEEP CONVOLUTIONAL NETWORKS
Ali Mousavi and Richard G. Baraniuk

DEEP CONVOLUTIONAL FRAMELETS: A GENERAL DEEP LEARNING FOR INVERSE PROBLEMS*
Jong chul ye ${ }^{+\dagger}$ and yoseob han-
One Network to Solve Them All — Solving Linear Inverse Problems using Deep Projection Models

```
J. H. Rick Chang, Chun-Liang Li, Barnabás Póczos, B. V. K. Vijaya Kumar,
```

            and Aswin C. Sankaranarayanan
            Carnegie Mellon University, Pittsburgh, PA
    Robust Single Image Super-Resolution via Deep Networks With Sparse Prior
Ding Liu, Student Member, IEEE, Zhaowen Wang, Member, IEEE, Bihan Wen, Student Member, IEEE, Jianchao Yang, Member, IEEE, Wei Han, and Thomas S. Huang, Fellow, IEEE

## Amortised MAP Inference for Image Super-ReSolution

Casper Kaae Sønderby ${ }^{12 *}$, Jose Caballero ${ }^{1}$, Lucas Theis ${ }^{1}$, Wenzhe Shi ${ }^{1}$ \& Ferenc Huszár ${ }^{1}$ casperkaae@gmail.com, \{jcaballero,ltheis, wshi,fhuszar\}@twitter.com casperkaae@gma1 ${ }^{2}$ University of Copenhagen, Denmark

## \#fakenews



Pixel Recursive Super Resolution

Ryan Dahl * Mohammad Norouzi Jonathon Shlens
Google Brain
$\{$ rld,mnorouzi, shlens\}@google.com


Figure 1: Illustration of our probabilistic pixel recursive super resolution model trained end-to-end on a dataset of celebrity faces. The left column shows $8 \times 8$ low resolution inputs from the test set. The middle and last columns show $32 \times 32$ images as predicted by our model $v s$. the ground truth. Our model incorporates strong face priors to synthesize realistic hair and skin details.

## Is this invertable?

## Is this invertable?

## CAN ONE HEAR THE SHAPE OF A DRUIV?

MARK KAC, The Rockefeller University, New York
To George Eugene Uhlenbeck on the occasion of his sixty-fifth birthday
"La Physique ne nous donne pas seulement l'occasion de résoudre des problèmes . . . , elle nous fait presentir la solution." H. Poincaré.

Before I explain the title and introduce the theme of the lecture I should like to state that my presentation will be more in the nature of a leisurely excursion than of an organized tour. It will not be my purpose to reach a specified destination at a scheduled time. Rather I should like to allow myself on many occasions the luxury of stopping and looking around. So much effort is being spent on streamlining mathematics and in rendering it more efficient, that a solitary transgression against the trend could perhaps be forgiven.

American Mathematical Monthly I966

## Is this invertable?

## CAN ONE HEAR THE SHAPE OF A DRUM?

MARK KAC, The Rockefeller University, New York
To George Eugene Uhlenbeck on the occasion of his sixty-fifth birthday
"La Physique ne nous donne pas seulement l'occasion de résoudre des problèmes . . . , elle nous fait presentir la solution." H. Poincaré.

Before I explain the title and introduce the theme of the lecture I should like to state that my presentation will be more in the nature of a leisurely excursion than of an organized tour. It will not be my purpose to reach a specified destination at a scheduled time. Rather I should like to allow myself on many occasions the luxury of stopping and looking around. So much effort is being spent on streamlining mathematics and in rendering it more efficient, that a solitary transgression against the trend could perhaps be forgiven.

## American Mathematical Monthly I966

In 1966, in a celebrated paper (Kac, 1966), Mark Kac formulated the famous question "Can one hear the shape of a drum?". This provocative question is of course to be understood mathematically as follows: Is it possible to find two (or more) non-isometric Euclidean simply connected domains for which the sets $\left\{E_{n} \| n \in \mathbb{N}\right\}$ of solutions of (1) with $\Psi_{\mid \text {Boundary }}=0$ are identical? More broadly, the question raises the issue of the inverse problem of retrieving information about a drum from knowledge of its spectral properties. As the spectroscopist A. Schuster put it in an 1882 report to the British Association for the Advancement of Science: "To find out the different tunes sent out by a vibrating system is a problem which may or may not be solvable in certain special cases, but it would baffle the most skillful mathematicians to solve the inverse problem and to find out the shape of a bell by means of the sounds which it is capable of sending out. And this is the problem which ultimately spectroscopy hopes to solve in the case of light. In the meantime we must welcome with delight even the smallest step in the desired direction." (Mehra and Rechenberg, 2000). Actually, it was known very early, from Weyl's formula, that one can "hear" the area of a drum and the length of its perimeter (see section V.A, and (Vaa et al., 2005) for a historical account of the problem). But could the shape itself be retrieved from the spectrum? That is, what kind of information on the geometry is it possible to gather from the knowledge of the spectrum, for instance, using semiclassical methods that allow investigation of the quantum-classical correspondence? And what kind of sufficient conditions allow the geometry to be entirely specified from the spectrum?

## Is this invertable?

its spectral properties. As the spectroscopist A. Schuster put it in an 1882 report to the British Association for the Advancement of Science: "To find out the different tunes sent out by a vibrating system is a problem which may or may not be solvable in certain special cases, but it would baffle the most skillful mathematicians to solve the inverse problem and to find out the shape of a bell by means of the sounds which it is capable of sending out. And this is the problem which ultimately spectroscopy hopes to solve in the case of light. In the meantime we must welcome with delight even the smallest step in the desired direction." (Mehra and Rechenberg, 2000).
cal? More broadly, inverse problem of from knowledge of oscopist A. Schuster tish Association for ad out the different is a problem which n special cases, but hematicians to solve he shape of a bell by able of sending out. nately spectroscopy In the meantime the smallest step in Rechenberg, 2000). om Weyl's formula, rum and the length d (Vaa et al., 2005) em). But could the pectrum? That is, letry is it possible to ctruh, for instance, or investigation of And what kind of sufficient conditions allow the geometry to be entirely specified from the spectrum?

## ENGINEERED PSF FOR 3D STORM

Astigmatic


Double Helix


Huang, B., Wang, W., Bates, M., \& Zhuang, X. (2008). Three-Dimensional Super-Resolution Imaging by Stochastic Optical Reconstruction Microscopy. Science, 3 I9(5864), 8I0-8। 3. http:// doi.org/ |0.1126/science. 1153529

Pavani, S. R. P., Thompson, M. A., Biteen, J. S., Lord, S. J., Liu, N., Twieg, R. J., ... Moerner, W. E. (2009). Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function. Proceedings of the National Academy of Sciences of the United States of America, I 06(9), 2995-2999. http://doi.org/ I 0. I O73/ pnas. 0900245 I 06

## LEARNING FROM SIMULATION

## LEARNING FROM SIMULATION

$$
\gamma=\left\{\left(w_{j}, \theta_{j}\right), \cdots\right\}
$$

## LEARNING FROM SIMULATION



## LEARNING FROM SIMULATION



## LEARNING FROM SIMULATION



## LEARNING FROM SIMULATION



## LEARNING FROM SIMULATION



## LEARNING FROM SIMULATION



## 2D COMPARISON


(a)




(c)


(d)


## ADCG: Alternating Descent Conditional Gradient

## Beyond Images to Graphs

Solving Spectroscopic Inverse Problems

## Beyond Images to Graphs

Solving Spectroscopic Inverse Problems


## Beyond Images to Graphs

Solving Spectroscopic Inverse Problems


## Beyond Images to Graphs

Solving Spectroscopic Inverse Problems

The forward problem


## Beyond Images to Graphs

## Solving Spectroscopic Inverse Problems

The forward problem


The inverse problem

Beyond Images to Graphs

## Beyond Images to Graphs



## Beyond Images to Graphs



## Beyond Images to Graphs



$\square$

## Computational Measurement

How can we design new measurement systems to be more interpretable / useful for Al?

## What is your query?

How many bits are you trying to extract from your system?

## What is your query?

How many bits are you trying to extract from your system?


## What is your query?

How many bits are you trying to extract from your system?


## What is your query?

How many bits are you trying to extract from your system?


## What is your query?

How many bits are you trying to extract from your system?


## What is your query?

How many bits are you trying to extract from your system?


1-bit measurement system

Tumor?<br>[] Yes<br>[] No

## What is your query?

How many bits are you trying to extract from your system?


1-bit measurement system

Tumor?<br>[] Yes<br>[] No



## What is your query?

How many bits are you trying to extract from your system?


1-bit measurement system

Tumor?<br>[] Yes<br>[] No



JPEG


170 kB

## What is your query?

 How many bits are you trying to extract from your system?

1-bit measurement system

Tumor?<br>[] Yes<br>[] No



JPEG


Is it possible to only collect 170 kB of data?

## Sgin Processing Compressive Sensing

SENSTNG, 2AMPTNG AND COMJPiESSJON
ALL FOR ONEOR
ONE FORALL?

LOCALITYYSENSITIVE HASHING
SP WITH BELIEF PROPAGATION WHY GAUSSIANIITY

IEEE Signal Processing Magazine March 2008


(a)

(b)

(c)
[FIG2] Single-pixel photo album. (a) $256 \times 256$ conventional image of a black-and-white R. (b) Single-pixel camera reconstructed image from $M=1,300$ random measurements ( $50 \times$ sub-Nyquist). (c) $256 \times 256$ pixel color reconstruction of a printout of the Mandrill test image imaged in a low-light setting using a single photomultiplier tube sensor, RGB color filters, and $M=6,500$ random measurements.

## DiffuserCam

## Single-shot 3D acquisition



Nick Antipa, Grace Kuo, Reinhard Heckel, Ben Mildenhall, Emrah Bostan, Ren Ng, and Laura Waller,
"DiffuserCam: lensless single-exposure 3D imaging," Optica 5, 1-9 (2018)

## DiffuserCam

## Single-shot 3D acquisition



## How can AI help?



## Active Learning

How can Al guide experimentation and measurement?

How to run an experiment

## How to run an experiment

## Passive Observation



## How to run an experiment

## Passive Observation



Optimal Experiment Design

| 0 | 0 | 0 | 0 | 0 | 0 |
| :--- | :--- | :--- | :--- | :--- | :--- |
| 1 | 1 | 1 | 1 | 1 | 1 |
| 2 | 2 | 2 | 2 | 2 | 2 |
| 3 | 3 | 3 | 3 | 3 | 3 |
| 4 | 4 | 4 | 4 | 4 | 4 |
| 5 | 5 | 5 | 5 | 5 | 5 |
| 6 | 6 | 6 | 6 | 6 | 6 |
| 7 | 7 | 7 | 7 | 7 | 7 |
| 8 | 8 | 8 | 8 | 8 | 8 |
| 9 | 9 | 9 | 9 | 9 | 9 |



## How to run an experiment

## Passive Observation



Optimal Experiment Design
$\begin{array}{llllll}0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 2 \\ 3 & 3 & 3 & 3 & 3 & 3 \\ 4 & 4 & 4 & 4 & 4 & 4 \\ 5 & 5 & 5 & 5 & 5 & 5 \\ 6 & 6 & 6 & 6 & 6 \\ 7 & 7 & 7 & 7 & 7 \\ 8 & 8 & 8 & 8 & 8 & 8\end{array}$

Active Learning


## Measuring Receptive Fields

## Measuring Receptive Fields



## Measuring Receptive Fields



## Measuring Receptive Fields



## Measuring Receptive Fields



## Measuring Receptive Fields



## Measuring Receptive Fields



## Active Learning Receptive Fields

## Active Learning Receptive Fields



## Active Learning Receptive Fields



B


## Active Learning Receptive Fields



## Active Learning Receptive Fields



- Implementation is hard and complex!



## Active Learning Receptive Fields



- Implementation is hard and complex!
- Even though theory bounds error, theory is based on a model (which may be wrong!)


B

62.82

400
trials


1000
trials

## Active Learning Receptive Fields



- Implementation is hard and complex!
- Even though theory bounds error, theory is based on a model (which may be wrong!)
- Easier to just spend \$3B and scale up experiments


B

62.82

400
trials

57.29

1000
trials



## Finding the sweet spot

## How much Al progress is necessary?

## How much Al progress is necessary?

## How much Al progress is necessary?

"How can I
analyze my data?"

## How much Al progress is necessary?

"Existing algorithms
aren't quite right"

```
"How can I analyze my data?"
```


## How much Al progress is necessary?

"Existing algorithms
aren't quite right"
"How can I
analyze my data?"
"Can you replace my grad student?"

## How much Al progress is necessary?

"Existing algorithms
aren't quite right"

## Regression

Sentience
"How can I
analyze my data?"
"Can you replace my grad student?"

## How much Al progress is necessary?

"Existing algorithms
aren't quite right"
Regression

| "How can I |
| :---: |
| analyze my |
| data?" |

sweet spot for us
my grad student?"

## How much Al progress is necessary?

"Existing algorithms aren't quite right"

## Regression

Sentience
"How can I analyze my data?"

- ML provides fundamentally new
"Can you replace my grad student?" capabilities but is "mostly there" already


## How much Al progress is necessary?



- Creative ideas beyond existing work - rethinking what's possible


## How much Al progress is necessary?

## "Existing algorithms aren't quite right"

## Regression

"How can I analyze my data?"
sweet spot for us

- ML provides fundamentally new capabilities but is "mostly there" already
- Creative ideas beyond existing work - rethinking what's possible
- Existing baselines so you understand how much progress can be made
"Can you replace my grad student?"

What is a robot?

## What is a robot?



## What is a robot?



## What is a robot?



## What is a robot?



What is a robot?

## What is a robot?



TELESCOPES

## What is a robot?



TELESCOPES


MR IMAGERS

## What is a robot?



TELESCOPES


## What is a robot?



TELESCOPES


## What is a robot?



TELESCOPES


ACCELERATORS

## What is a robot?



TELESCOPES


MR IMAGERS


SPECTROMETERS


CYTOMETERS


ACCELERATORS

...AND MORE

## The future of $\mathrm{Al}+$ Measurement

Every scientific instrument is a robot and can be smarter

Tremendous opportunities for collaboration across U. Chicago and National Labs
Getting people in the same room is just a start - they have to speak each other's


ACCELERATORS language!

## Extra Slides

## Lack of automated analysis inhibits scientific advancement

Contemporary computational approaches for spectral analysis are effectively library lookups
(thus can only find "known knowns")
Yet there are $>10^{60}$ possible small molecules!
(obviously impossible to build comprehensive libraries)

## Inhibits scale

$80 \%$ of human small metabolites are still unknown [1]
(even worse for other organisms)
Crude oil can have over 1,000,000 unique compounds [2] and its composition is still a mystery [3]

## Inhibits Robotic Laboratories

Automated synthetic chemistry and wet lab platforms are coming online

But how do we know what they made? Can't close the loop if you can't measure the output

## Inhibits new instrumentation

Quantum Sensing and other molecule sensing techniques are increasingly viable

Often have fundamentally different tradeoffs from classical instrumentation, resulting in tremendous data interpretation challenges

## Key insight:

Spectroscopy is an inverse problem

What is an inverse problem?

## What is an inverse problem?

## Model System

Physical properties, unknowns

## Observables

Measurements
and data

## What is an inverse problem?

The forward problem


## What is an inverse problem?

The forward problem


The inverse problem
(hard)

## What is an inverse problem?



## 20th Century Measurement

Linear, continuous inverse problems transformed measurement in the latter half of the 20th century


## Molecular spectroscopy as an inverse problem



## Molecular spectroscopy as an inverse problem

The forward problem


## Forward

Calculating the spectrum for a given structure

## SCF/DFT solves the forward problem for many modalities of interest

Calling this "easy" is a stretch - performance is cubic in the number of atoms and many aspects of experimental setup (conformational diversity, salvation, etc.) are still challenging.

## Molecular spectroscopy as an inverse problem

The forward problem


The inverse problem
(hard)

## Forward

Calculating the spectrum for a given structure
SCF/DFT solves the forward problem for many modalities of interest

Calling this "easy" is a stretch - performance is cubic in the number of atoms and many aspects of experimental setup (conformational diversity, salvation, etc.) are still challenging.

## Inverse

Deducing the structure for a given spectrum
This is incredibly challenging, a longstanding open problem

Highly nonlinear forward model
Combinatorial solution space
Single correct structure!

The forward problem


The inverse problem
(hard)

## The forward problem



The inverse problem

## Al advances make this possible

We're developing Al techniques to solve this.

Physics-informed deep learning and graph neural networks let us generate millions of synthetic spectra for training data



Structured prediction via deep imitation learning lets us learn to build molecules consistent with observed spectra

Deep latent variable models let us model and understand physical measurement processes where ab initio techniques fail.


## The forward problem



The inverse problem
(hard)

## Al advances make this possible

We're developing Al techniques to solve this.

Physics-informed deep learning and graph neural networks let us generate millions of synthetic spectra for training data



Structured prediction via deep imitation learning lets us learn to build molecules consistent with observed spectra

Deep latent variable models let us model and understand physical measurement processes where ab initio techniques fail.


## Research Objectives

We've had early success in NMR and are moving into mass spec

NMR

Fast model to simulate spectra: DFT accuracy in milliseconds

Can predict correct structure with high accuracy on a wide variety of compounds from 1D 13C spectrum

## MS

Computational forward model still an open research challenge!

Next big challenge

## Methodologies

We learn to build molecules from spectroscopic data by first taking them apart

## 1. We need a lot of training data!



Solution: build a fast approximation to the forward mode that lets you generate 100 M synthetic spectra
(Bootstrapped from 30k experimental spectra)
[1] Jonas, Kuhn. Rapid prediction of NMR spectral properties with quantified uncertainty. Journal of Cheminformatics, 11(1): 2019.

## Methodologies

We learn to build molecules from spectroscopic data by first taking them apart

## 1. We need a lot of training data!

${ }^{13} \mathrm{C}$


Solution: build a fast approximation to the forward model that lets you generate 100 M synthetic spectra
(Bootstrapped from 30k experimental spectra)
[1] Jonas, Kuhn. Rapid prediction of NMR spectral properties with quantified uncertainty. Journal of Cheminformatics, 11(1): 2019.

## 2. Generating complete molecules is hard!

Solution: Construct molecule incrementally - use deep imitation learning to learn to place the next bond of a partial molecule [2]


[^1]
## Methodologies

We learn to build molecules from spectroscopic data by first taking them apart

## 1. We need a lot of training data!



Solution: build a fast approximation to the forward model that lets you generate 100 M synthetic spectra
(Bootstrapped from 30k experimental spectra)
[1] Jonas, Kuhn. Rapid prediction of NMR spectral properties with quantified uncertainty. Journal of Cheminformatics, 11(1): 2019.

## 2. Generating complete molecules is hard!

Solution: Construct molecule incrementally - use deep imitation learning to learn to place the next bond of a partial molecule [2]

## 3.Generate candidate structures



Search tree generating candidate structures from observed spectrum using learned function


[^2]
## Methodologies

We learn to build molecules from spectroscopic data by first taking them apart

## 1. We need a lot of training data!

Solution: build a fast approximation to the forward model that lets you generate 100 M synthetic spectra
(Bootstrapped from 30k experimental spectra)
[1] Jonas, Kuhn. Rapid prediction of NMR spectral properties with

quantified uncertainty. Journal of Cheminformatics, 11(1): 2019.

## 2. Generating complete molecules is hard!

Solution: Construct molecule incrementally - use deep imitation learning to learn to place the next bond of a partial molecule [2]


[^3]
## 3.Generate candidate structures



Search tree generating candidate structures from observed spectrum using learned function

## 4. Use fast forward model to validate!



We predict the right structure $68 \%$ of the time (up from $56 \%$ in 2020)

## Three Phases

Build fast precise forward model

Real-time solution to
Inverse problem

Compute optimal next measurement

## Three Phases

## Build fast precise forward model

## Real-time solution to Inverse problem

## Compute optimal next measurement

Why start with commodity spectroscopic modalities like NMR and MS?
Existing Data

NMRShiftDB : 50k NMR exp
NIST-17 : 250k GC/EI-MS exp
SDBS: 20k NMR exp
MassBank: 50k LC-MS/MS exp

## Reliable Ubiquitous Hardware

At UChicago we have:
7 NMR specrometers
(Bruker 400MHz+)
8 MS instruments
(GC-EI/MS, QTOF, incoming Thermo Orbitraps)

## Existing Platforms are programmable

Custom real-time pulse sequence design via Bruker hardware

MS HW enables programmatic control over collision energies
and peak selection for
fragmentation

# Developing new techniques for acquisition 

Modern Al requires tons of data
We are developing solutions to generate massive quantities of training data called

# Developing new techniques for acquisition 

## Modern AI requires tons of data

We are developing solutions to generate massive quantities of training data called
A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures

## Developing new techniques for acquisition

## Modern AI requires tons of data

We are developing solutions to generate massive quantities of training data called
A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures


Completely solve MS forward model for \$100k

## Developing new techniques for acquisition

## Modern AI requires tons of data

A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures


Generate 10k known reaction products in a day for $\$ 5 \mathrm{k}$

We are developing solutions to generate massive quantities of training data called

Completely solve MS forward model for \$100k

## Developing new techniques for acquisition

## Modern AI requires tons of data

A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures


Generate 10k known reaction products in a day for $\$ 5 \mathrm{k}$

We are developing solutions to generate massive quantities of training data called
"Shotgun spectrometry"

Completely solve MS forward model for \$100k

## Developing new techniques for acquisition

## Modern Al requires tons of data

A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures


Generate 10k known reaction products in a day for \$5k

We are developing solutions to generate massive quantities of training data called
"Shotgun spectrometry"

1. Generate a mixture of 10 k molecules with maximal molecular weight diversity
2. Perform LC/MS/MS on the combined mixture at each level
3. Train our new models on the resulting lightly-separated mixtures - some overlap will exist and that's ok

Completely solve MS forward model for \$100k

## Developing new techniques for acquisition

## Modern Al requires tons of data

A modern MS machine can select a single $\mathrm{m} / \mathrm{z}$ before fragmentation - useful for complex mixtures


Generate 10k known reaction products in a day for \$5k

We are developing solutions to generate massive quantities of training data called
"Shotgun spectrometry"

1. Generate a mixture of 10 k molecules with maximal molecular weight diversity
2. Perform LC/MS/MS on the combined mixture at each level
3. Train our new models on the resulting lightly-separated mixtures - some overlap will exist and that's ok

Can potentially scale CombChem to 100k per batch

Completely solve MS forward model for \$100k

## Research Plan



## Self-driving spectrometers

## Designing Algorithms, Software, and Systems to Measure Every Molecule



Scalable measurement of complex mixtures

## laboratory automation



Ultimately allowing novel spectroscopic techniques


Better fast forward models via physicsinformed Graph NNs


## Structured prediction via

Deep Imitation Learning
By new Al techniques:



Optimal experimentation via real-time active learning


[^0]:    "A common belief in model-free reinforcement learning is that methods based on random search in the parameter space of policies exhibit significantly worse sample complexity than those that explore the space of actions. We dispel such beliefs by introducing a random search method for training static, linear policies for continuous control problems, matching state-of-theart sample efficiency on the benchmark MuJoCo locomotion tasks.'

[^1]:    [2] Jonas, Deep Imitation learning for Molecular Inverse Problems, NeurIPS 2019

[^2]:    [2] Jonas, Deep Imitation learning for Molecular Inverse Problems, NeurIPS 2019

[^3]:    [2] Jonas, Deep Imitation learning for Molecular Inverse Problems, NeurIPS 2019

