TENSOR SPIN OBSERVABLES

WORKSHOP

Modern Measurement Tools

And the algorithms changing everything…

D Keller

UVA - Spin

ECT*
EUROPEAN CENTRE
FOR THEORETICAL STUDIES
IN NUCLEAR PHYSICS AND RELATED AREAS

Contents

- Seems to be something going on in AI
- Examples
- Positioning for the greatest leaps forward
- Why its essential for spin-1

ARTIFICIAL INTELLIGENCE TERMS

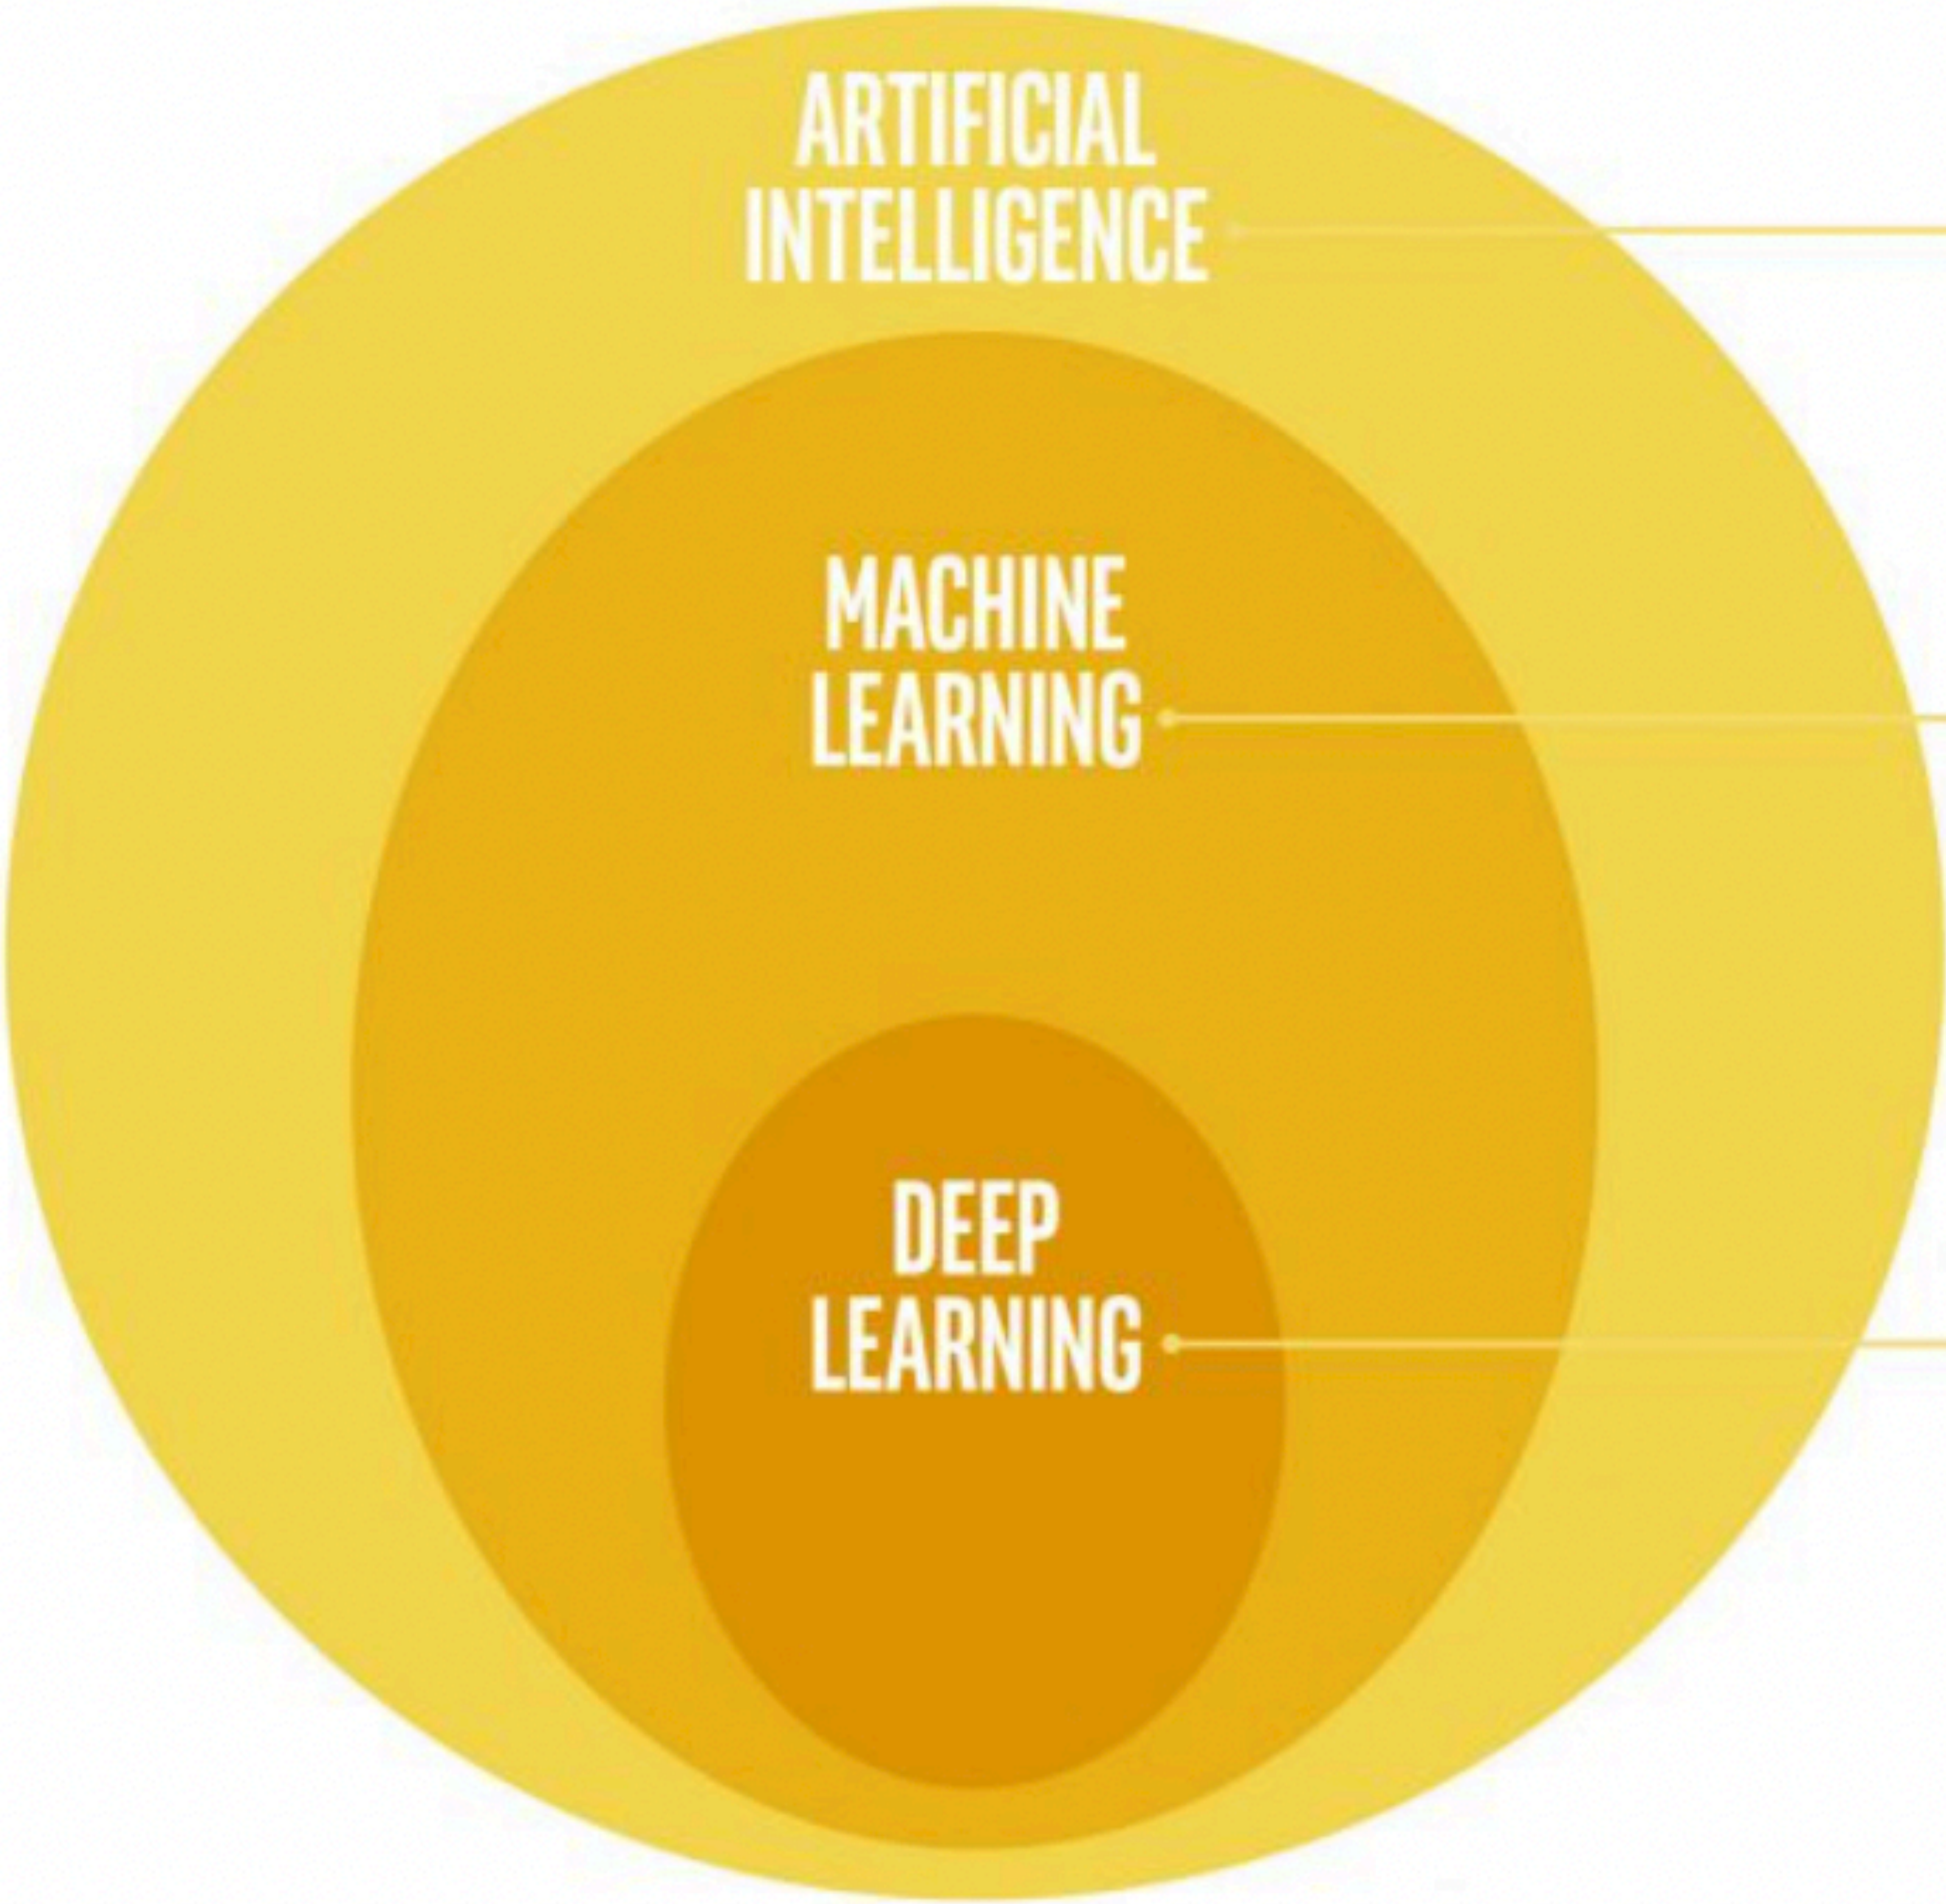

- **AI** is an umbrella term for machines capable of perception, logic, and learning.
- **Machine learning** employs algorithms that learn from data to make predictions or decisions, and whose performance improves when exposed to more data over time.
- **Deep learning** uses many-layered neural networks to build algorithms that find the best way to perform tasks on their own, based on vast sets of data.

First Perceptron Paved the way for AI

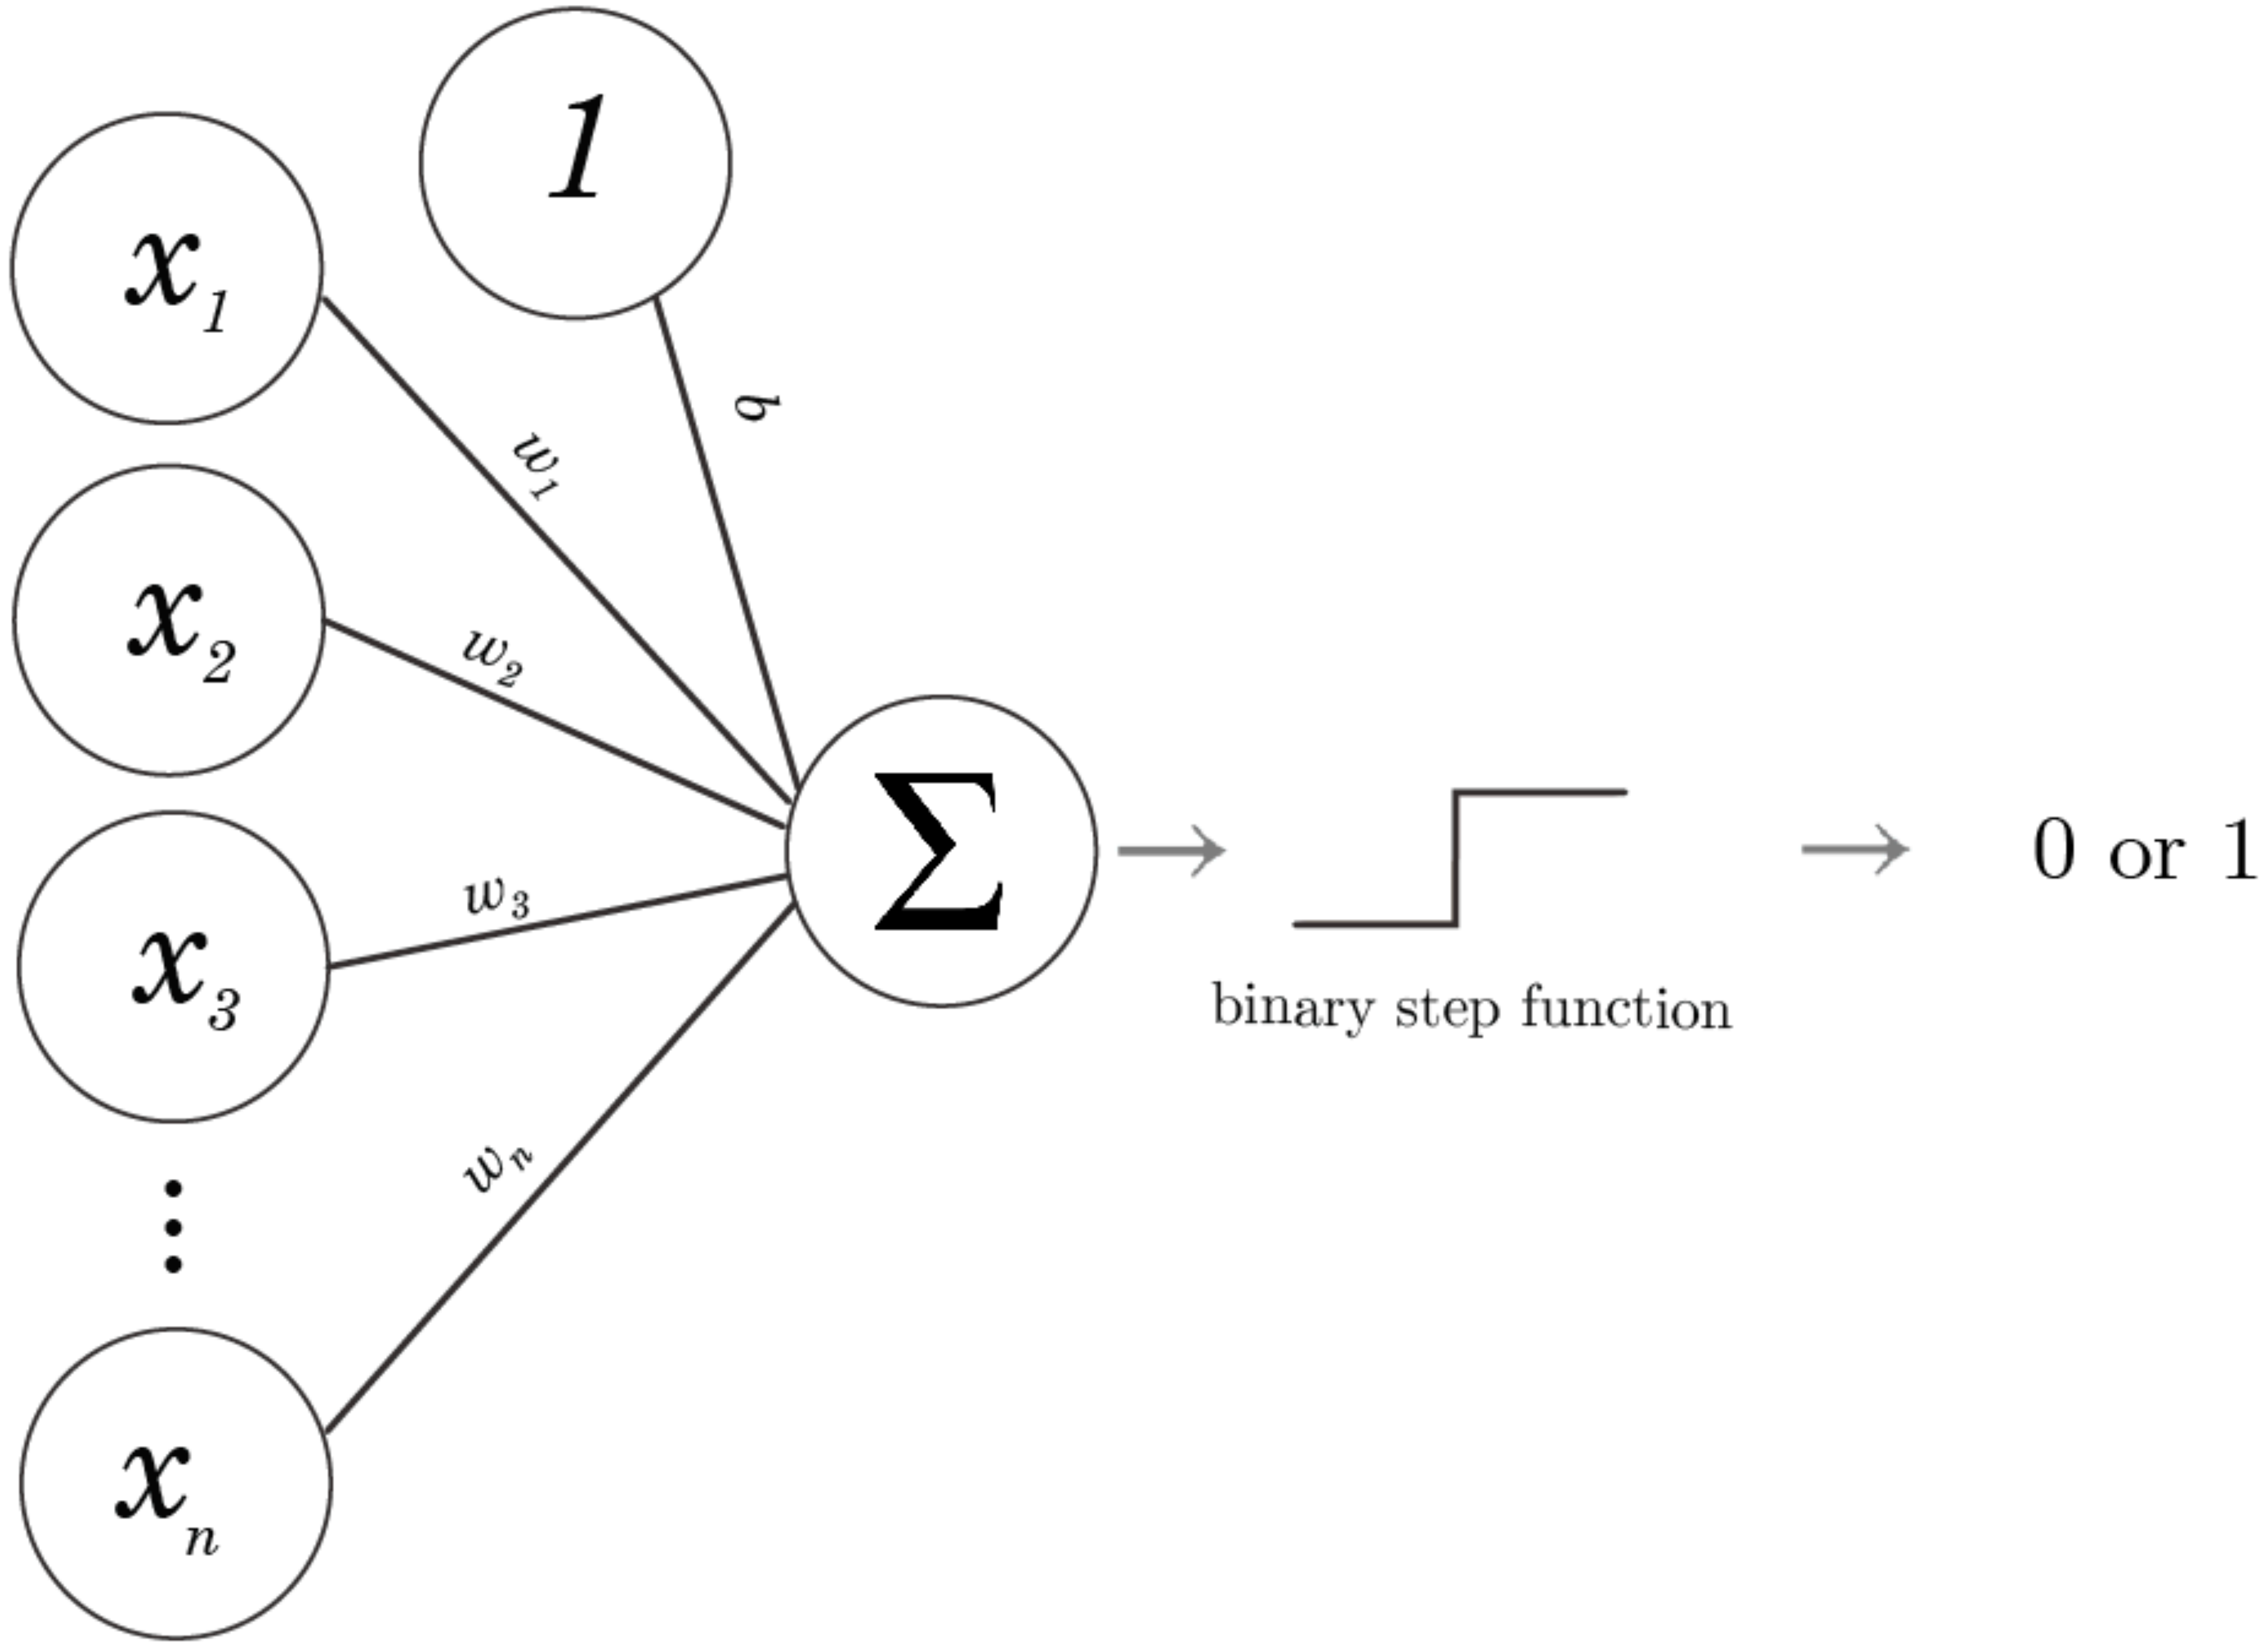

Frank Rosenblatt in July 1958, the U.S. Office of Naval Research

Birth of Modern Algorithm

In the mid 80s Geoff Hinton discovered how to make deep ANNs

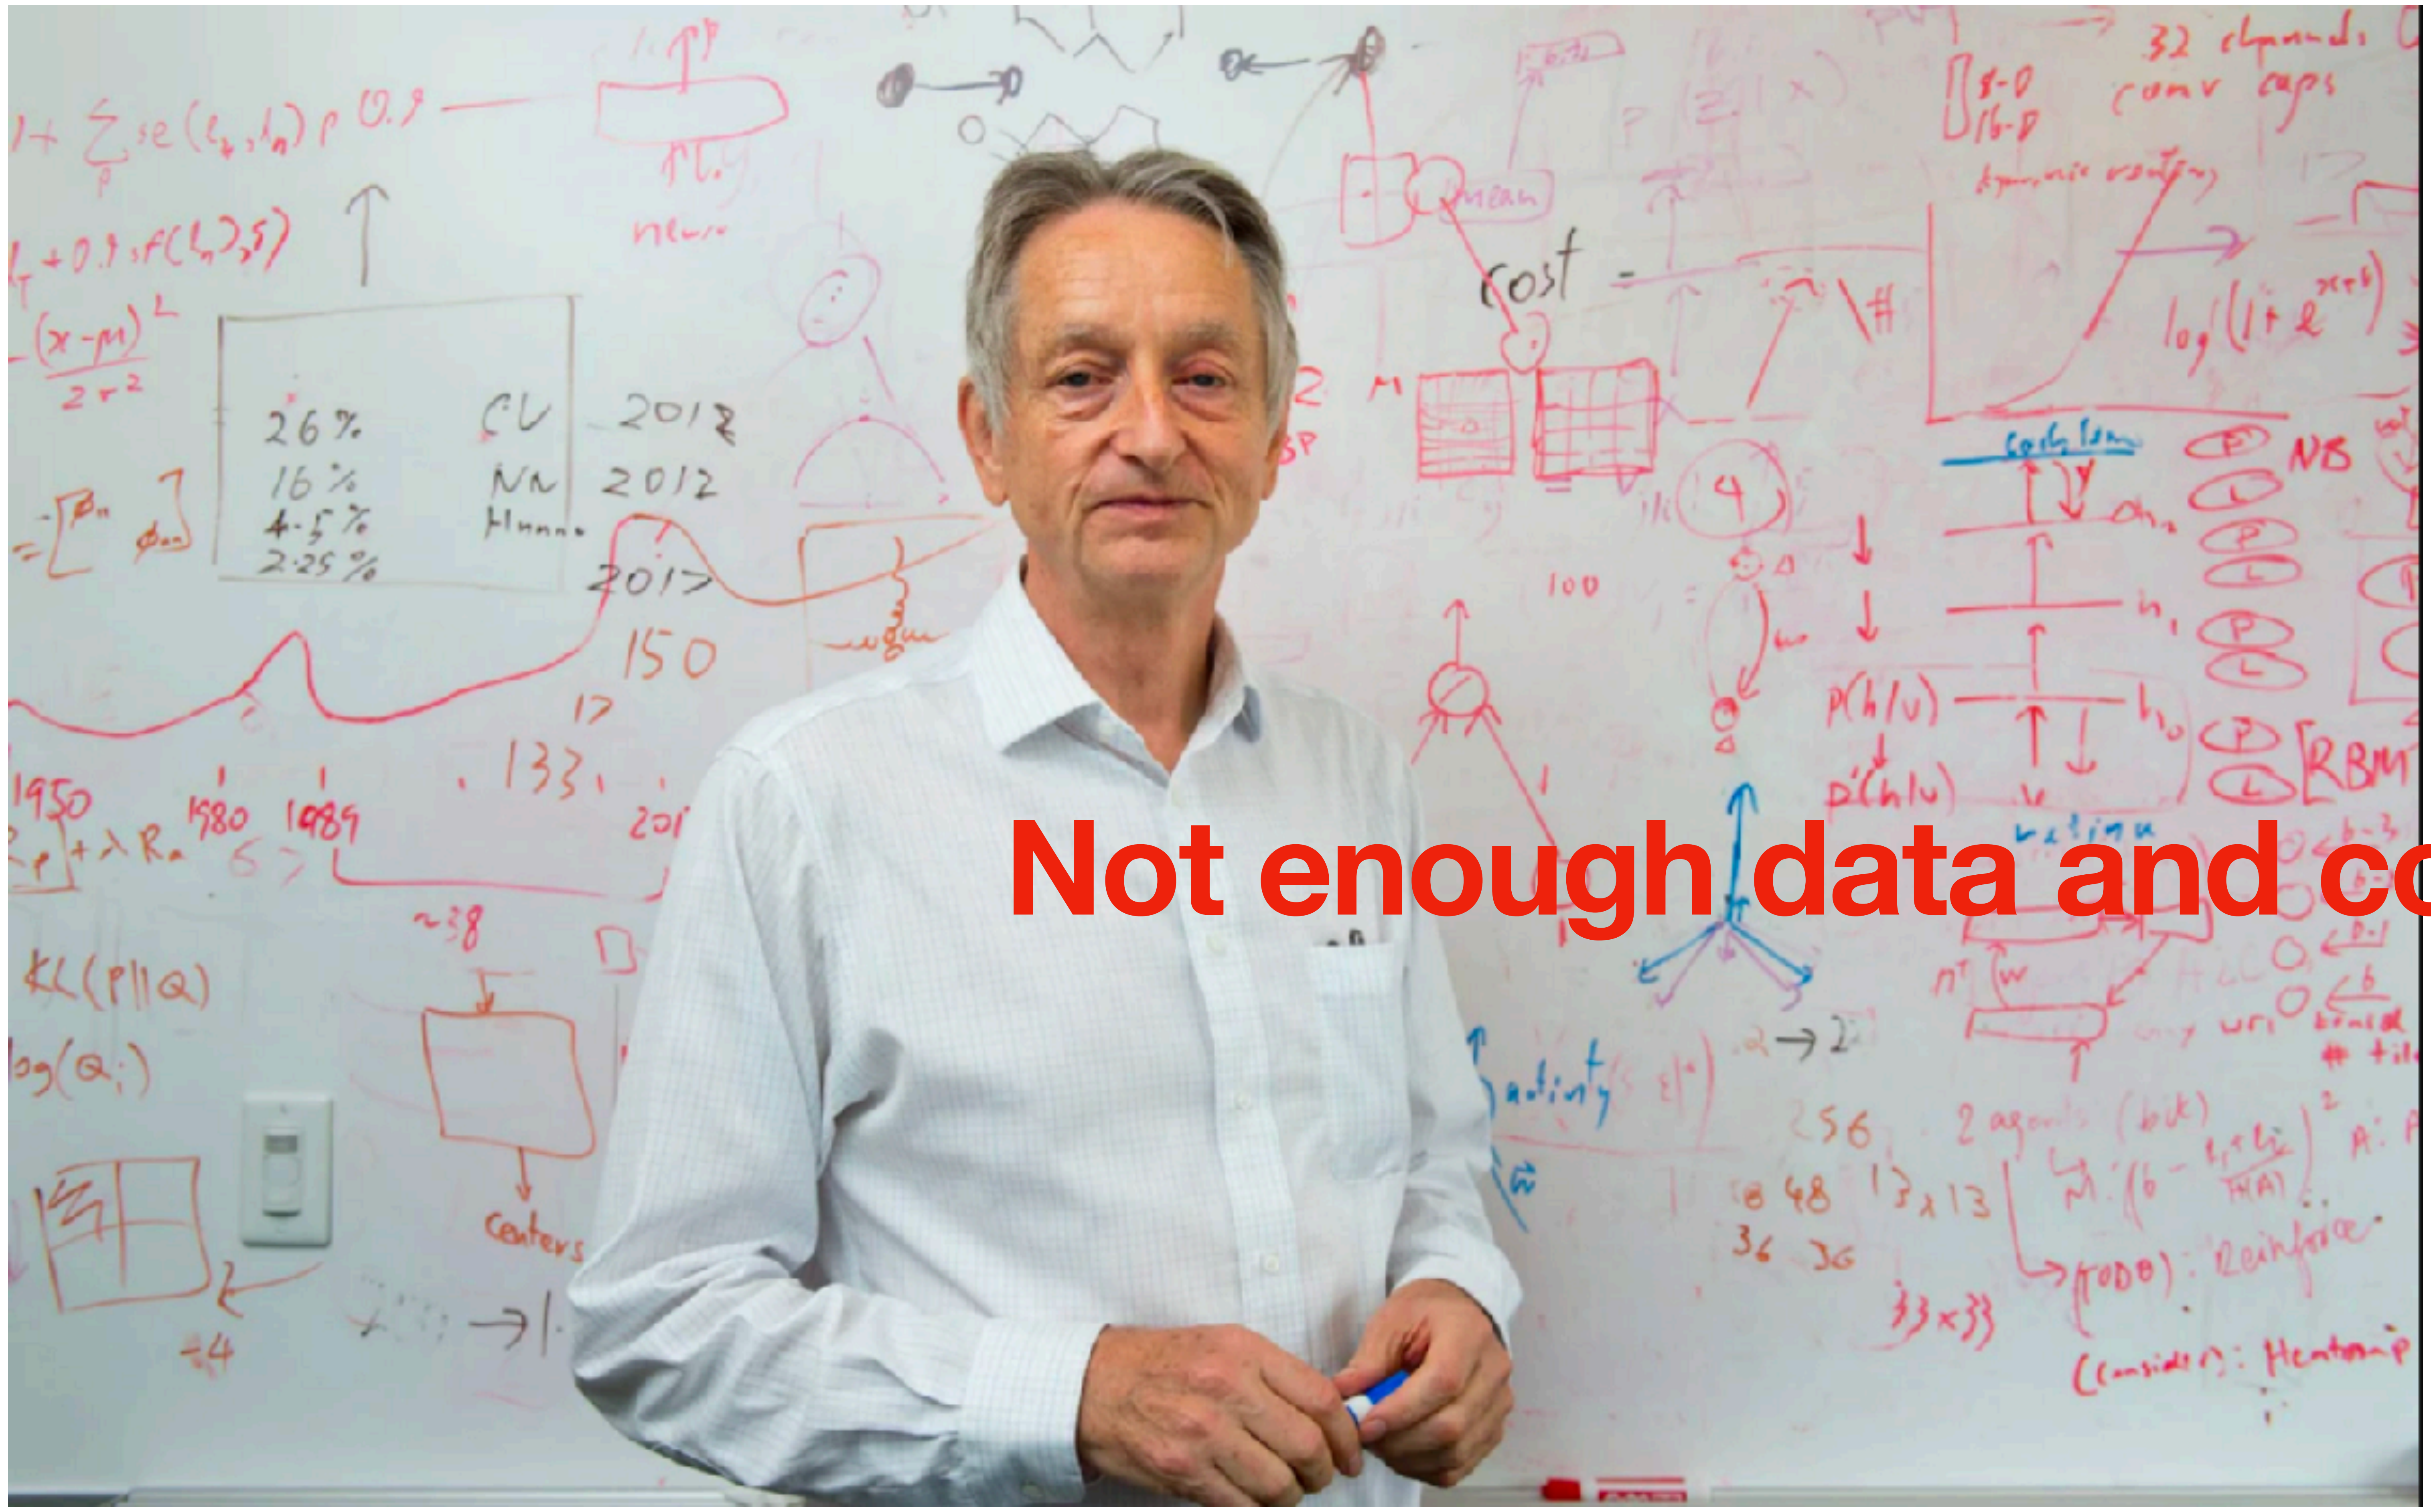

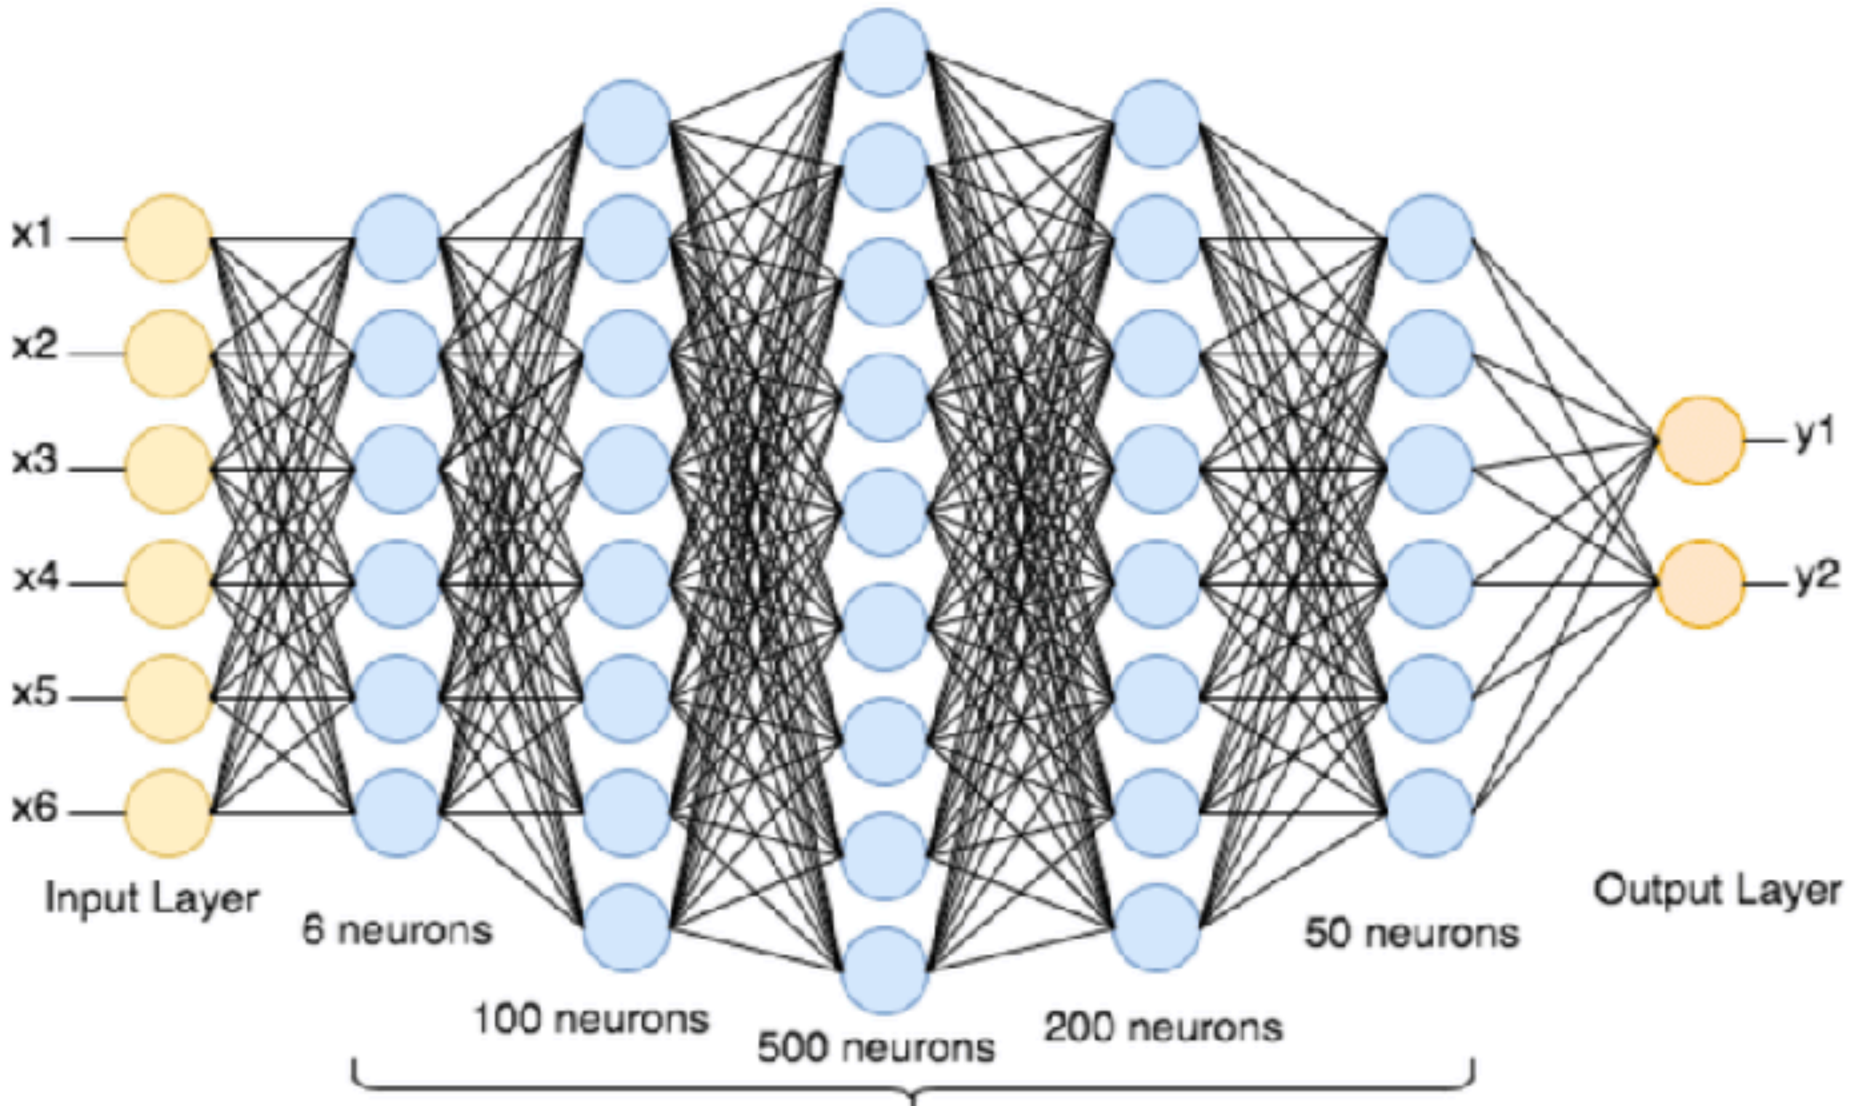

Not enough data and computing power

- Dean Pomerleau of Carnegie Mellon (1988 first AI self driving car)
- Yann LeCun (1989 handwritten digit classification)
- …

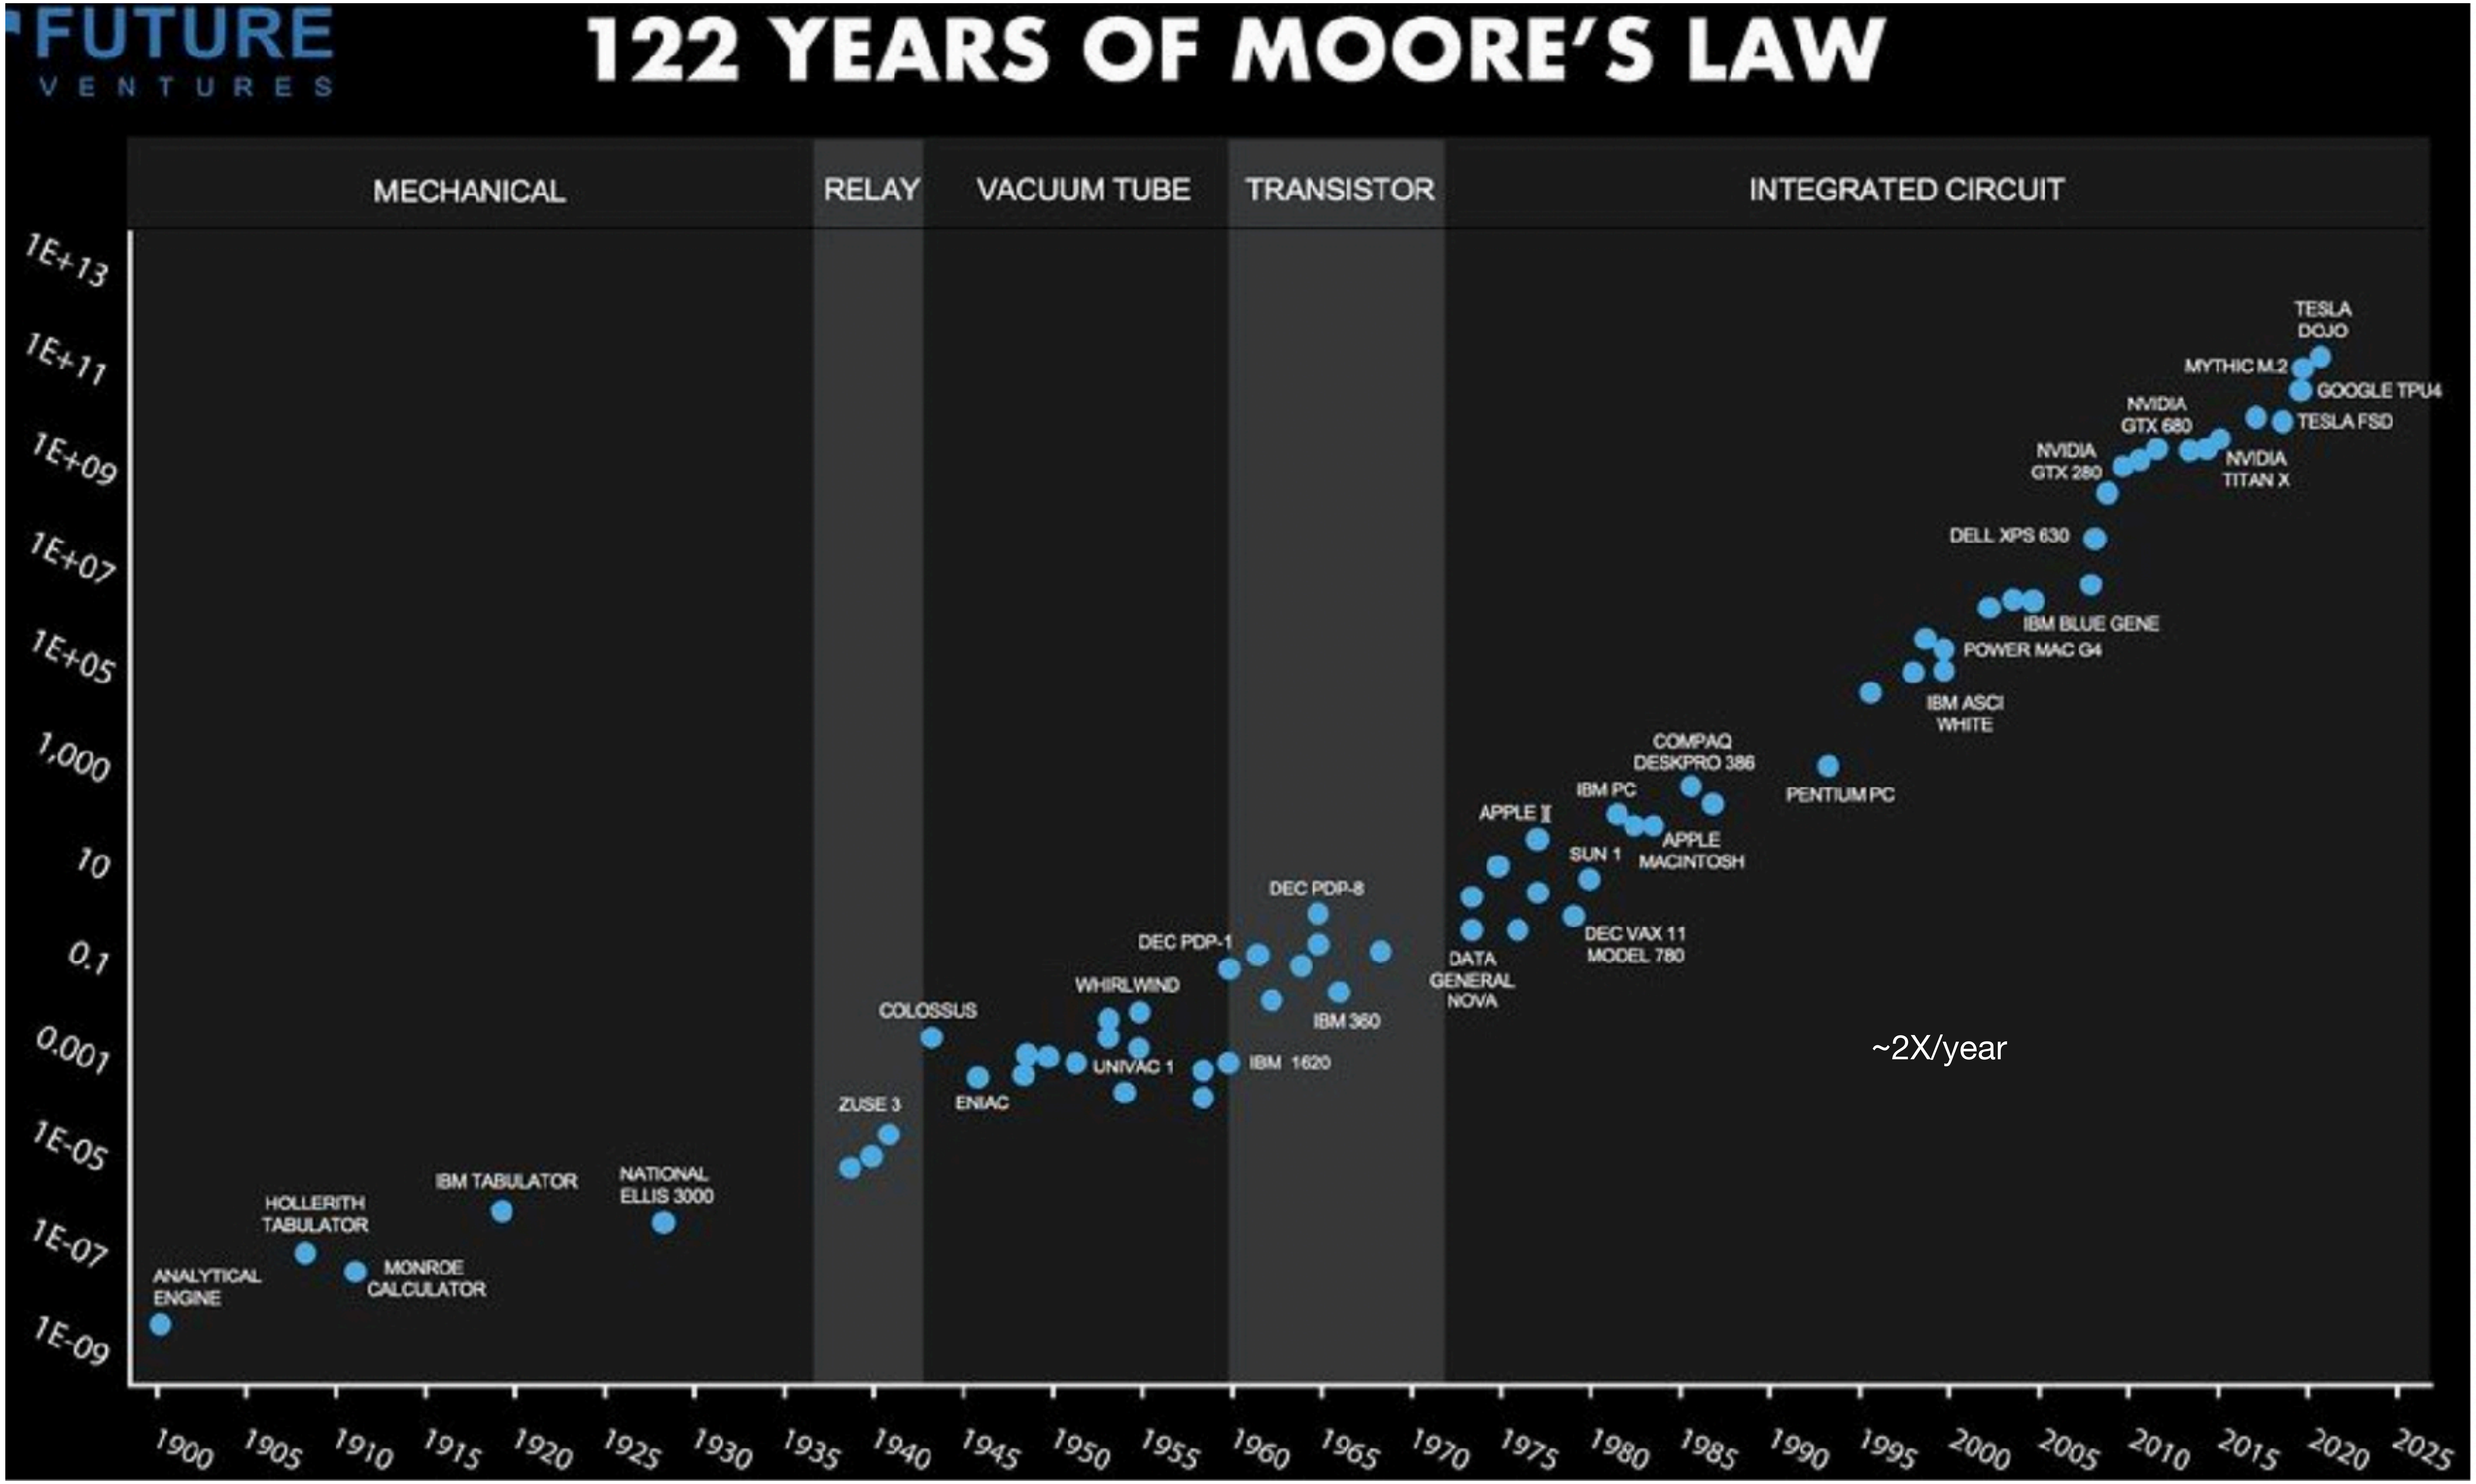

FUTURE VENTURES
122 YEARS OF MOORE'S LAW
MECHANICAL
RELAY
VACUUM TUBE
TRANSISTOR
INTEGRATED CIRCUIT
1E+13
1E+11
1E+09
1E+07
1E+05
1,000
10
0.1
0.001
1E-05
1E-07
1E-09
ANALYTICAL ENGINE
HOLLERITH TABULATOR
MONROE CALCULATOR
IBM TABULATOR
NATIONAL ELLIS 3000
ZUSE 3
COLOSSUS
ENIAC
WHIRLWIND
UNIVAC 1
IBM 1620
DEC PDP-1
DEC PDP-8
IBM 360
DATA GENERAL NOVA
APPLE II
SUN 1
IBM PC
DEC VAX 11 MODEL 780
APPLE MACINTOSH
COMPAQ DESKPRO 386
PENTIUM PC
IBM ASCI WHITE
POWER MAC G4
IBM BLUE GENE
DELL XPS 630
NVIDIA GTX 280
NVIDIA GTX 680
NVIDIA TITAN X
MYTHIC M.2
TESLA DOJO
GOOGLE TPU4
TESLA FSD
~2X/year
1900 1905 1910 1915 1920 1925 1930 1935 1940 1945 1950 1955 1960 1965 1970 1975 1980 1985 1990 1995 2000 2005 2010 2015 2020 2025

Mores Law Caught up in 2006

Algorithm Adapted to modern computing

A fast learning algorithm for deep belief nets *

Geoffrey E. Hinton and **Simon Osindero**
Department of Computer Science University of Toronto
10 Kings College Road
Toronto, Canada M5S 3G4
{hinton, osindero}@cs.toronto.edu

Yee-Whye Teh
Department of Computer Science
National University of Singapore
3 Science Drive 3, Singapore, 117543
tehyw@comp.nus.edu.sg

- Advancement in Processors, Memory, Storage
- Essential high speed systems become available (parallelism)
- Large scale datasets (the world becomes teaming with data)

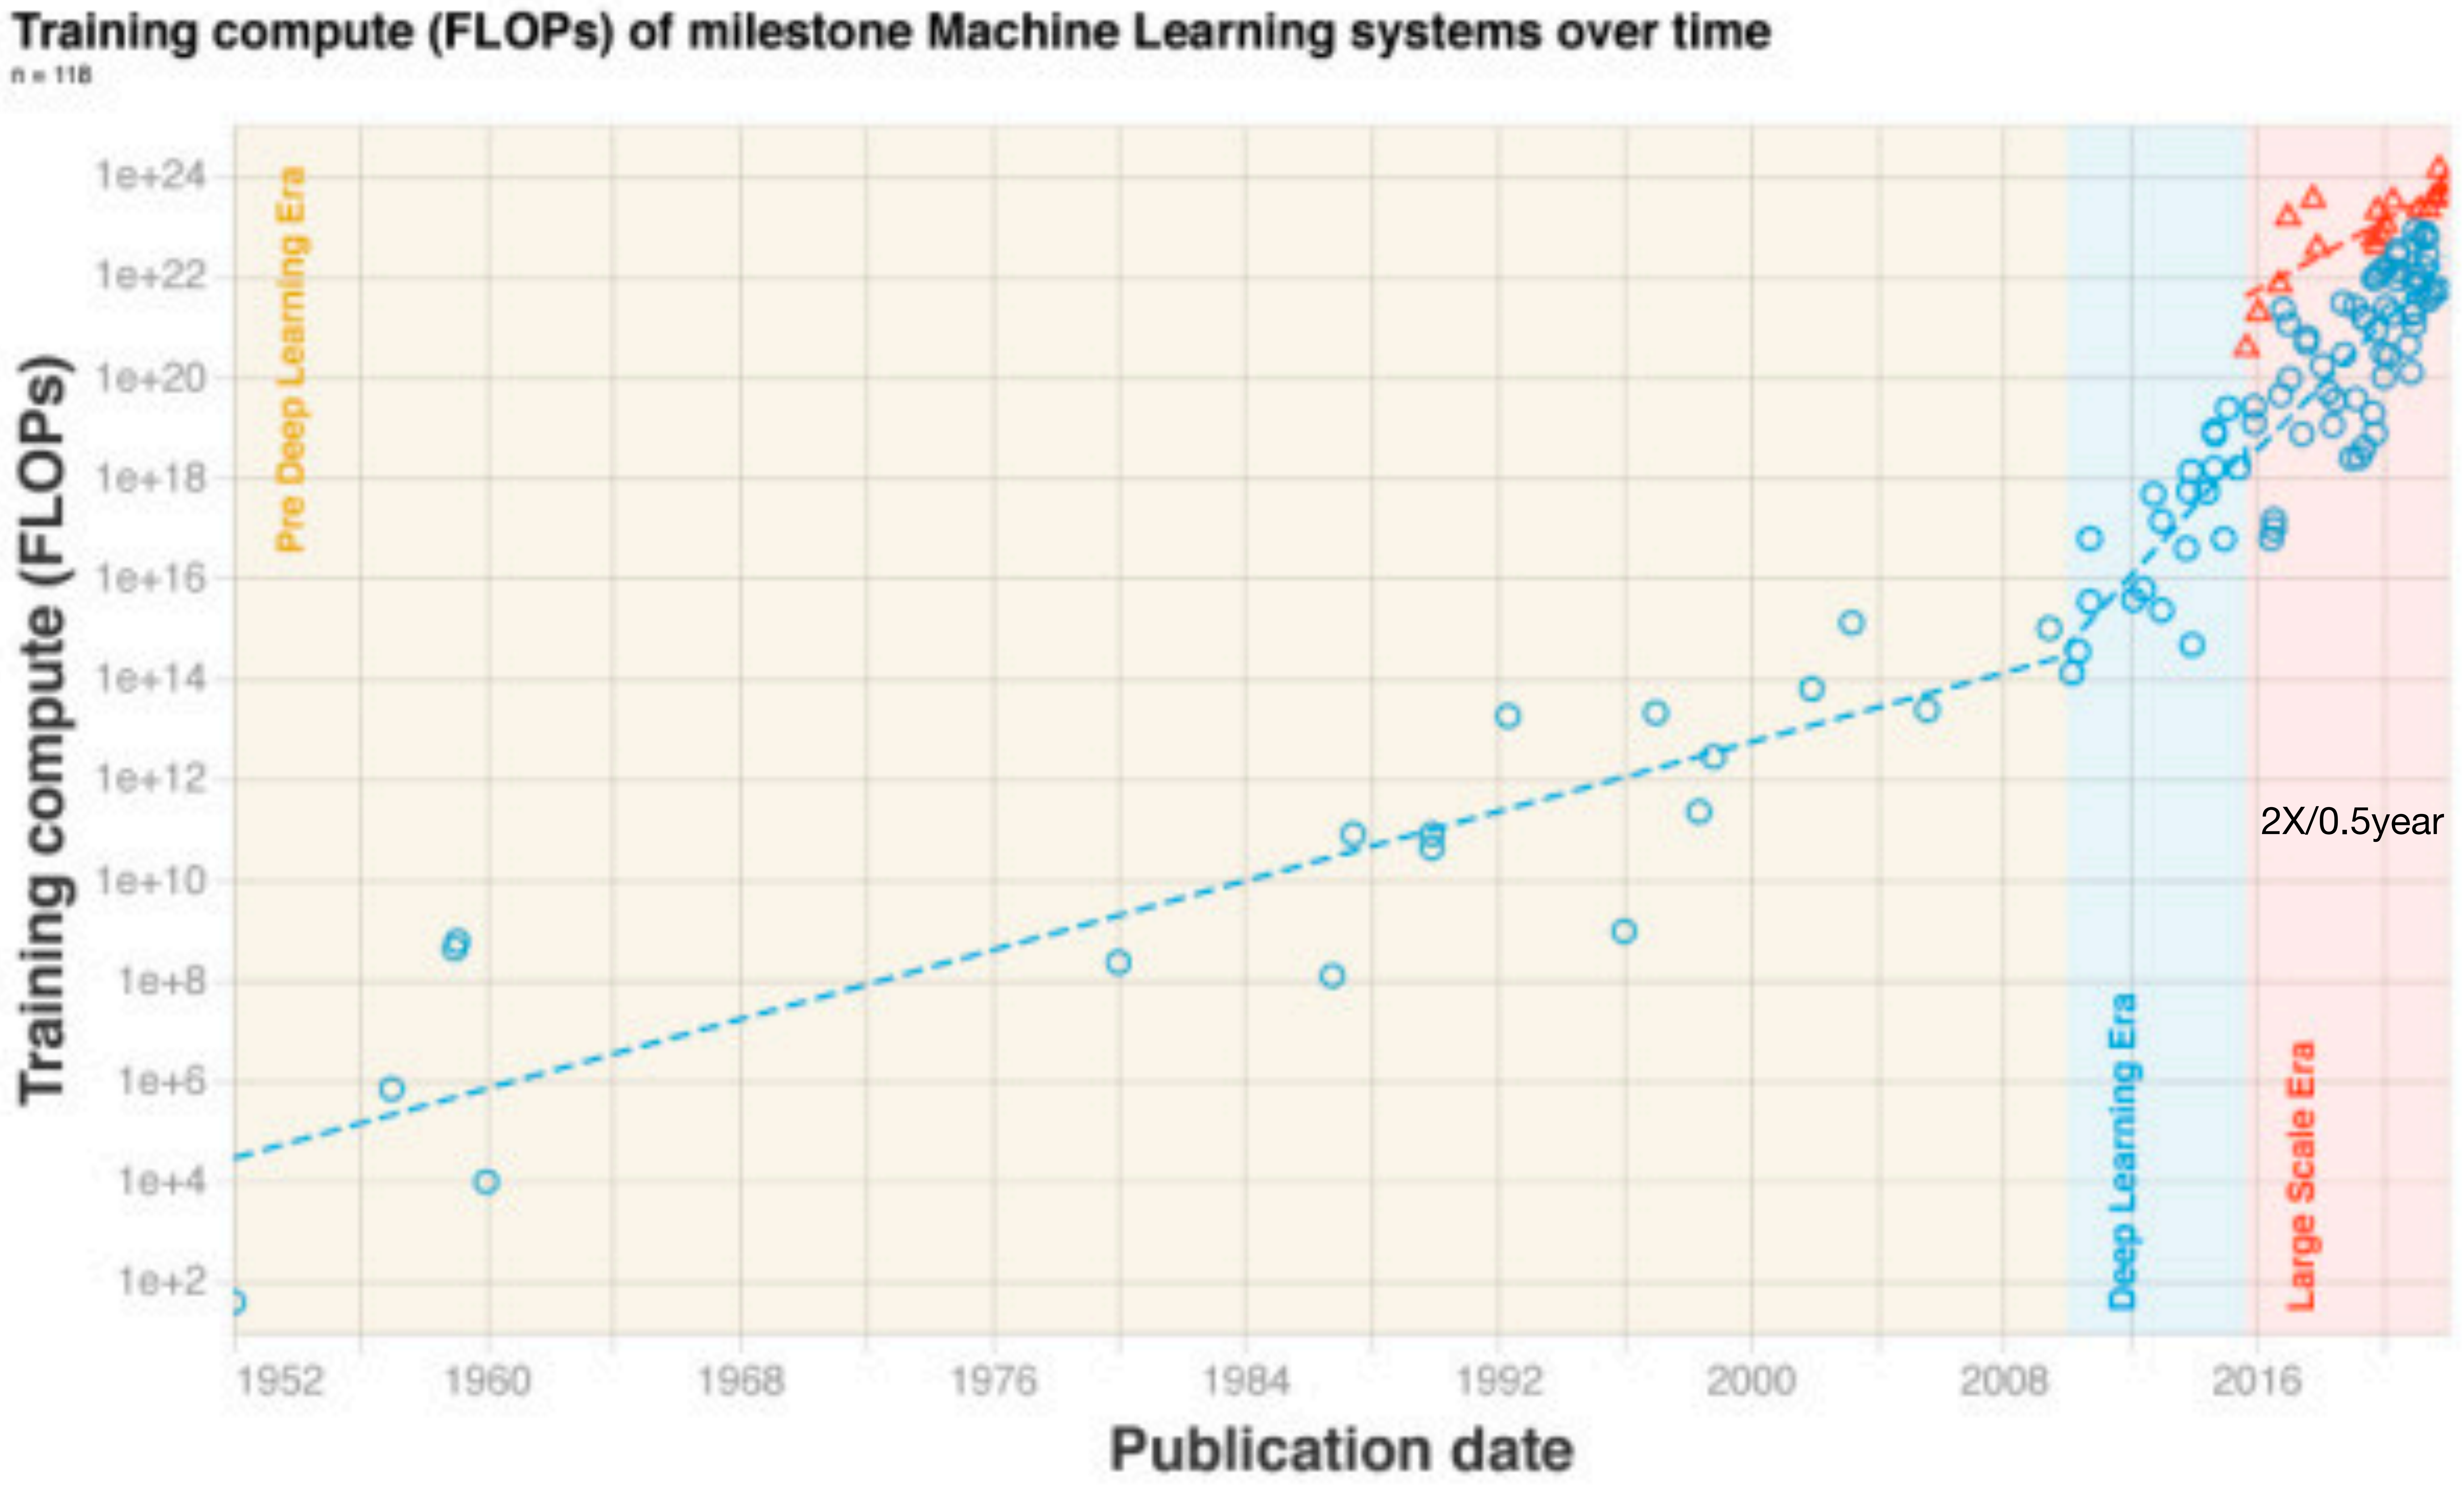
Training compute (FLOPs) of milestone Machine Learning systems over time
n = 118
Training compute (FLOPs)
1e+24
1e+22
1e+20
1e+18
1e+16
1e+14
1e+12
1e+10
1e+8
1e+6
1e+4
1e+2
Pre Deep Learning Era
Deep Learning Era
Large Scale Era
2X/0.5year
1952
1960
1968
1976
1984
1992
2000
2008
2016
Publication date

Evolution Timeline

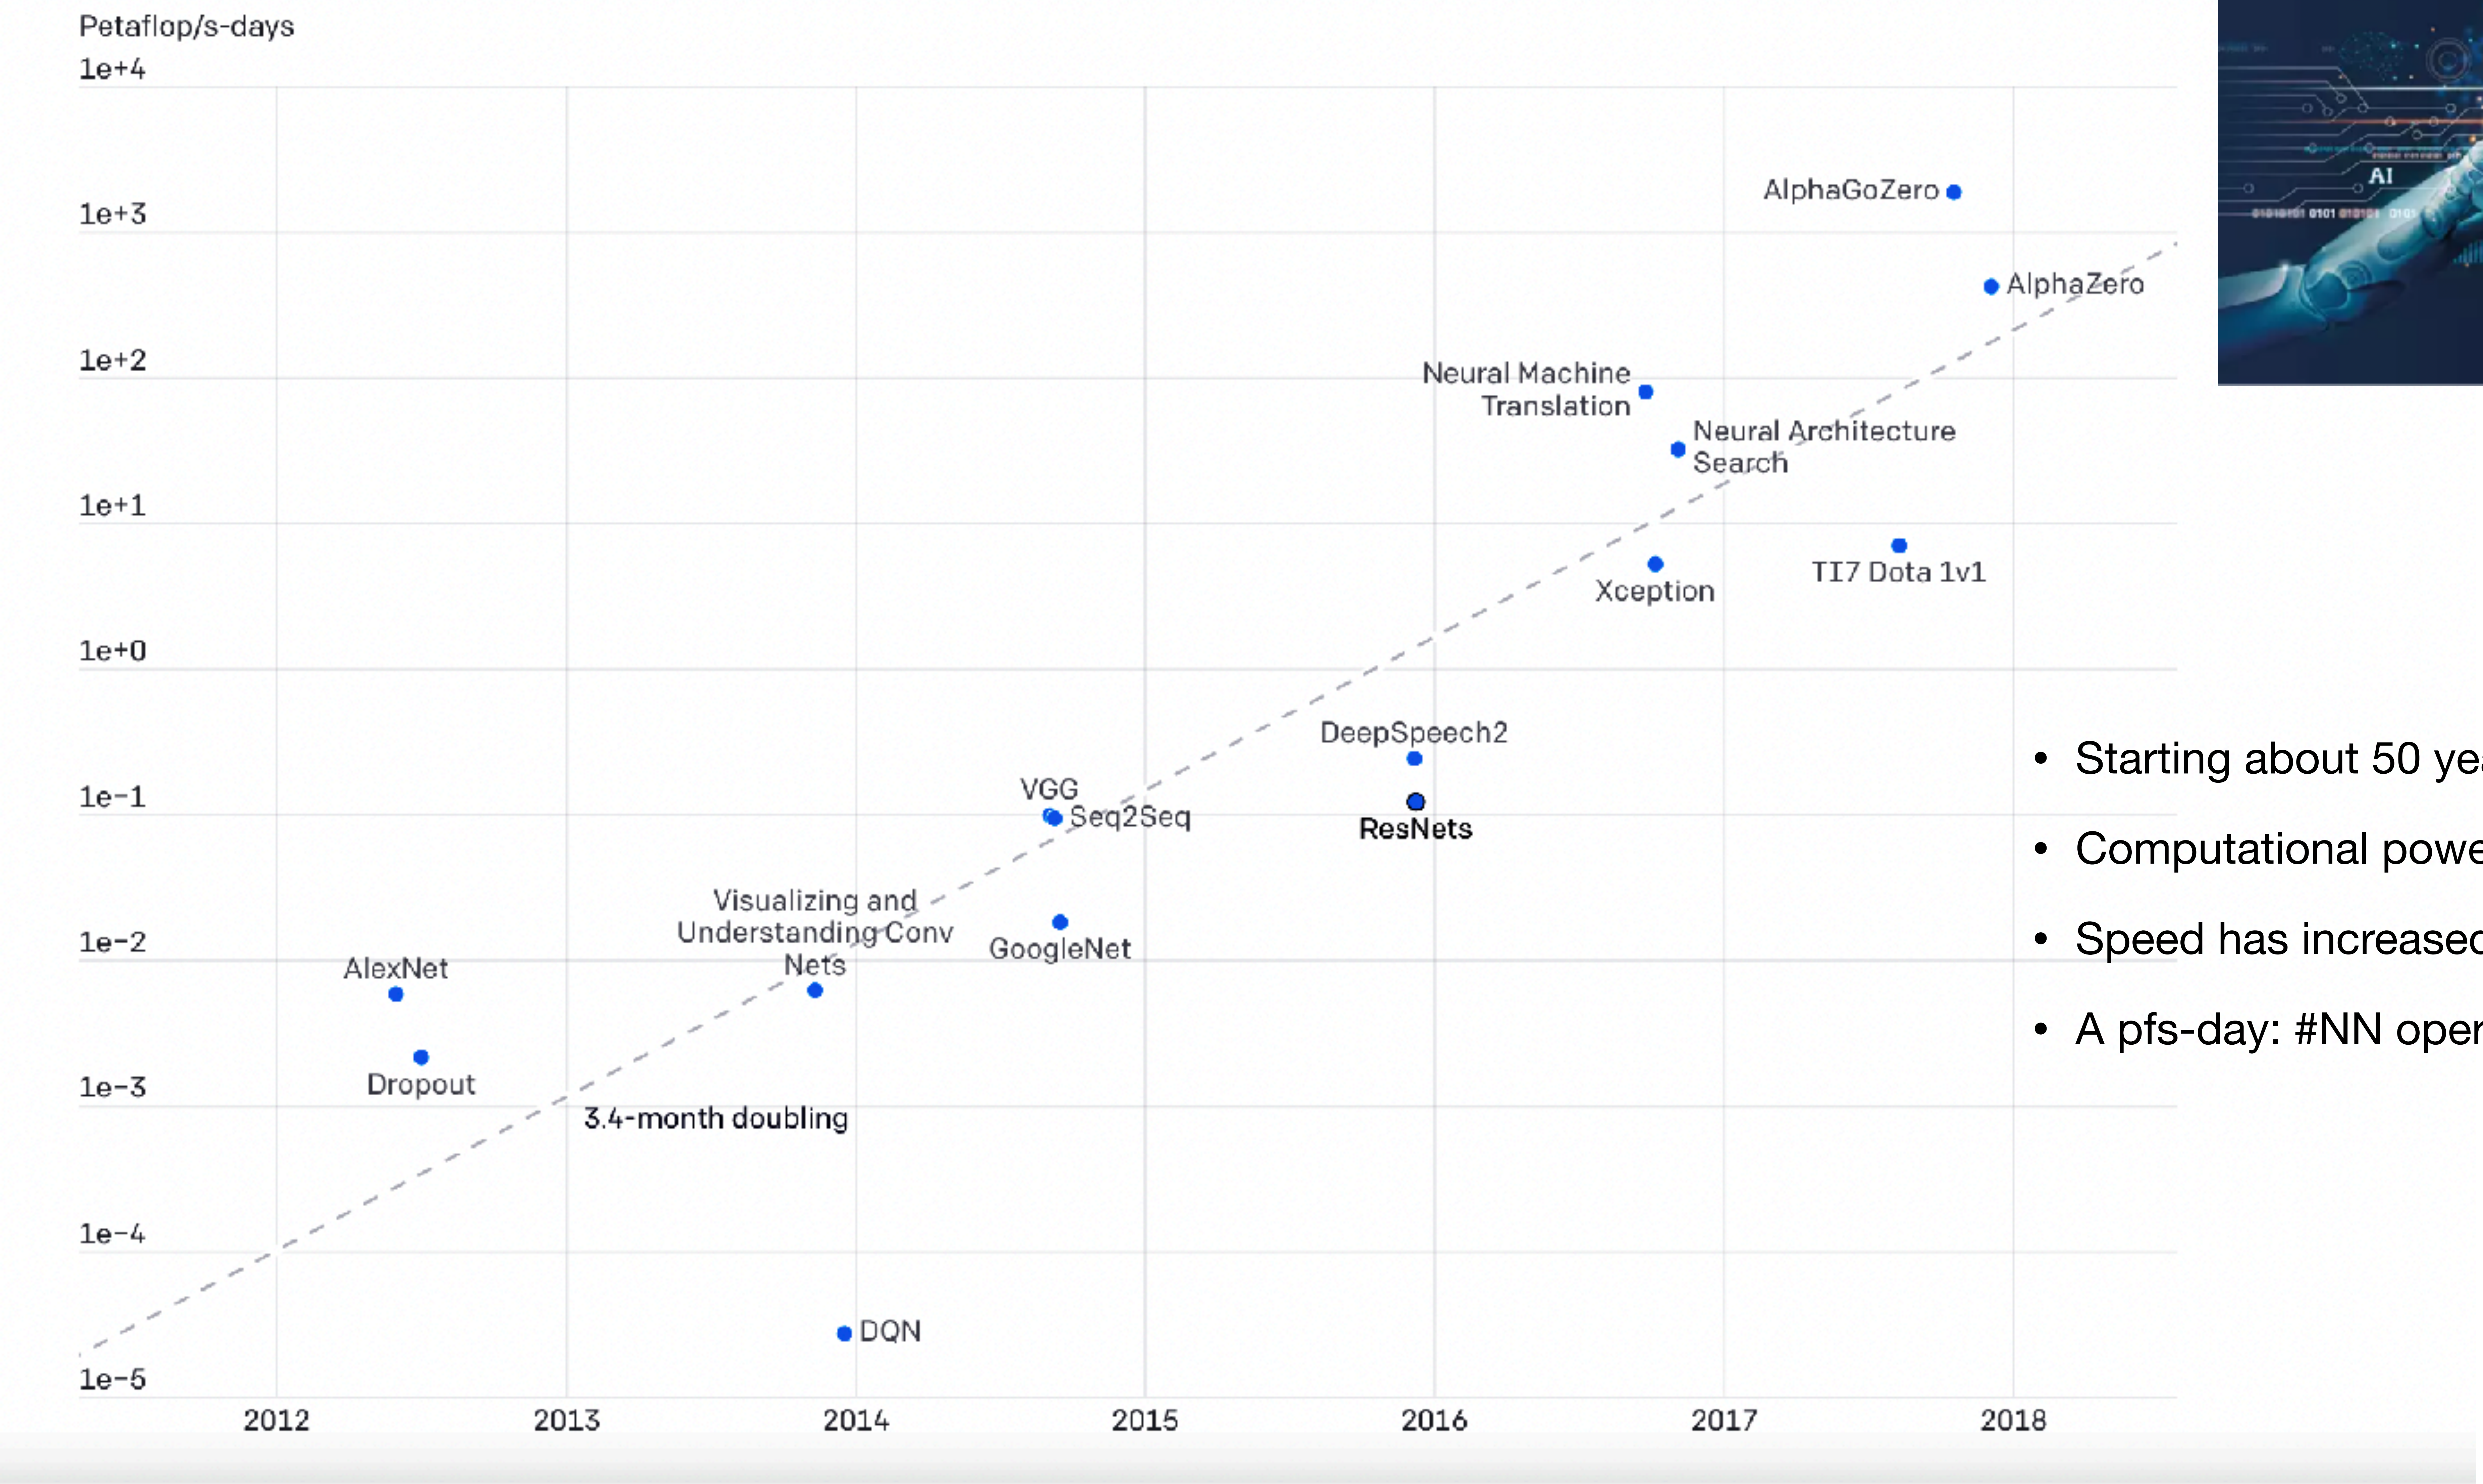

2045

- Starting about 50 years after the perceptron
- Computational power doubling 100 days
- Speed has increased 300K from ‘12
- A pfs-day: #NN operations/s for 1 day

R. Kurzweil


Dimensionality reduction
Meaningful Compression
Structure Discovery
Big data Visualistaion
Dimensionality Reduction
Feature Elicitation
Recommender Systems
Unsupervised Learning
Clustering
Targetted Marketing
Customer Segmentation
Machine Learning
Image Classification
Customer Retention
Idenity Fraud Detection
Classification
Diagnostics
Supervised Learning
Advertising Popularity Prediction
Weather Forecasting
Regression
Population Growth Prediction
Market Forecasting
Estimating life expectancy
Real-time decisions
Game AI
Reinforcement Learning
Robot Navigation
Skill Acquisition
Learning Tasks
cluster
1
2
3
CLASSIFICATION
REGRESSION
Environment
Reward
Best Action
State
Selection of Algorithm
Agent

Basic DNN Concepts

That are changing the way we do physics…

- Extract more information with more inputs (neurons, layers, variables, data)
- Lightning fast inference (specialized hardware to parallelize)
- Ultimately adaptive and accurate (Universal Approximation Theorem)

Machine Learning

In Nuclear and Particle Physics

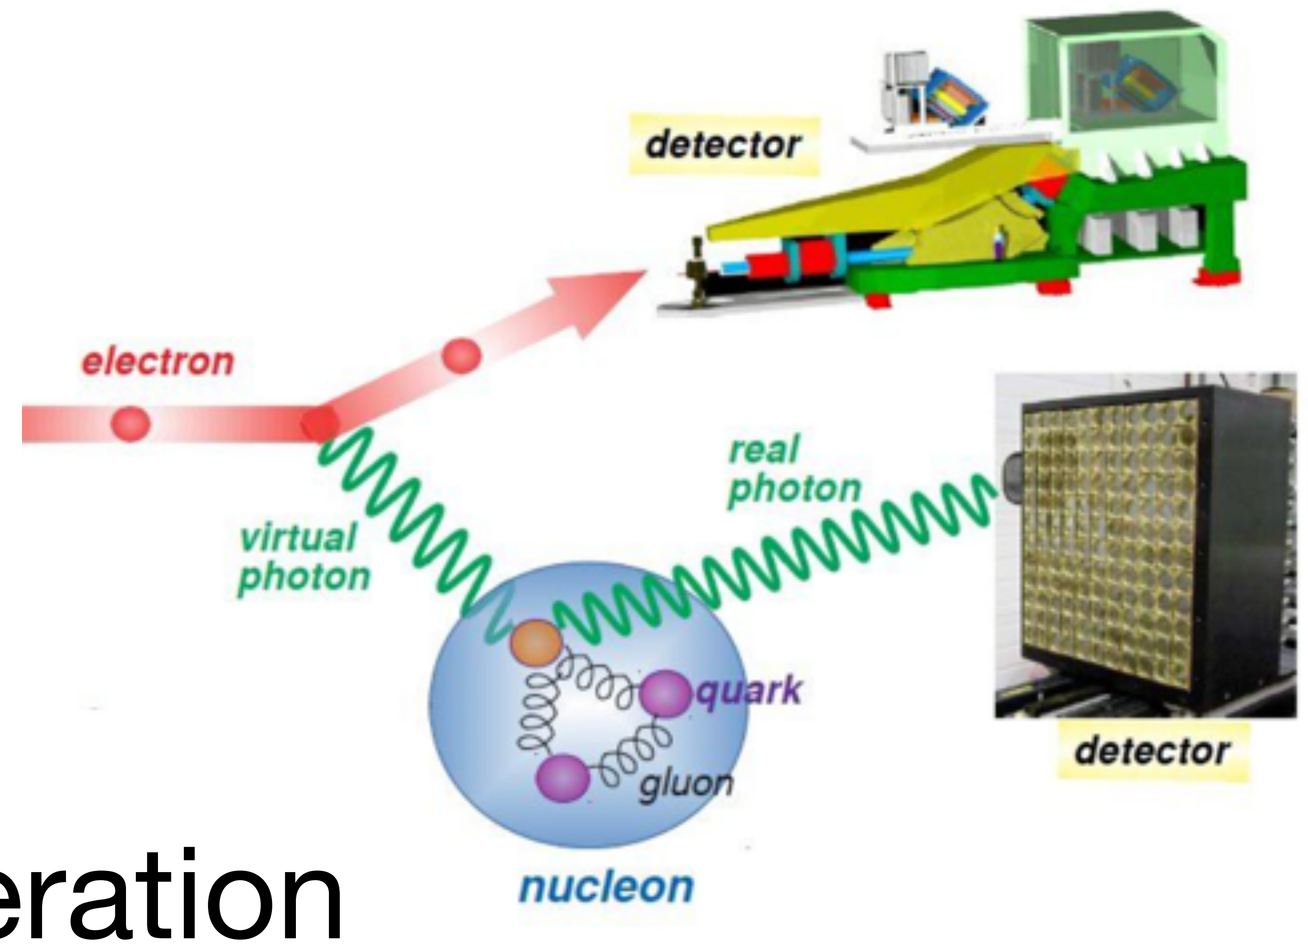

- Event-level Classification
- Trigger and Pattern Recognition
- Tracking/Event Reconstruction
- Cluster Reconstruction in Calorimeters
- Regression of detector drifts
- Phenomenological Extraction

- Simulation acceleration
- Instrumentation Automation
- Detector/cryostat Design
- Online Monitoring/Optimization
- Detector Readout Optimization
- Modeling *

Configure Experiments for DNN optimization
One project at a time

- Detector characterization and performance tuning with DNNs
  - calibration, efficiencies, acceptance
- Raw data (ACD/TDC) straight to Physics
  - Suggests the concept of Intelligent Detectors
  - Optimization of beam time
- Inference using GPU is fast
  - Keep up with data rates
  - Specialize I/O pipeline

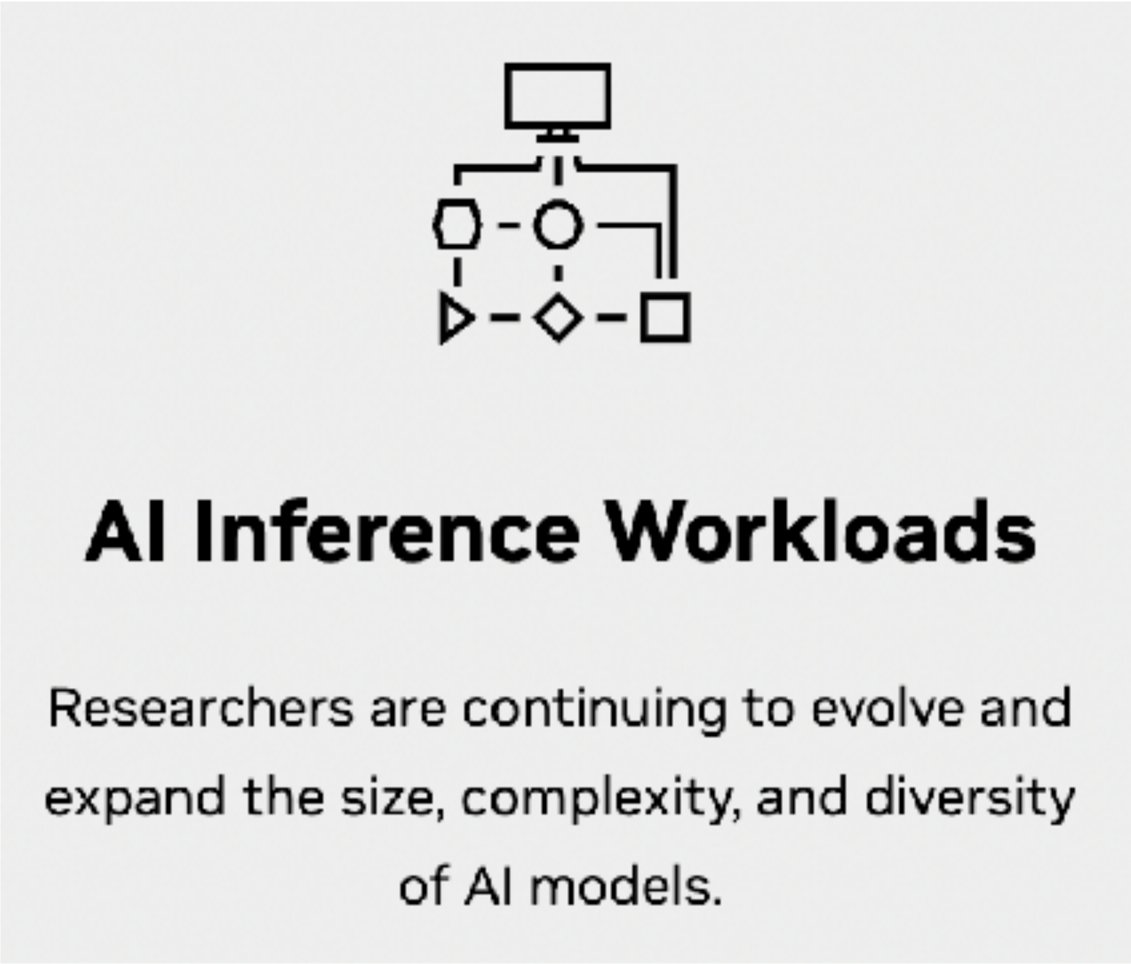

Rethinking the Way We do Research

Getting the most from computing means *Redesigning Everything*

- Raw information is more advantageous for DNN analysis
  - Need to start at the DAQ and tracking level
  - Very fast online monitoring can now perform full chain analysis
  - Anomaly detection on incoming data
  - Real time systematic quantification

Some Examples

In Polarized Targets

- Traveling-wave electron linac
- Irradiated to $10^{17}\ e^{-}$/cm^2
- 14 GeV 10 μA under Liquid Argon (~87 K)
- Proton knocked out to from free radicals
- Also form color centers
- Material color is correlated to the dose
- Optimized for field and temperature

Irradiation Performed at NIST (MIRF Accelerator)

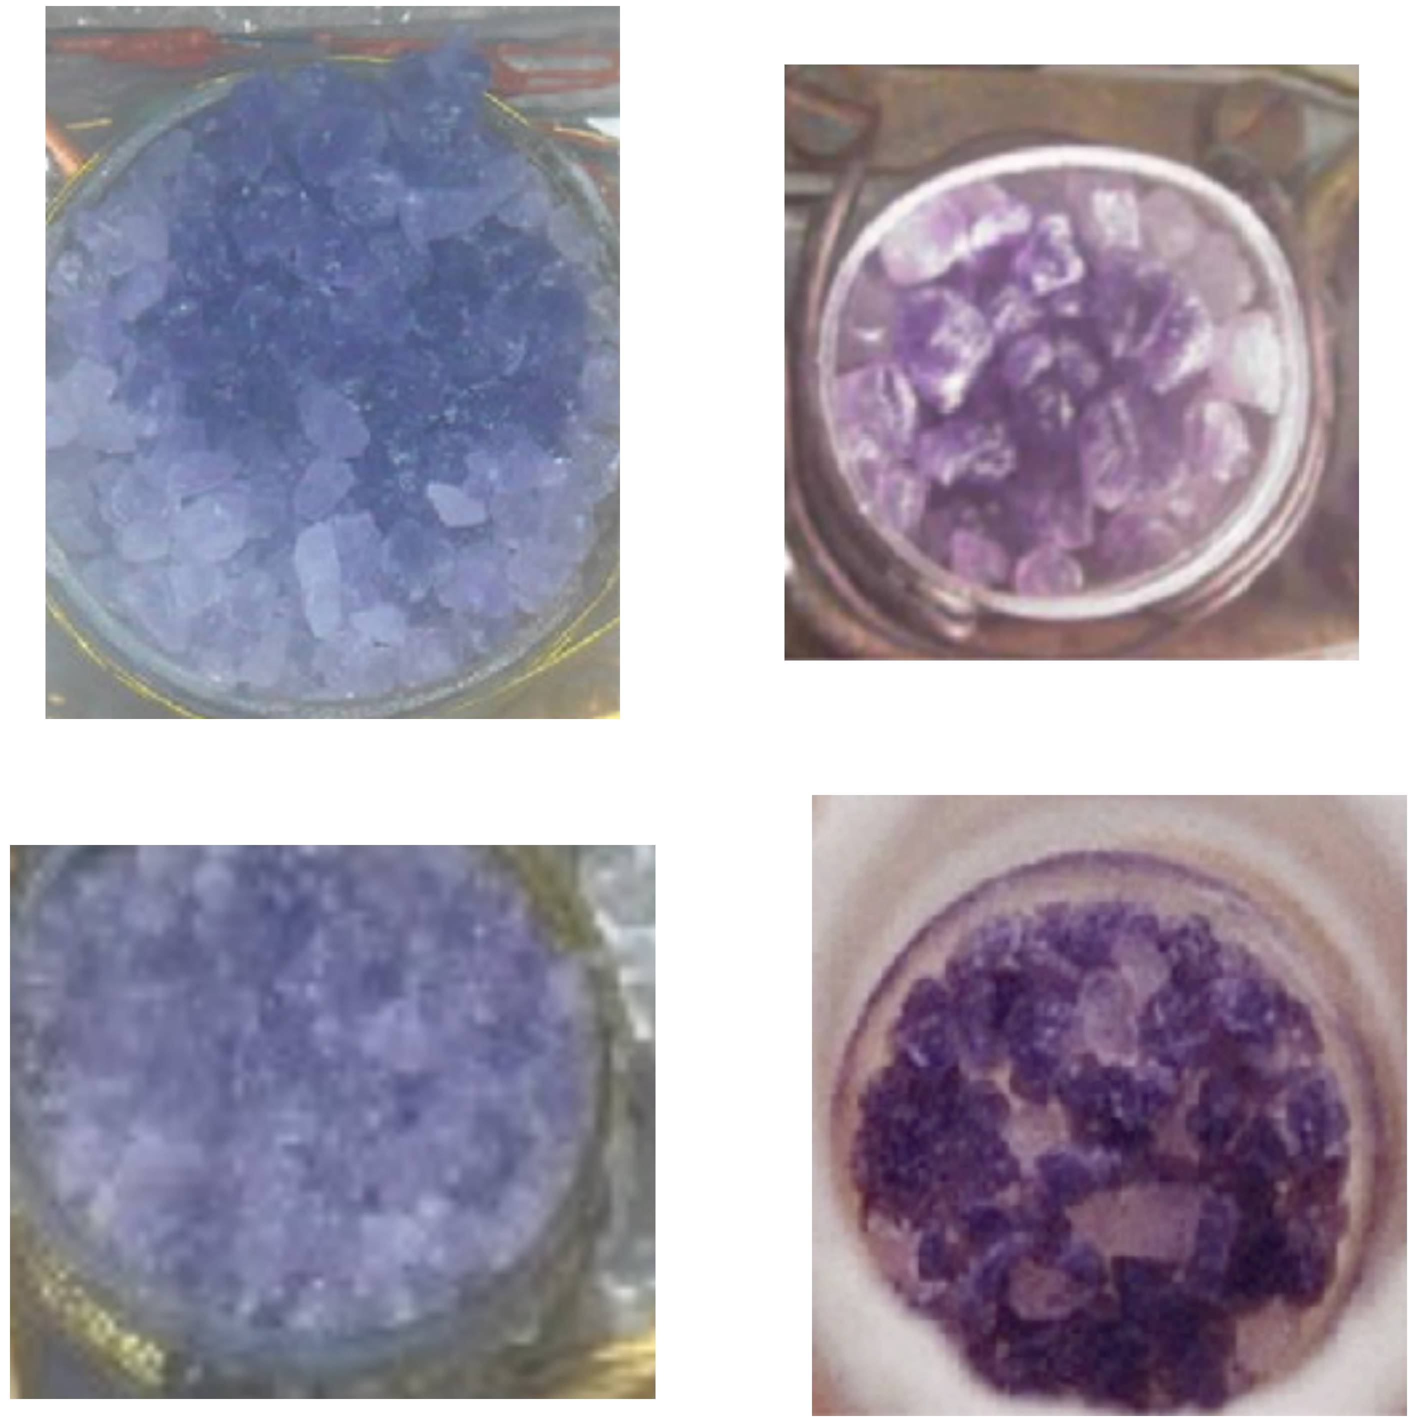

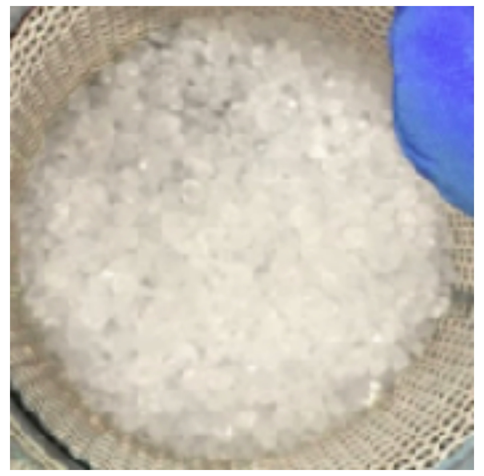

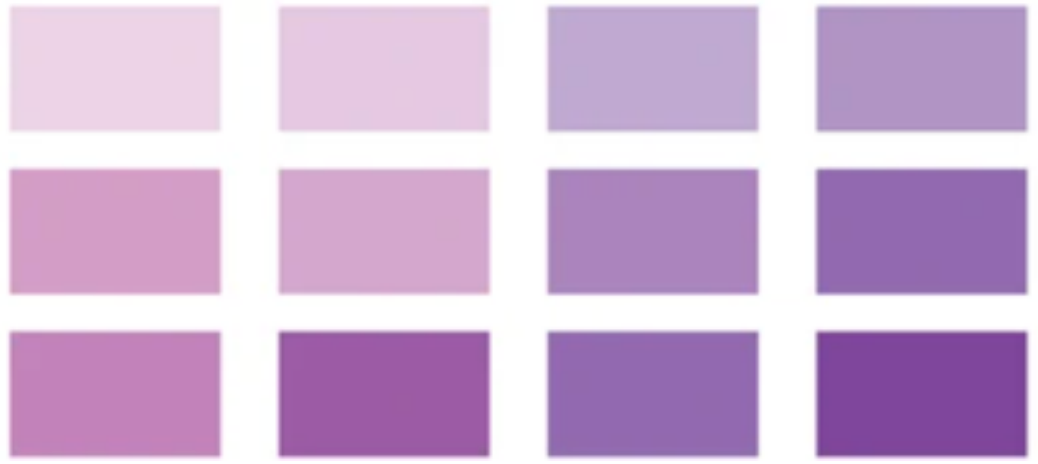

Material Photo

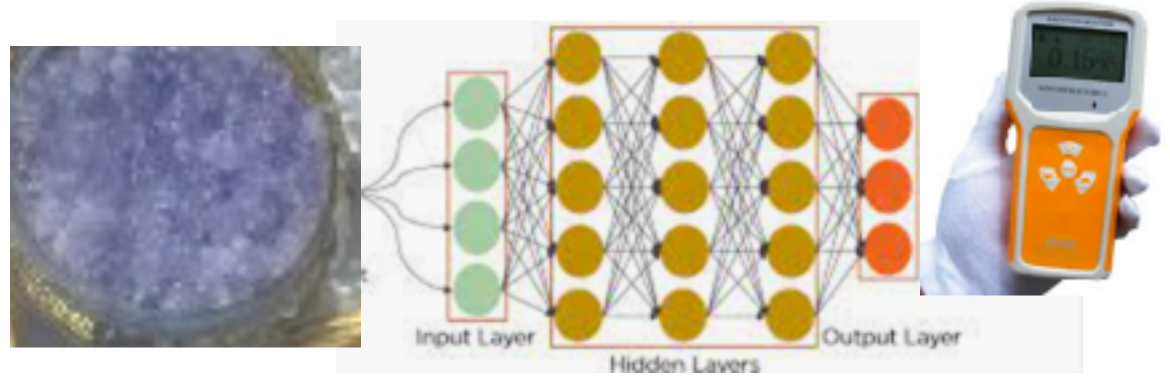

Accumulated dose

AI in Online Monitoring

Fast Reconstruction

$$A = \frac{N_L - N_R}{N_L + N_R}$$

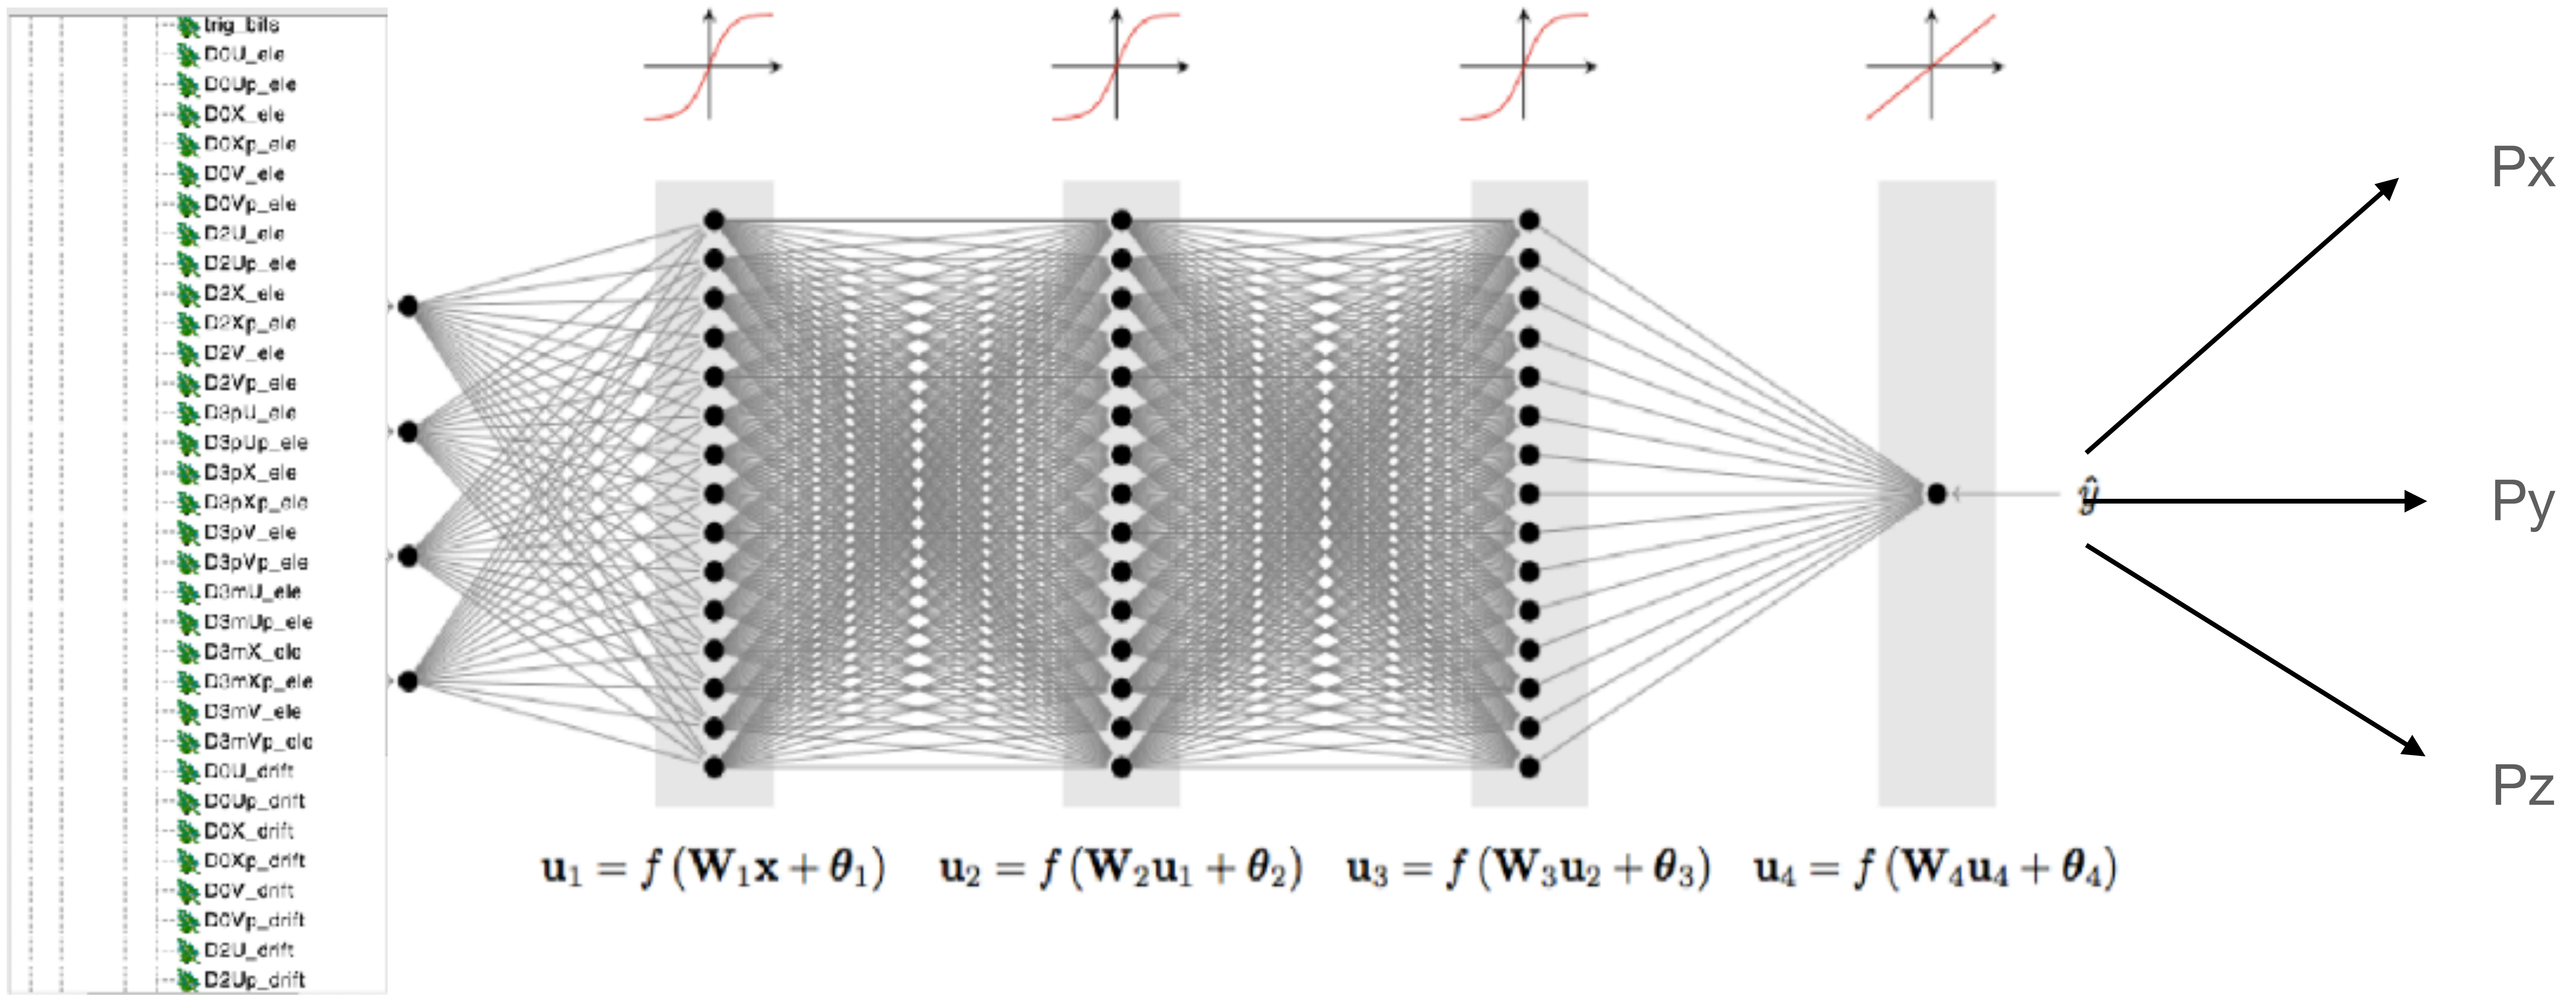

- Detector element failure
- Target-beam miss alignment
- Cell twisting
- Faulty position
- Target cell damage

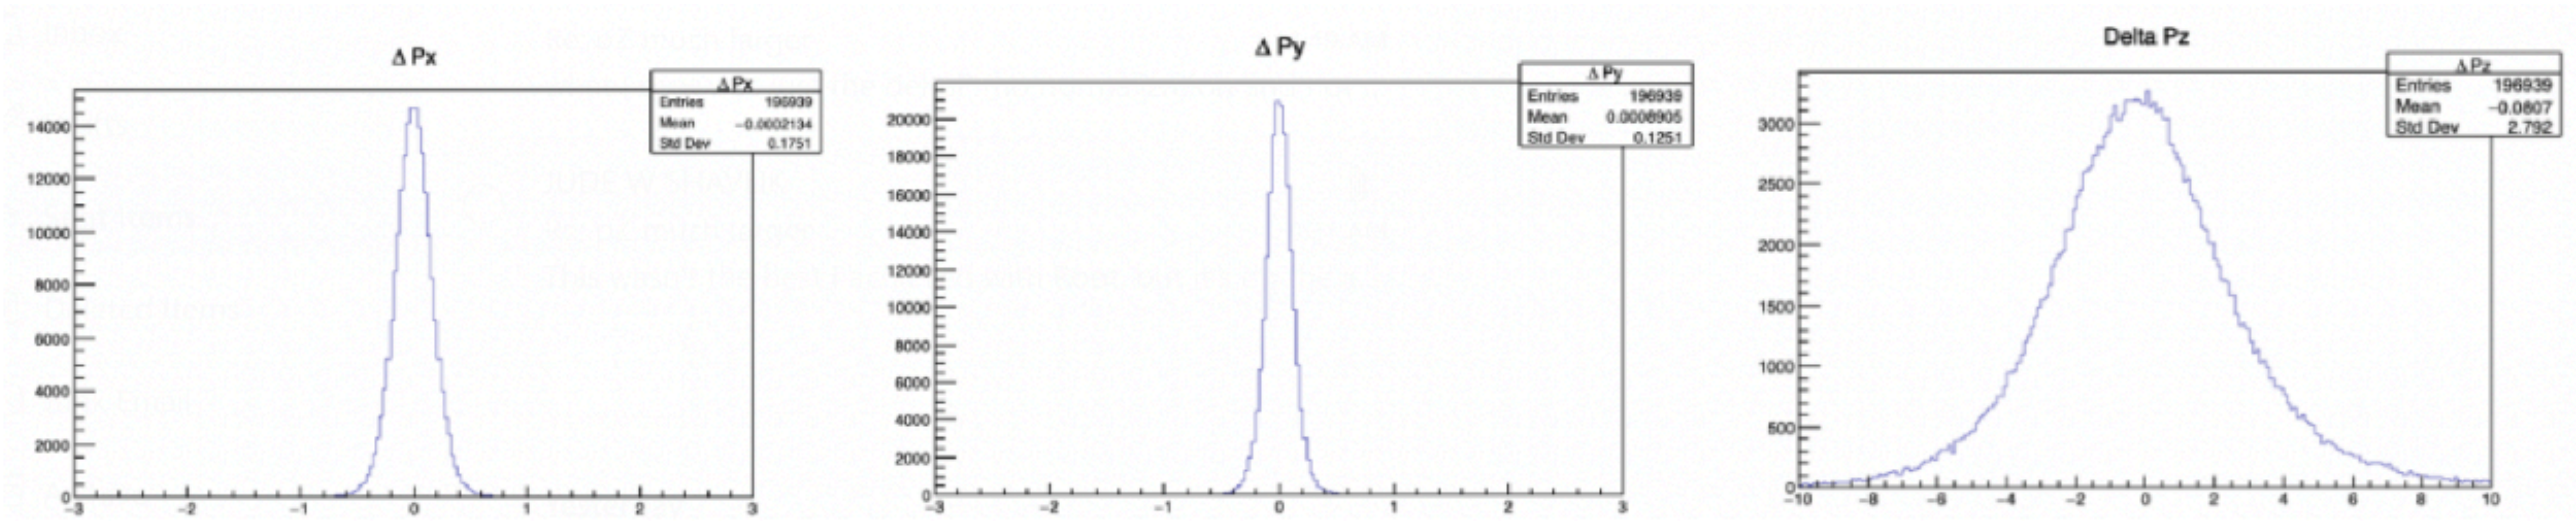

Arthur Conover

AI in Online Reconstruction

Fast Reconstruction

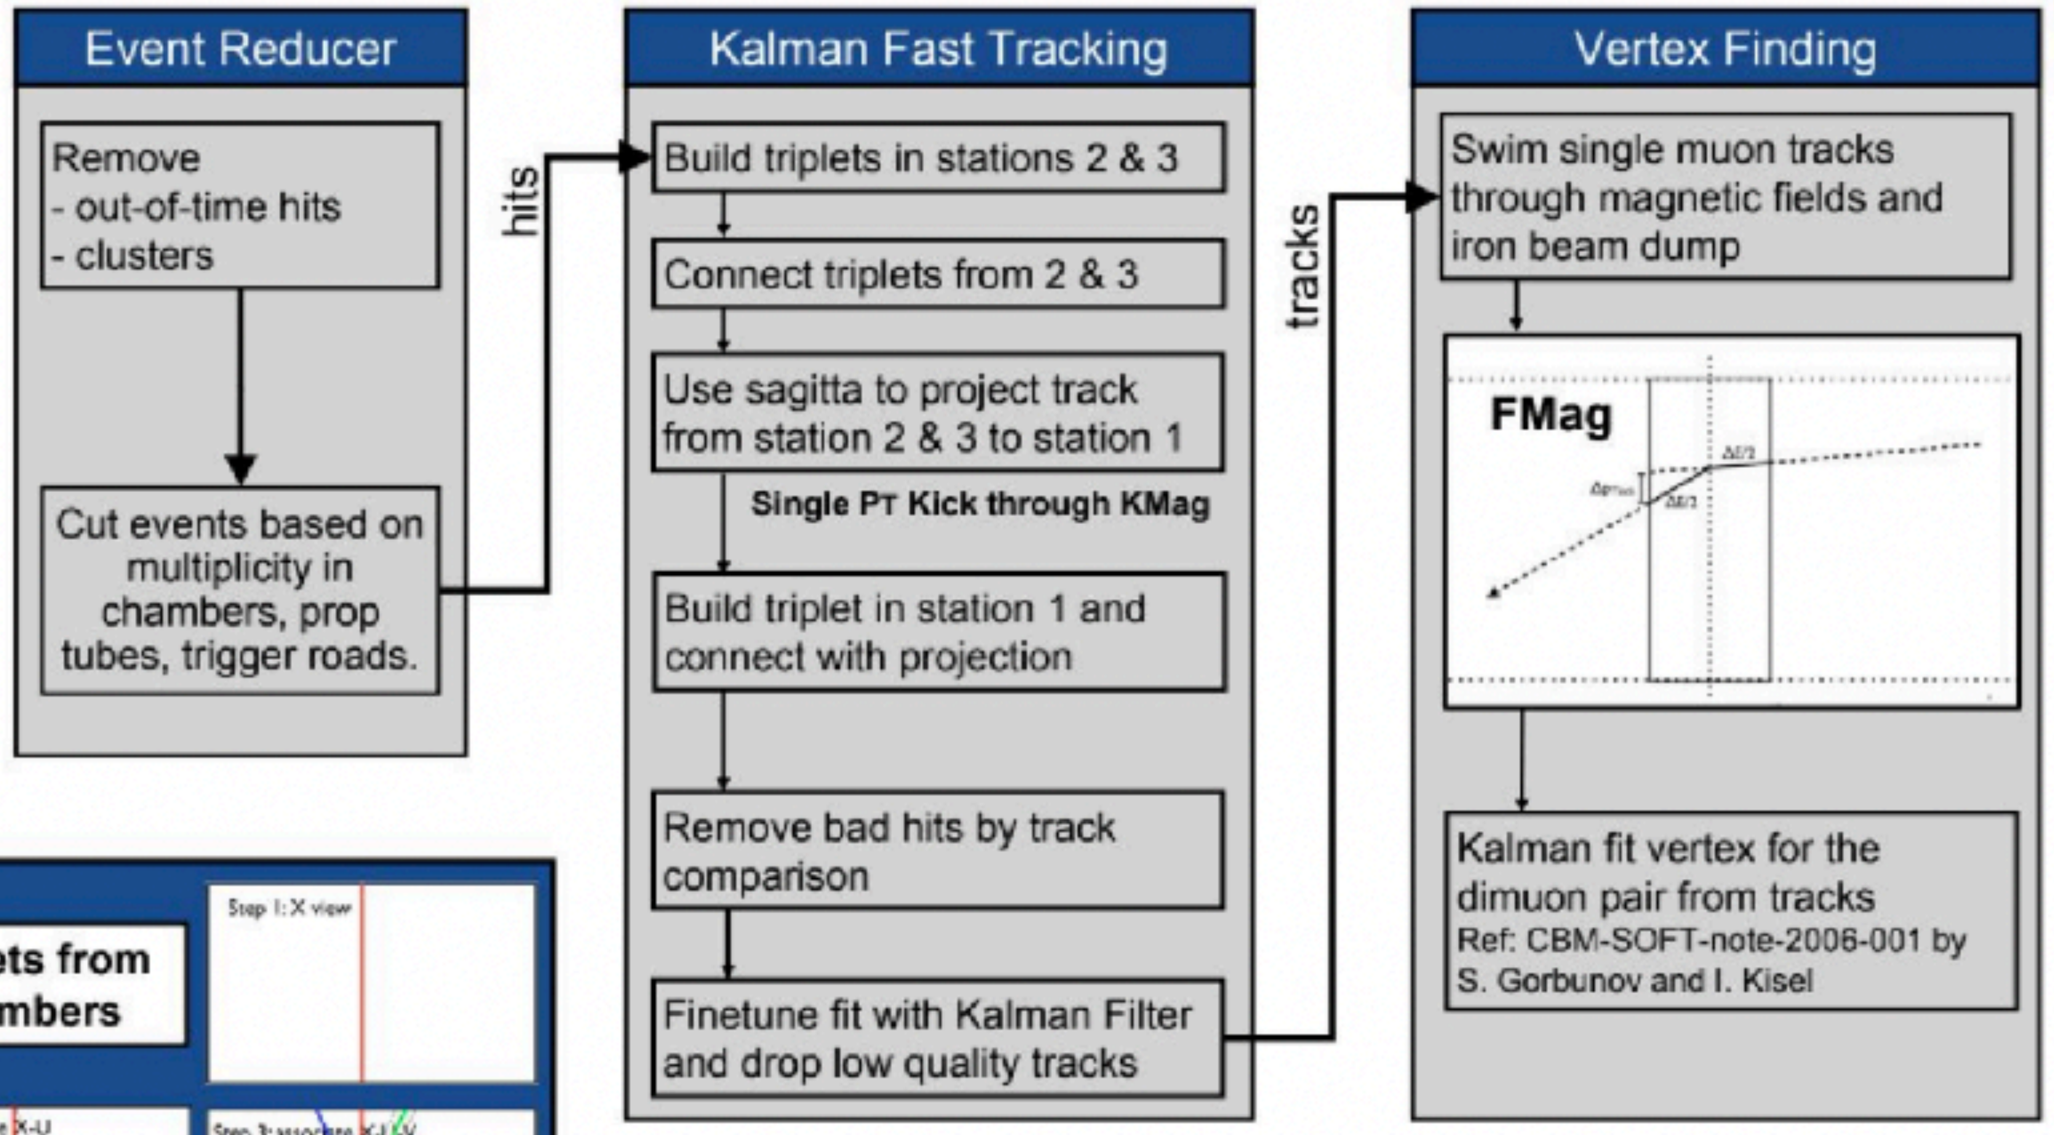

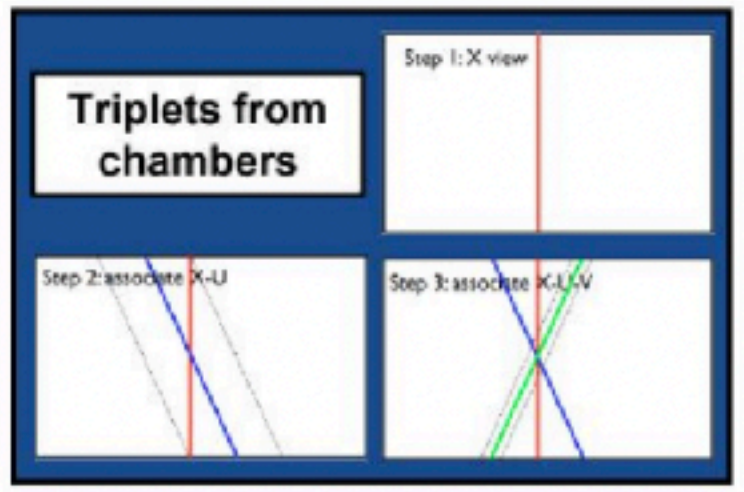

Noah Wurfel, University of Michigan

1 CPU: 110 minutes for 40K

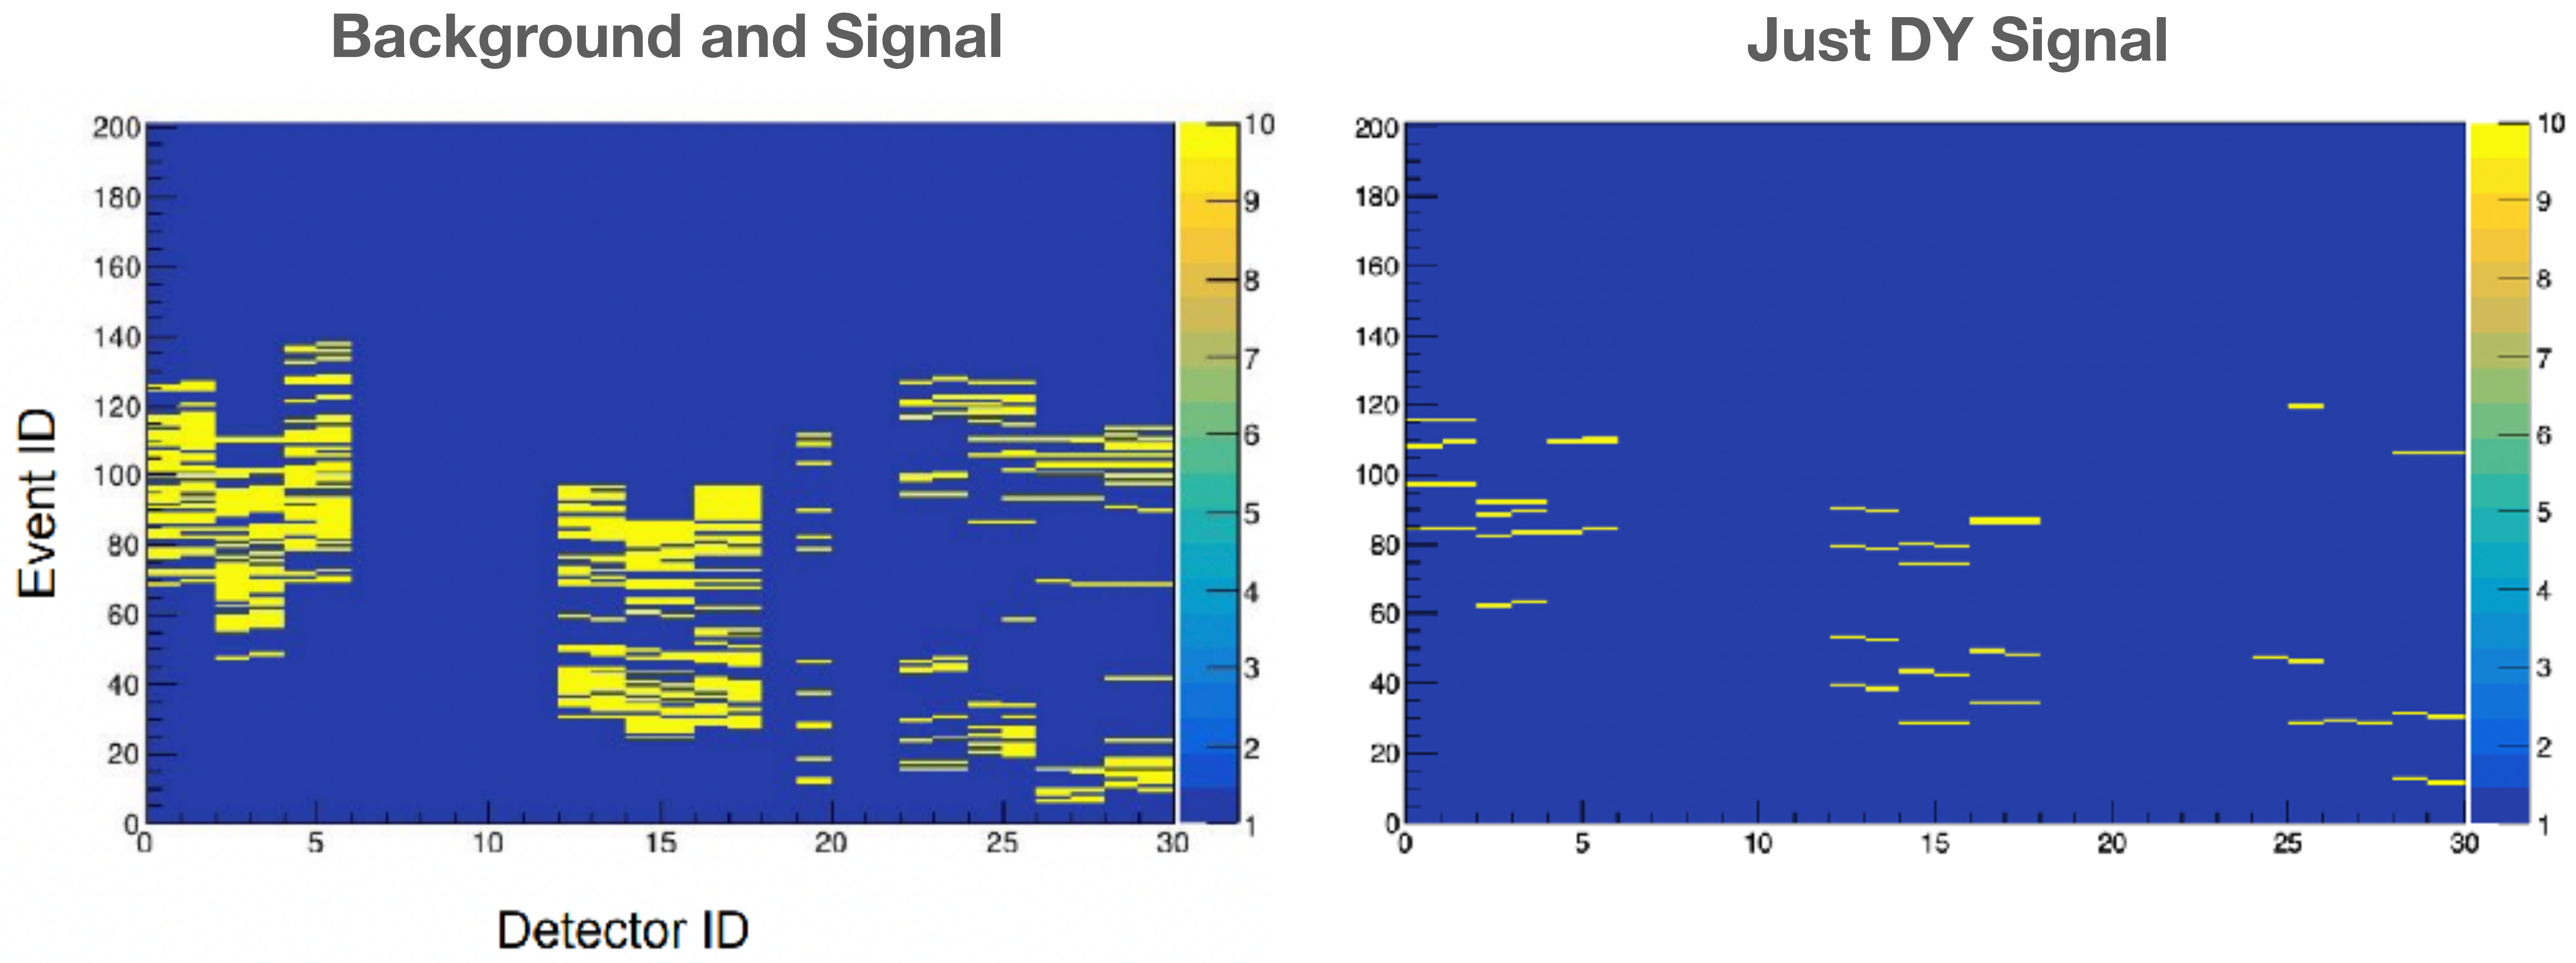

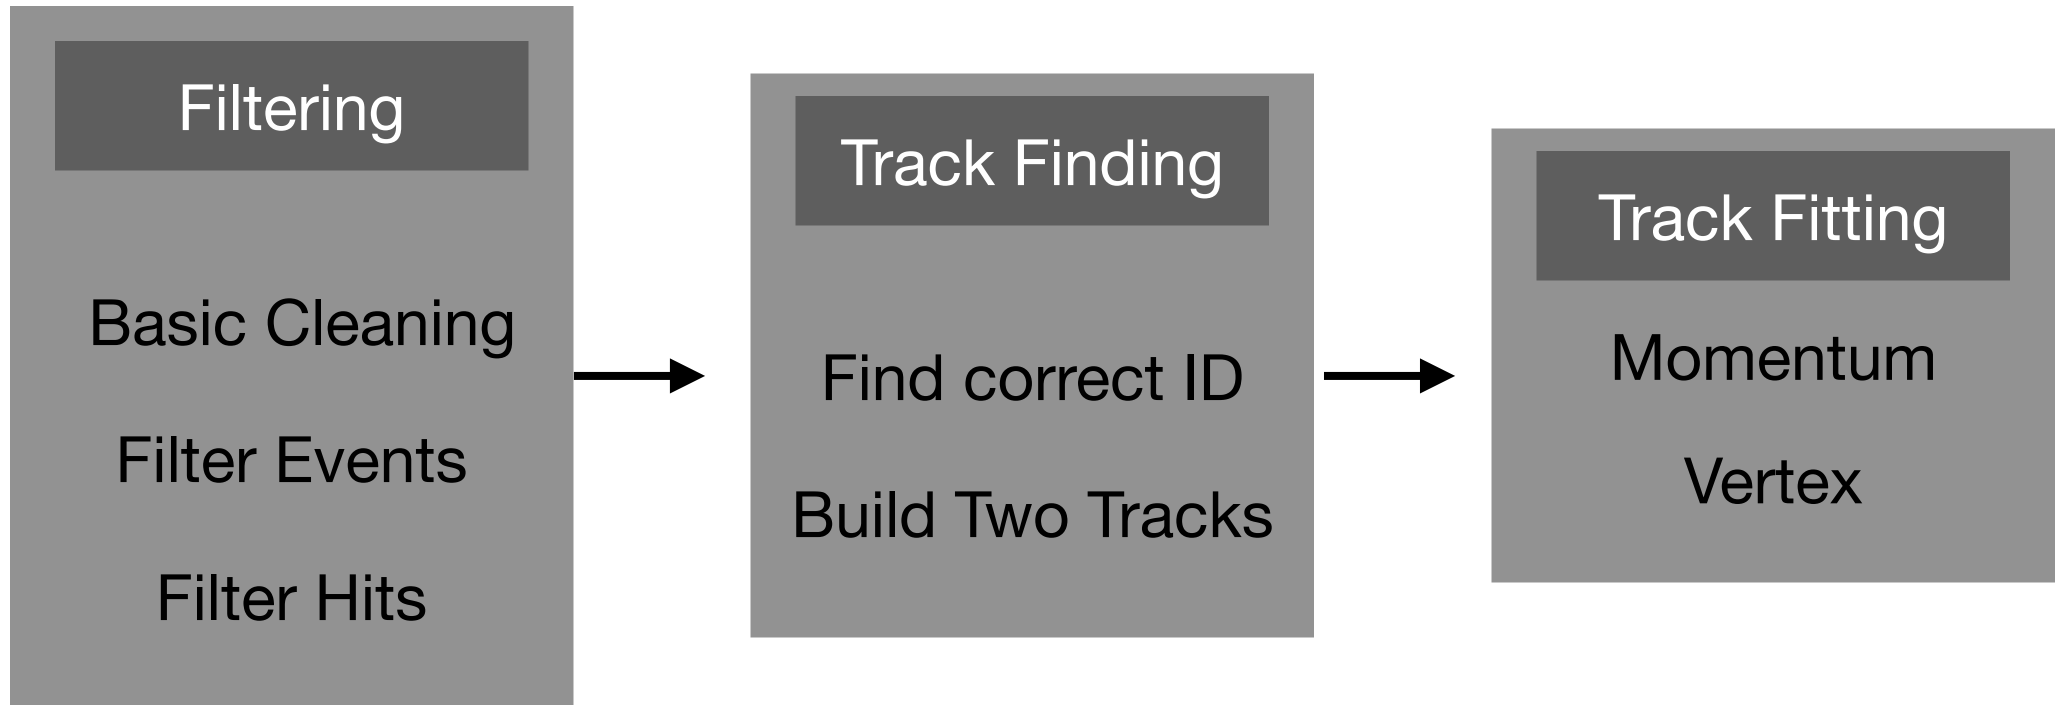

1 CPU: 48 seconds for 40K

Parton Distribution Functions

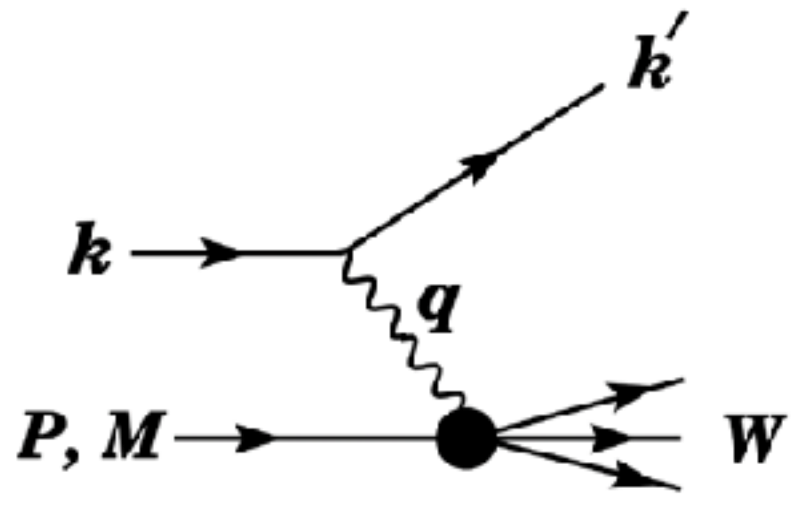

HERA and LHC programs

RHIC program

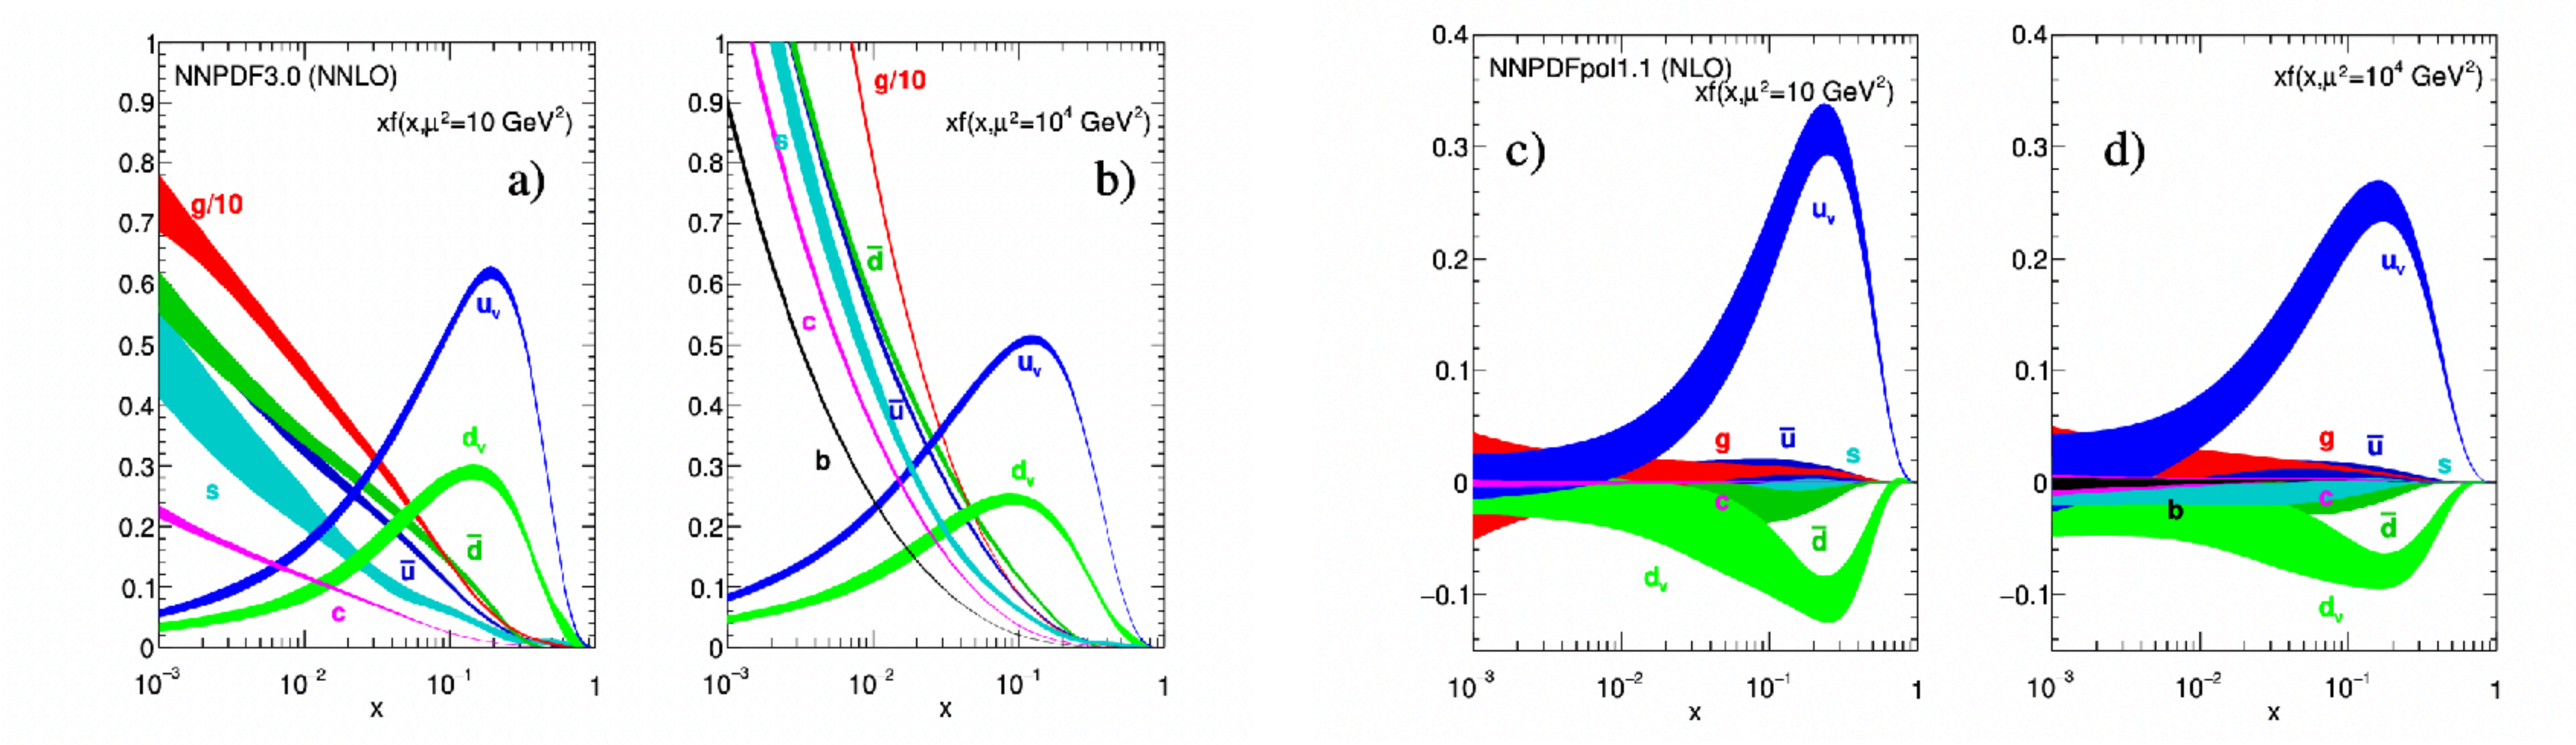

E. R. Nocera et al. (NNPDF), Nucl. Phys. B887, 276 (2014), [arXiv:1406.5539]

- In DIS resolving power goes like $\hbar/q$, $q = 100$ GeV → 0.002 fm
- Consider the parts in a collinear with the nucleon momentum so each parton caries x
- PDF represent probability density (number densities as they are normalized to the number of partons)

Some Issues in Spin

In Phenomenology (Understanding Femtoscale Dynamics)

- Inverse Problem: Determine definitive measures of proton structures using experimental information, Lattice Calculations, and Phenomenology
  - Extraction of GPDs while eliminating the reliance of model fits
  - Extraction of TMDs without assuming a Gaussian factorized form
- Curse of Dimensionality: Understanding the Mother Function (Wigner?) in terms of processes and physical observables (interpretation yields inherent sparsity)
  - How can we impose constraints at the higher-level to interpret dynamics and geometry
  - How do we best obtain information from experiments that gets us the farthest

Candidate Mother Distribution

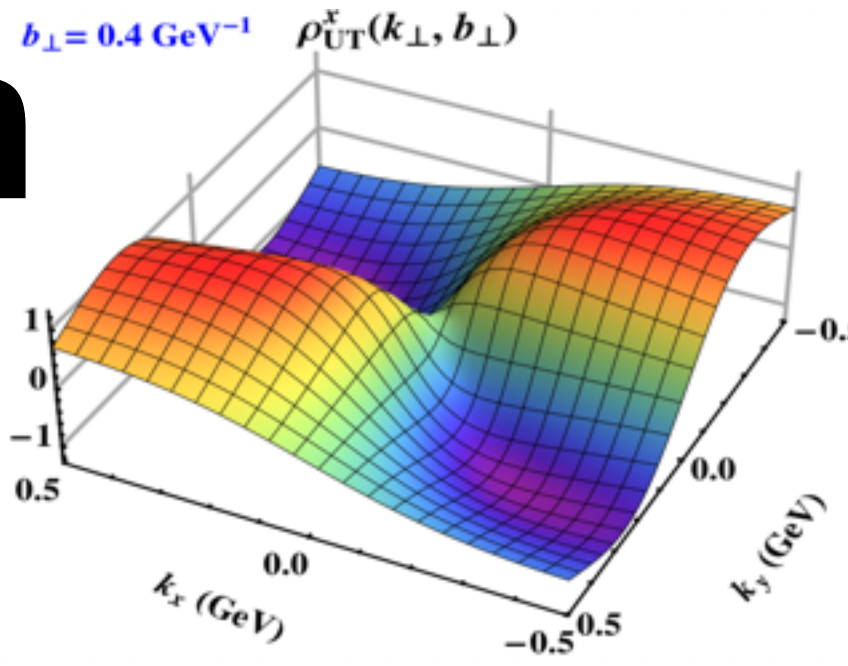

More, AM, Nair, PRD 95, 074039 (2017)

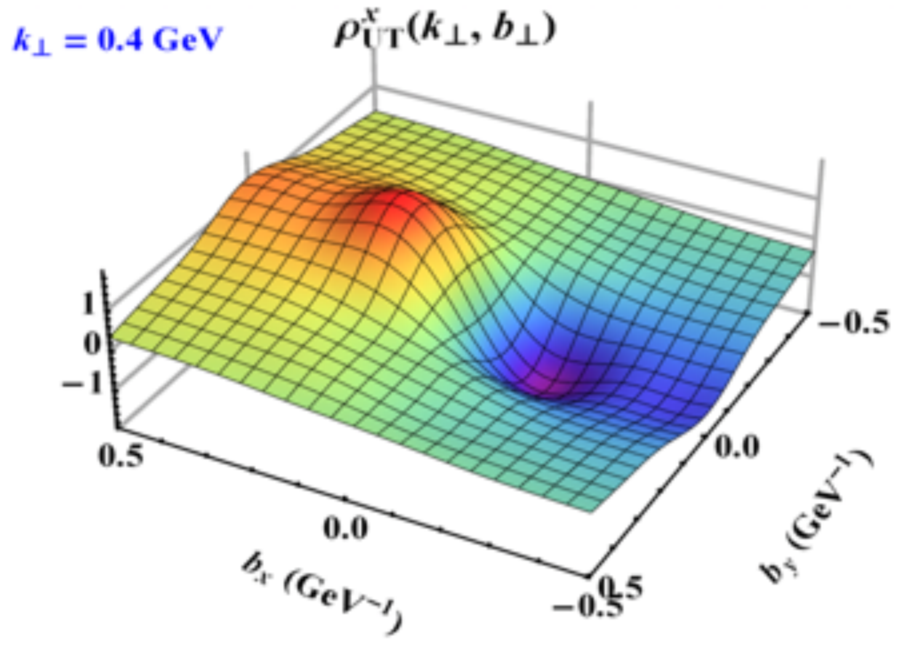

$W(x, b_T, k_T)$
Wigner Distribution

$\int d^2 b_T$ $\quad\quad$ $\int d^2 k_T$

TMDs $f(x, k_T)$ $\quad\quad$ $f(x, b_T)$ Impact Parameter

$\int d^2 k_T$ $\quad\quad$ $\int d^2 b_T$

$f(x)$
PDFs

(Liu and Ma, PRD 91, 034019 (2015))

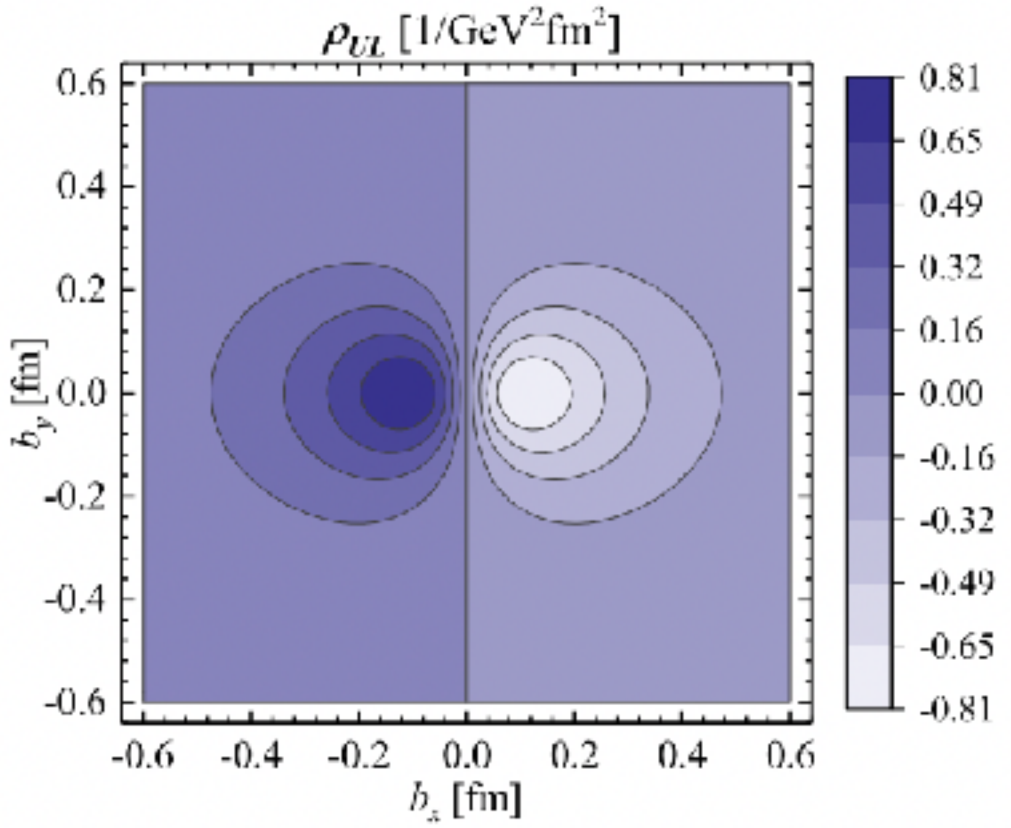

TMDs and Impact Parameters give complementary information about partons and are fundamentally connected to the Wigner Distribution

Lorce and Pasquini, PRD 84, 014015 (2011); Liu and Ma, PRD 91, 034019 (2015)

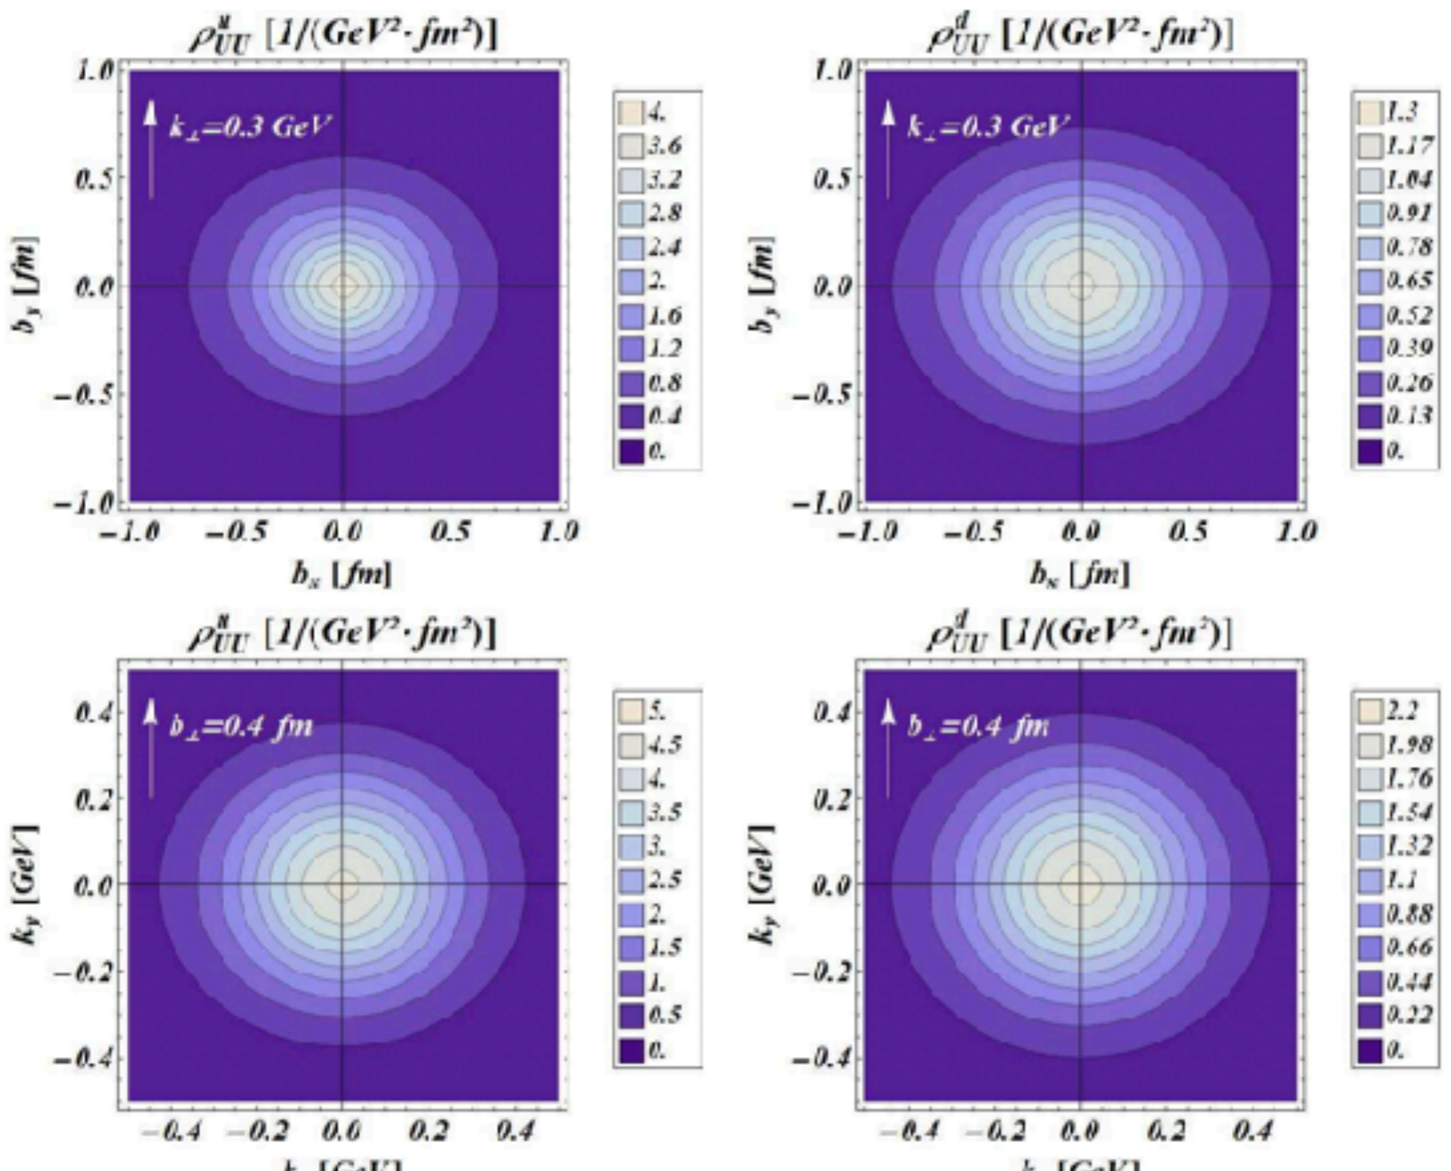

More, AM, Nair, PRD 95, 074039 (2017)

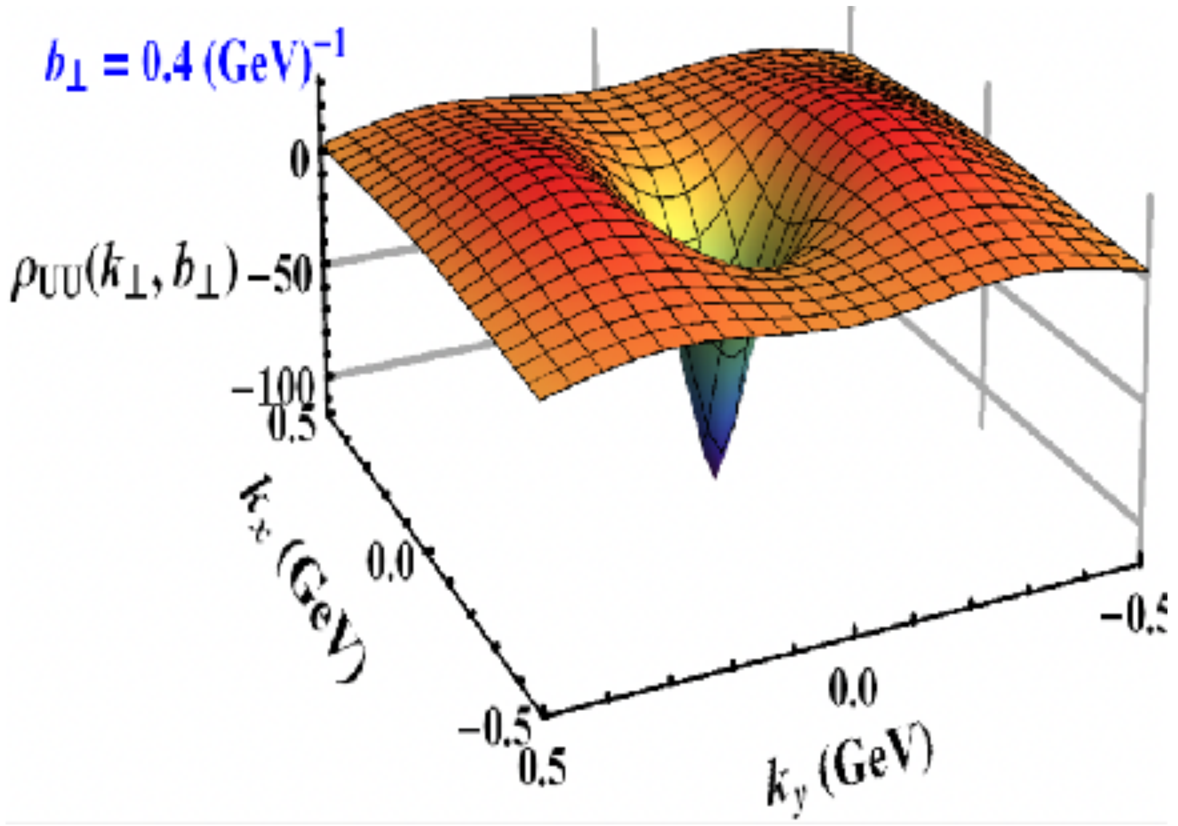

Husimi distributions have a Gaussian regularization factor in the integrand that keeps them positive in the entire range of transverse space coordinate

Zhi-Lei Ma and Zhun Lu
Phys. Rev. D **98**, 054024

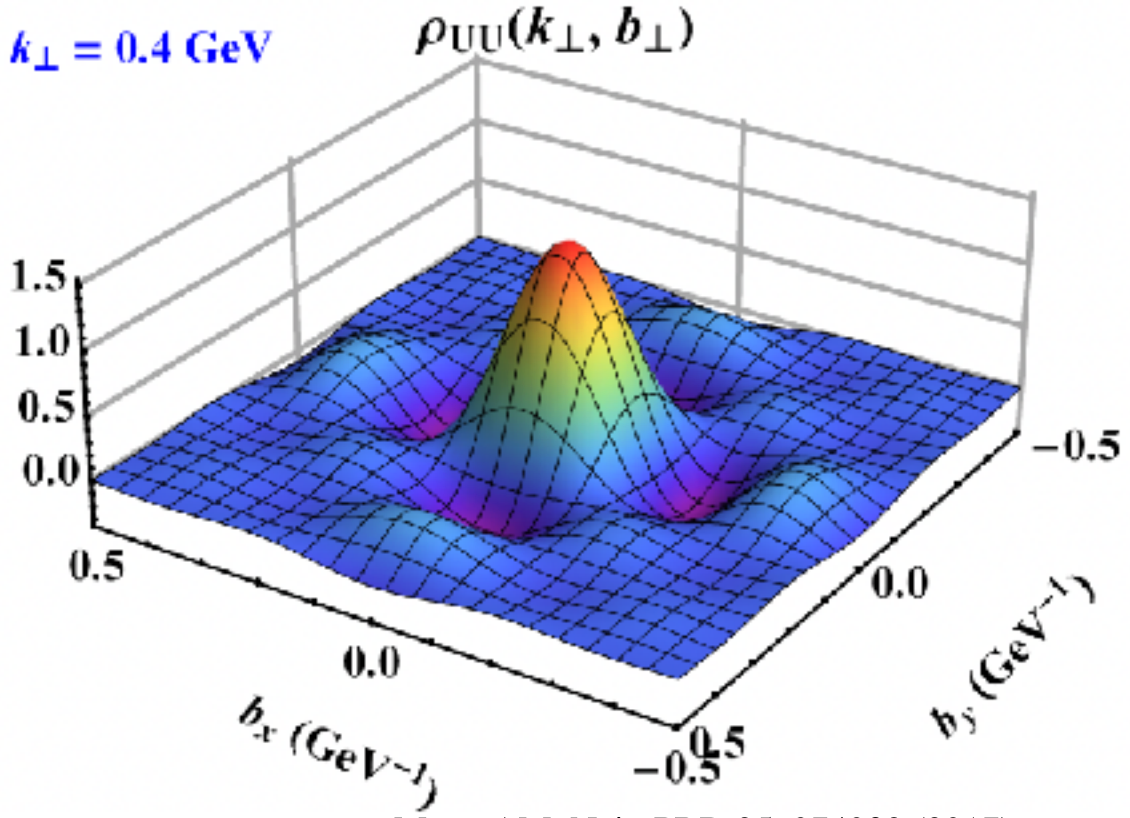

More, AM, Nair, PRD 95, 074039 (2017)

Nucleon Tomography

Or Femtography

Density in the Transverse-momentum plane

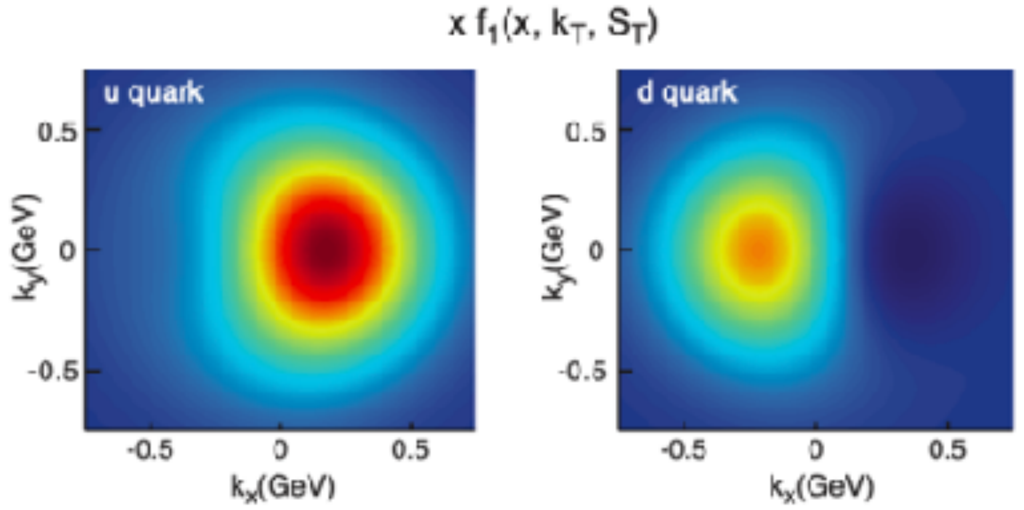

Deep red (blue) are large negative (positive)

quark polarization

nucleon polarization

TMD	U	L	T
U	f_1		$h_1^\perp$
L		g_{1L}	$h_{1L}^\perp$
T	$f_{1T}^\perp$	g_{1T}	$h_1,\ h_{1T}^\perp$

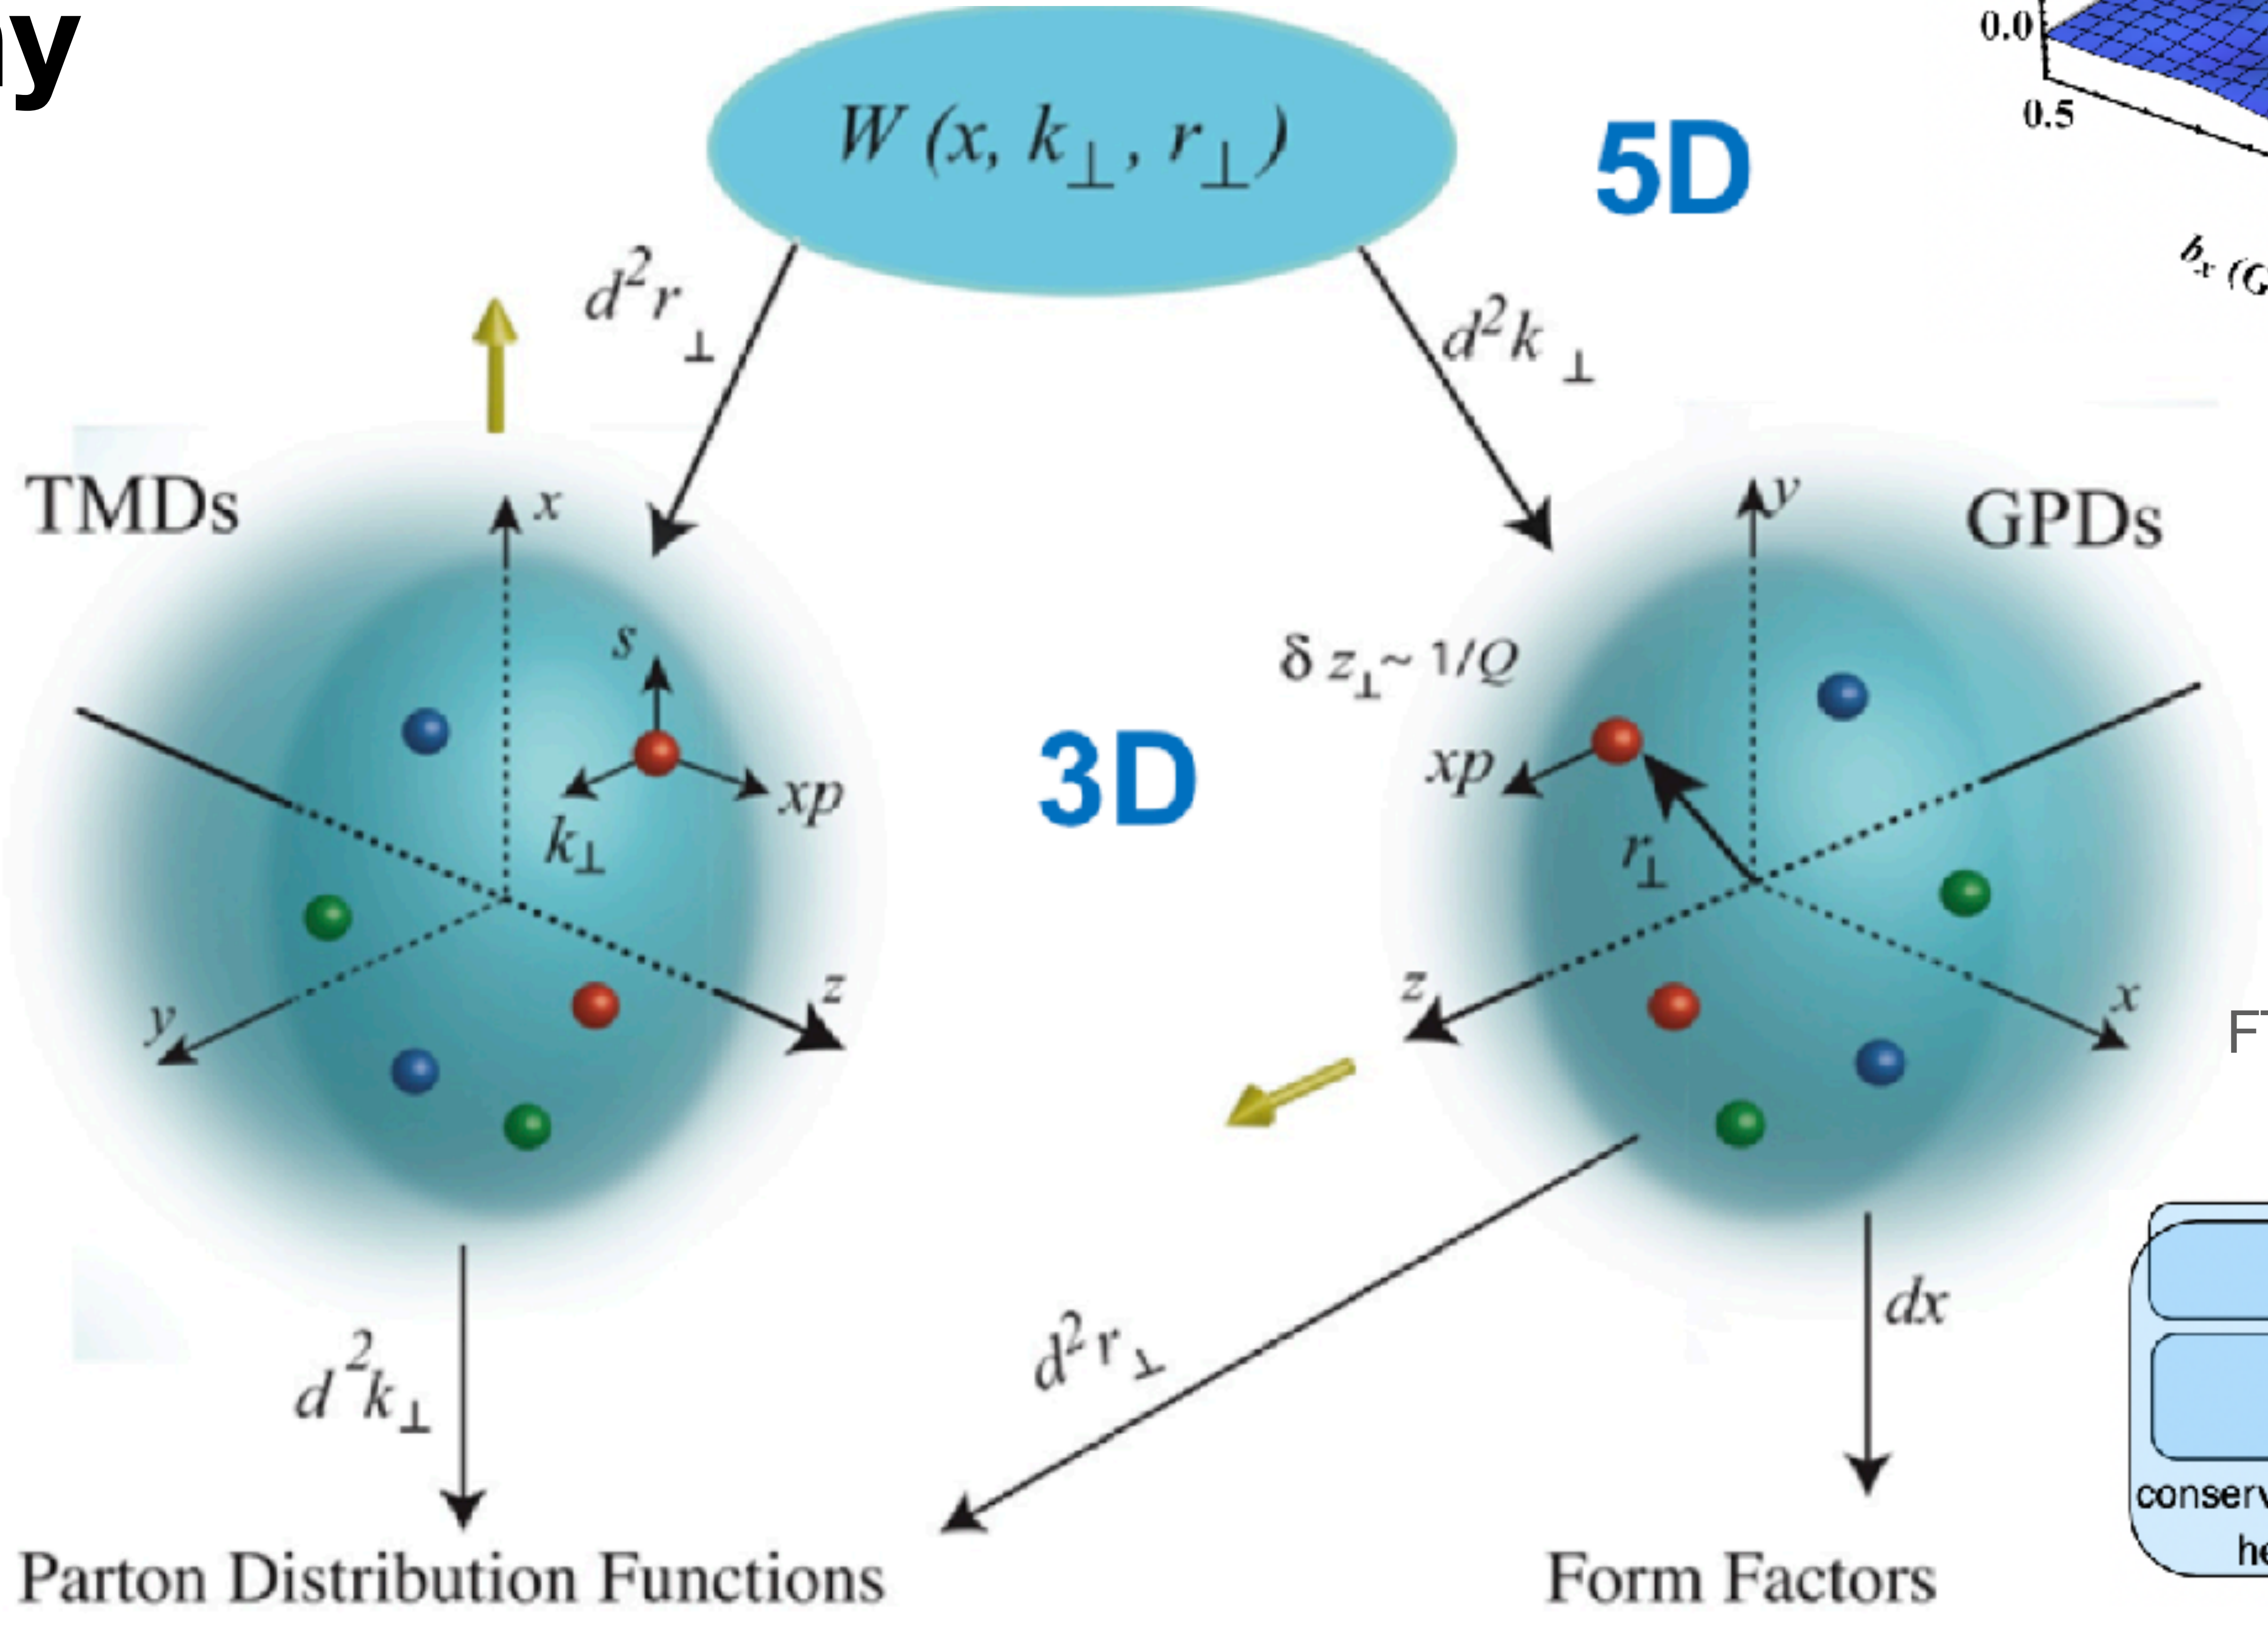

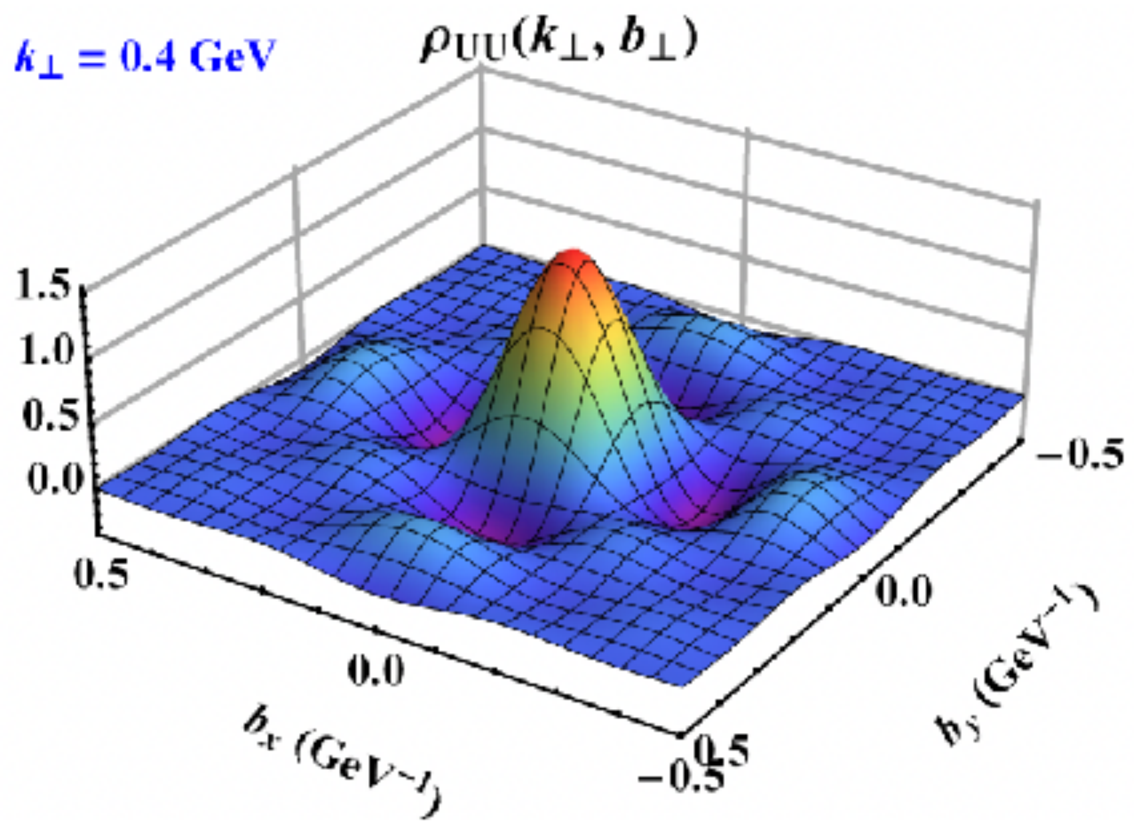

quark polarization

nucleon polarization

GPD	U	L	T
U	H		$\mathcal{E}_T$
L		$\tilde{H}$	$\tilde{E}_T$
T	E	$\tilde{E}$	$H_T,\ \tilde{H}_T$

FT with respect to transverse components leads To spatial distribution of pardons

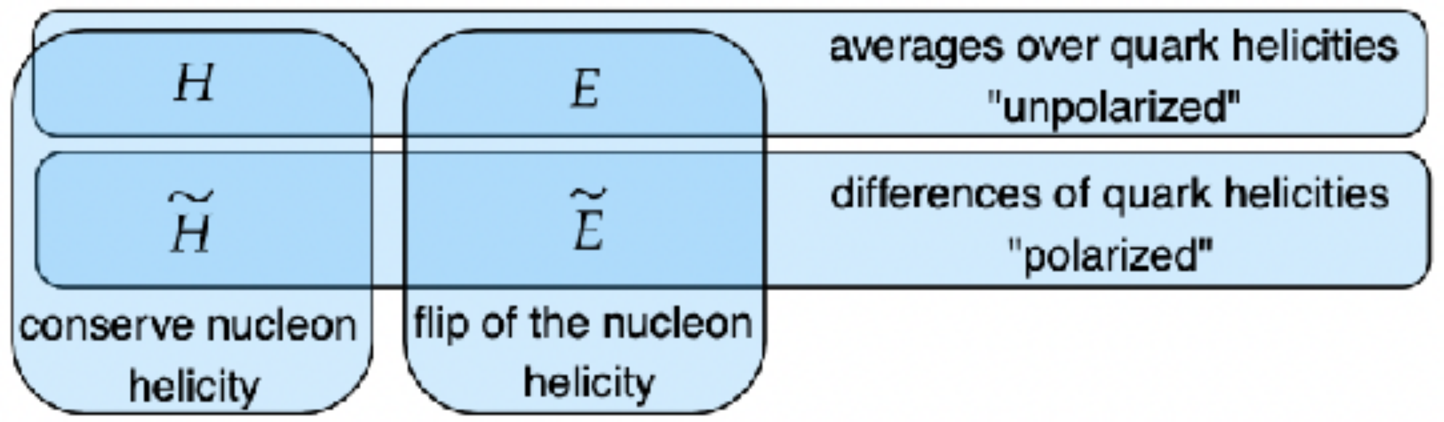

Transverse Momentum Distributions

- Transverse momentum structure
- Confined Motion of Partons

Generalized Parton Distributions

- Spatial Imaging of internal structure
- Form factors and generalized longitudinal structure

TMD Global Fitting

Pseudodata Extraction before Real Data

- Reduce assumptions
- Explore cuts
- Reduce Bias
- Still lots of improvement to go

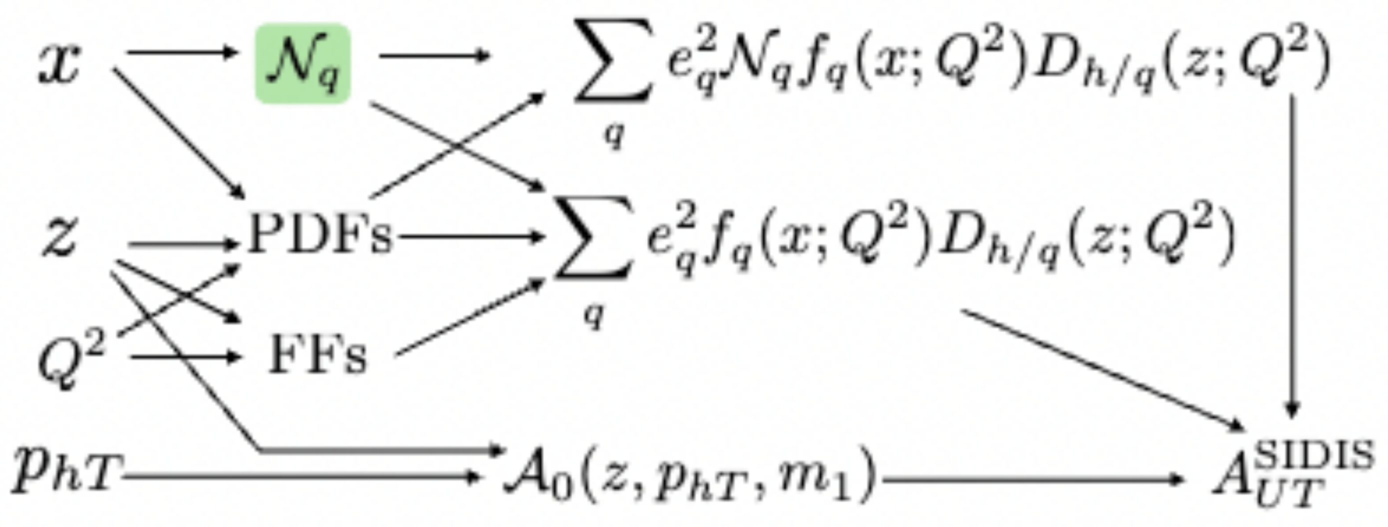

$\mathcal{N}_q^{\mathrm{DNN}}(x) h(k_\perp) f_{q/p}(x, k_\perp)$,

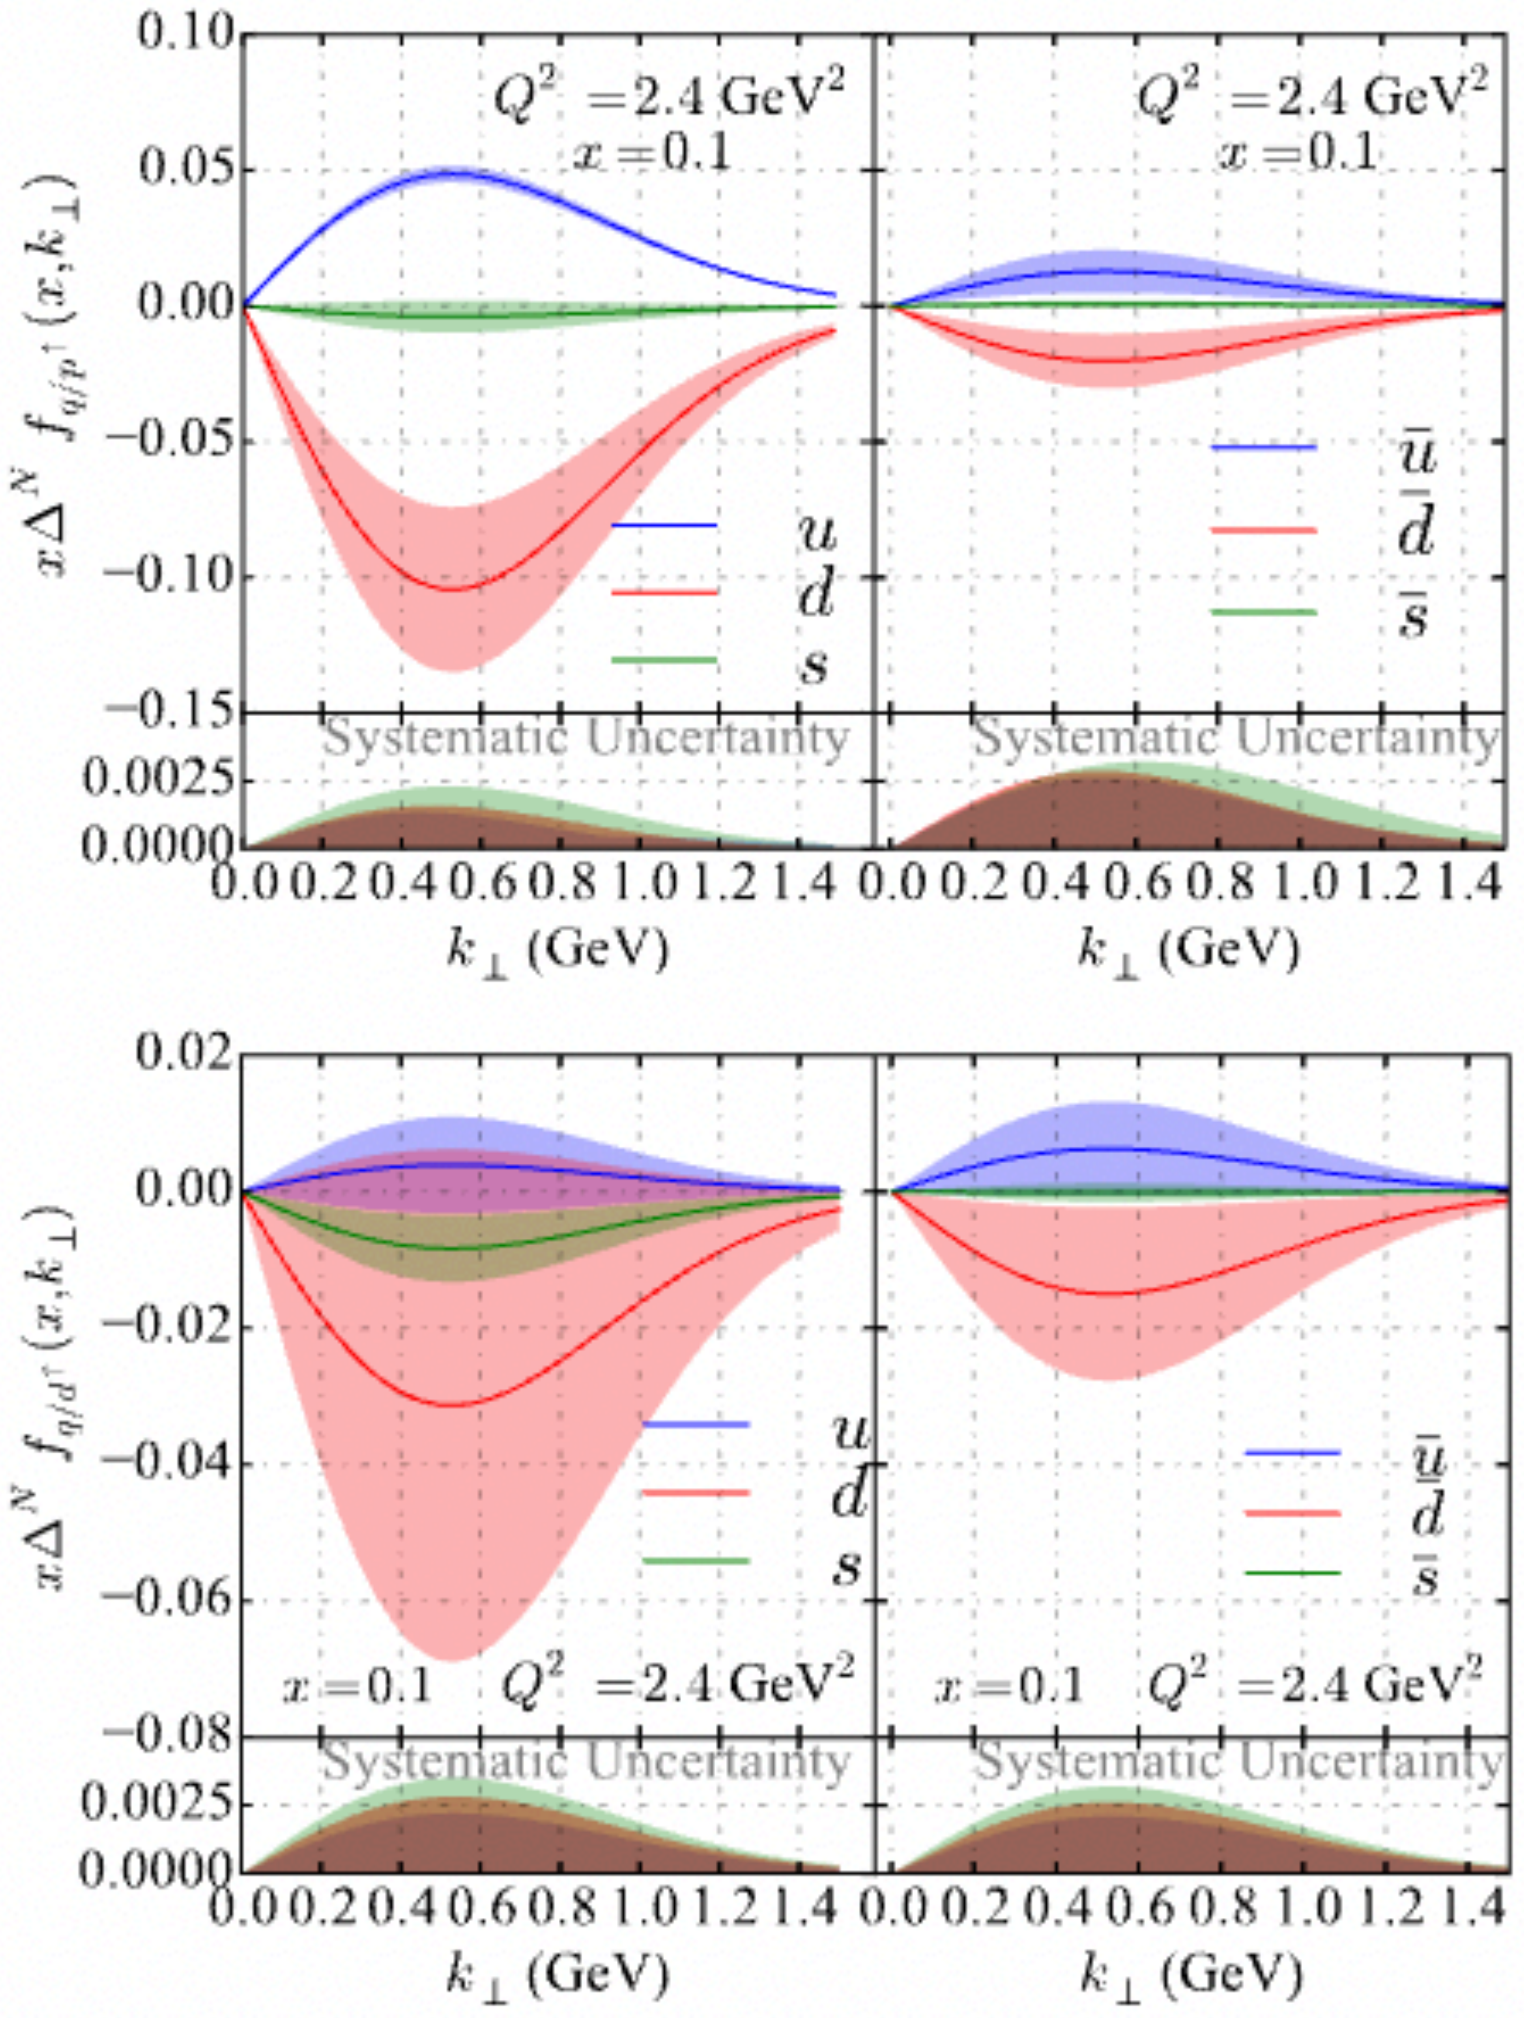

Ridiculously low quality data after decades of experiments

I. Fernando

CFF Extraction Assuming BKM

Pains taking Pseudodata Studies

- Still Assume formalism
- Very good local fits
- Improved global extraction

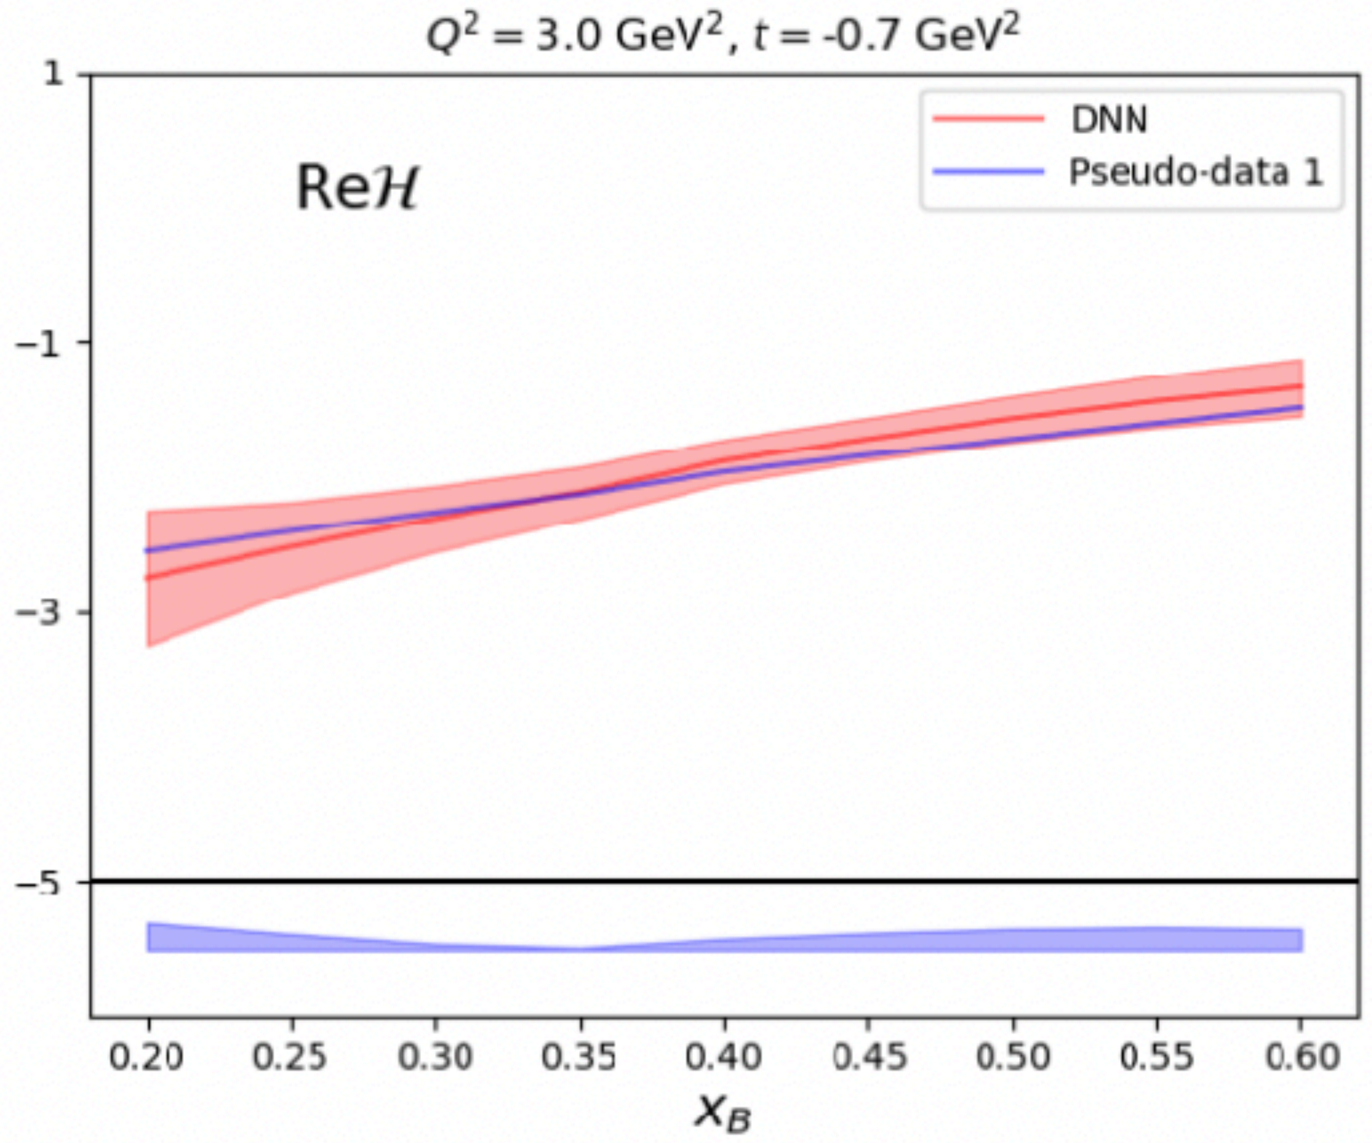

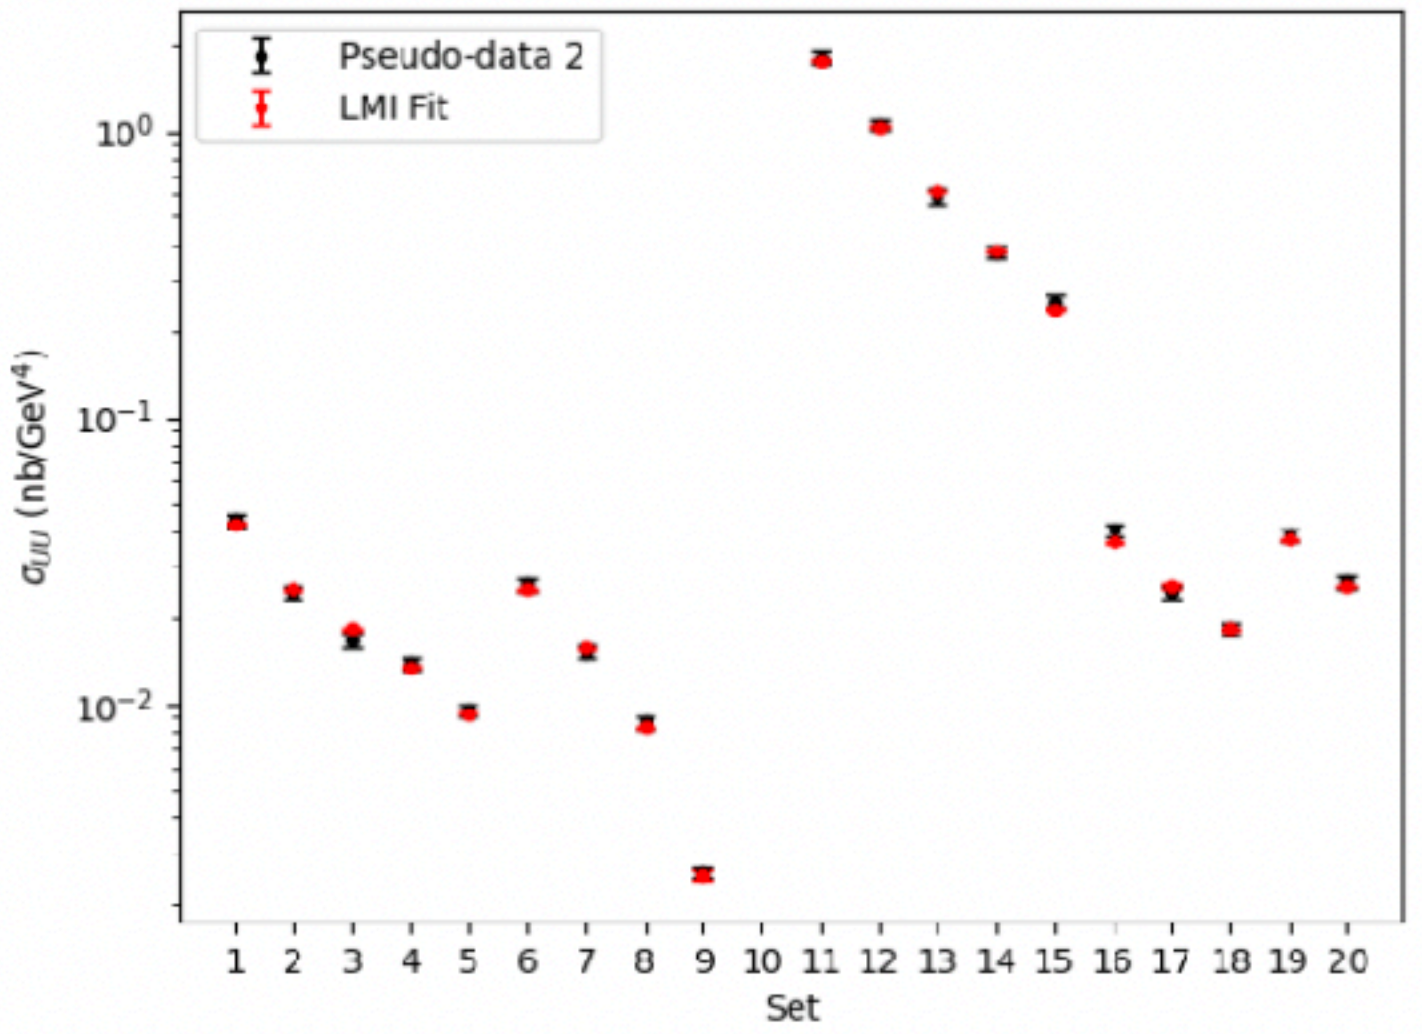

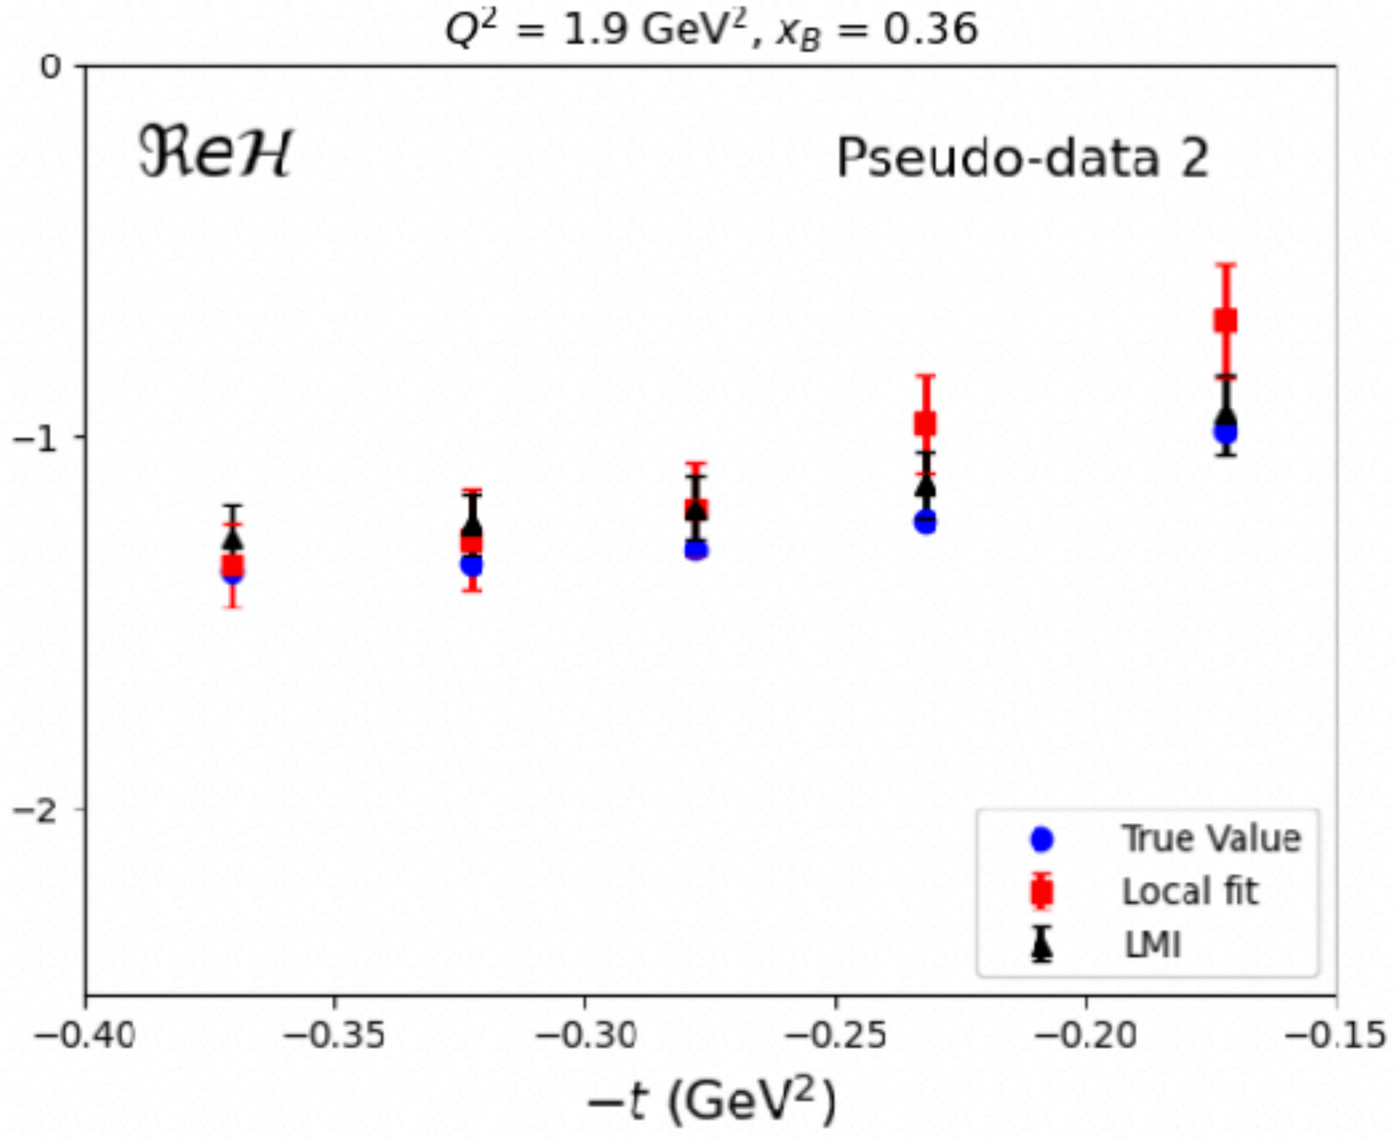

But doing local fitting in this traditional way doesn’t really make sense!!

Basic ANN Extraction

$f(x, y, z, a, b, c)$

Local Fit

- The experimental data is use to generate MC replica data
- Have collection of data with know f with unknown set of parameters
- Pass data and dependent variables through the network according to weights and activation function (first iteration weights are random values)
- The output is the first approximation of the desired parameters
- Use these output parameters to calculated f. The results are used in the loss function which is use to modify the network and adjusts weights so the error is decreasing

Long Standing Problems

Data should be globally accessible

- Once experiment analysis is performed
  - Stuck with kinematic bins
  - Little to no covariance information
  - Systematic errors inadequately
  - Raw data is lost
  - No storage for future use/analysis/fitting

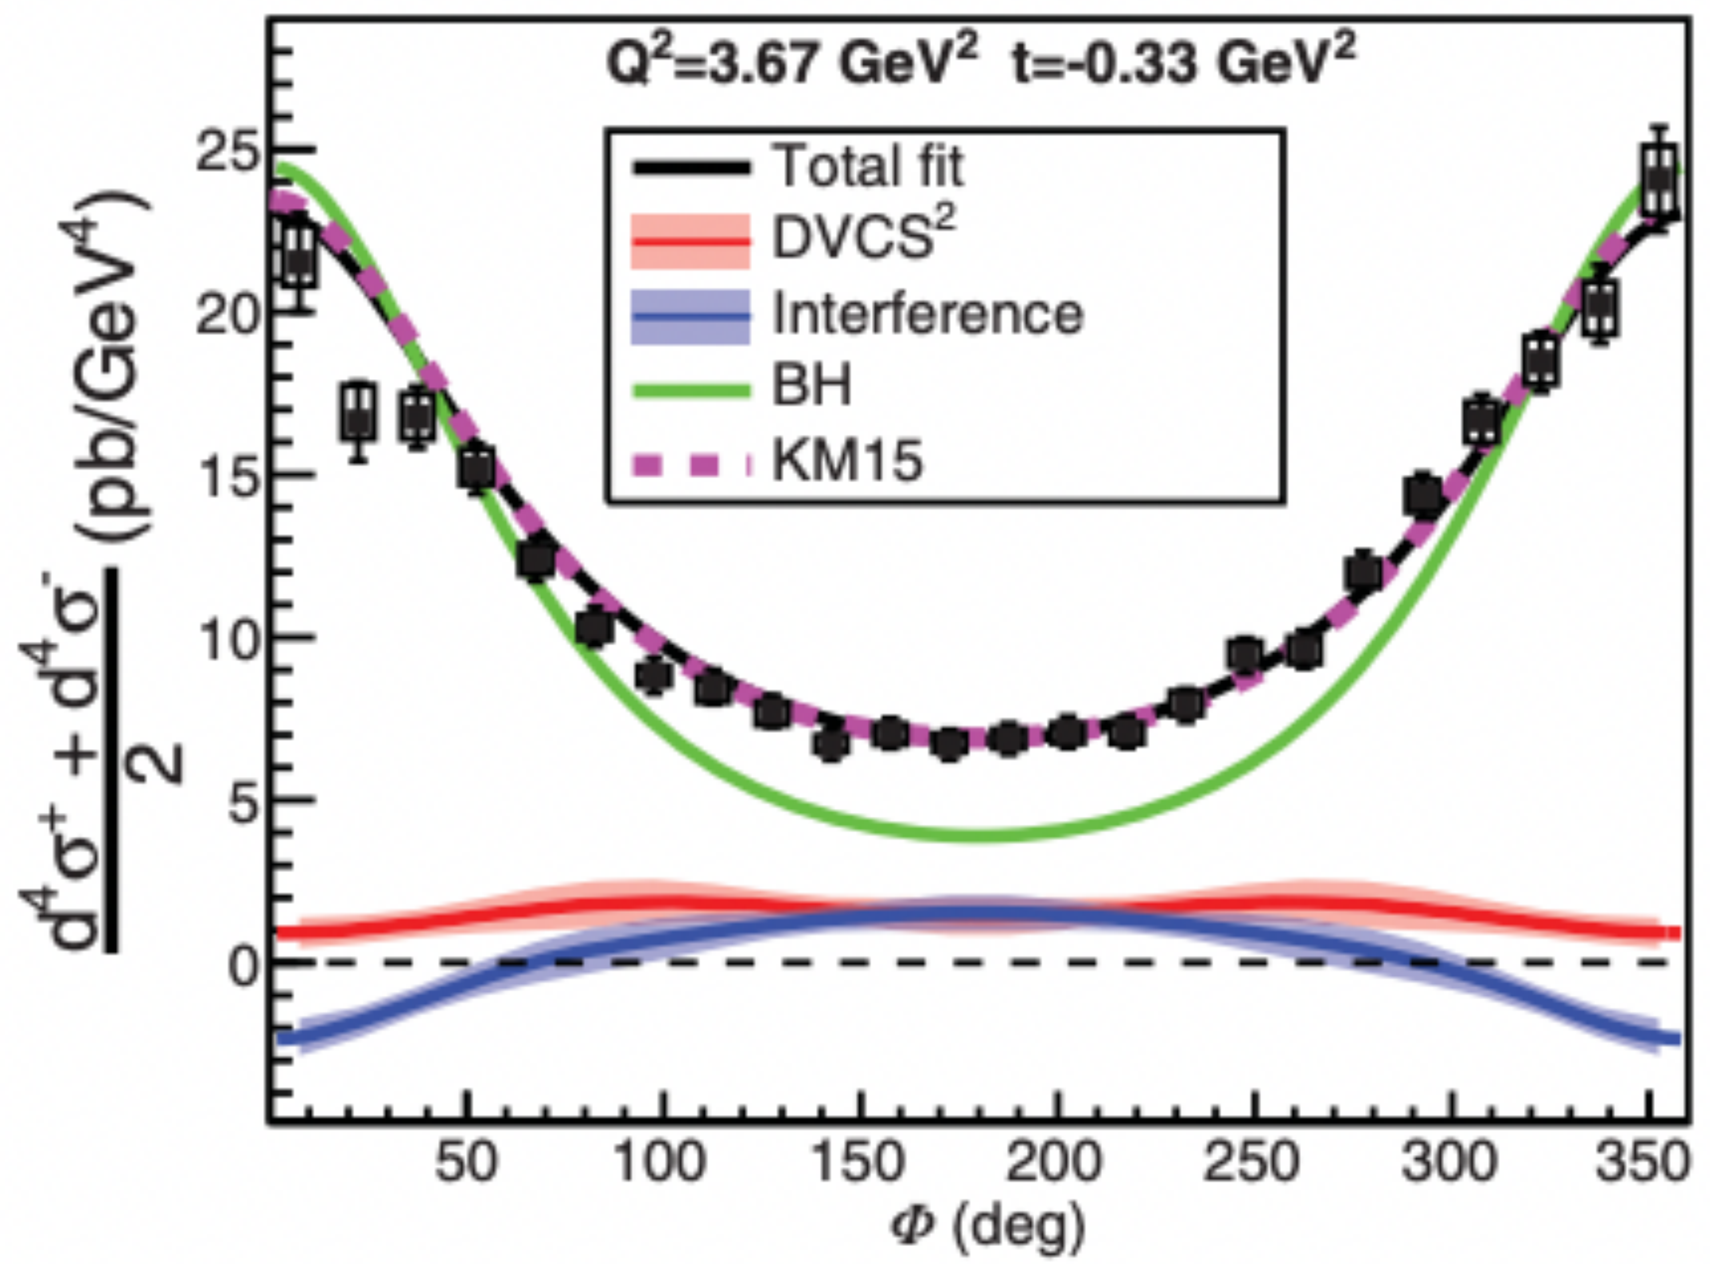

Georges, F. et al, Phys Rev Lett 128 252002

Same Issue with Polarized Target Data too!!

Storage Candidates

Needed to make data available to the world for training AI for everyone!!

- Compressed root file (Struct format)
- Compressed numpy arrays
- Hierarchical Data Format (HDF5)
- Reduce types: float16, int8, uint8
- Reduce structure root/npy: few dimensions and elements
- Saved raw but not too raw…
    - Level of information
    - I/O speed
    - Compressibility
    - Able to read publicly with standardized data dictionary

New Measurement Approach

How to Measure and What is the Error

- Measure
  - Assume TE and Boltzmann signal studies done well during calibration
  - Use the ***3-principles*** ss-RF extraction
  - Continue to sweep-measure/sweep-manipulate
- Uncertainty
  - Additional Error from modulating but have tools to improve

The Three Principles

For Enhanced Tensor Polarization that can be measured

A. Differential Binning

B. Spin Temperature Consistency

C. Rates Response

https://doi.org/10.1016/j.nima.2020.164504

https://doi.org/10.1016/j.nima.2023.168177

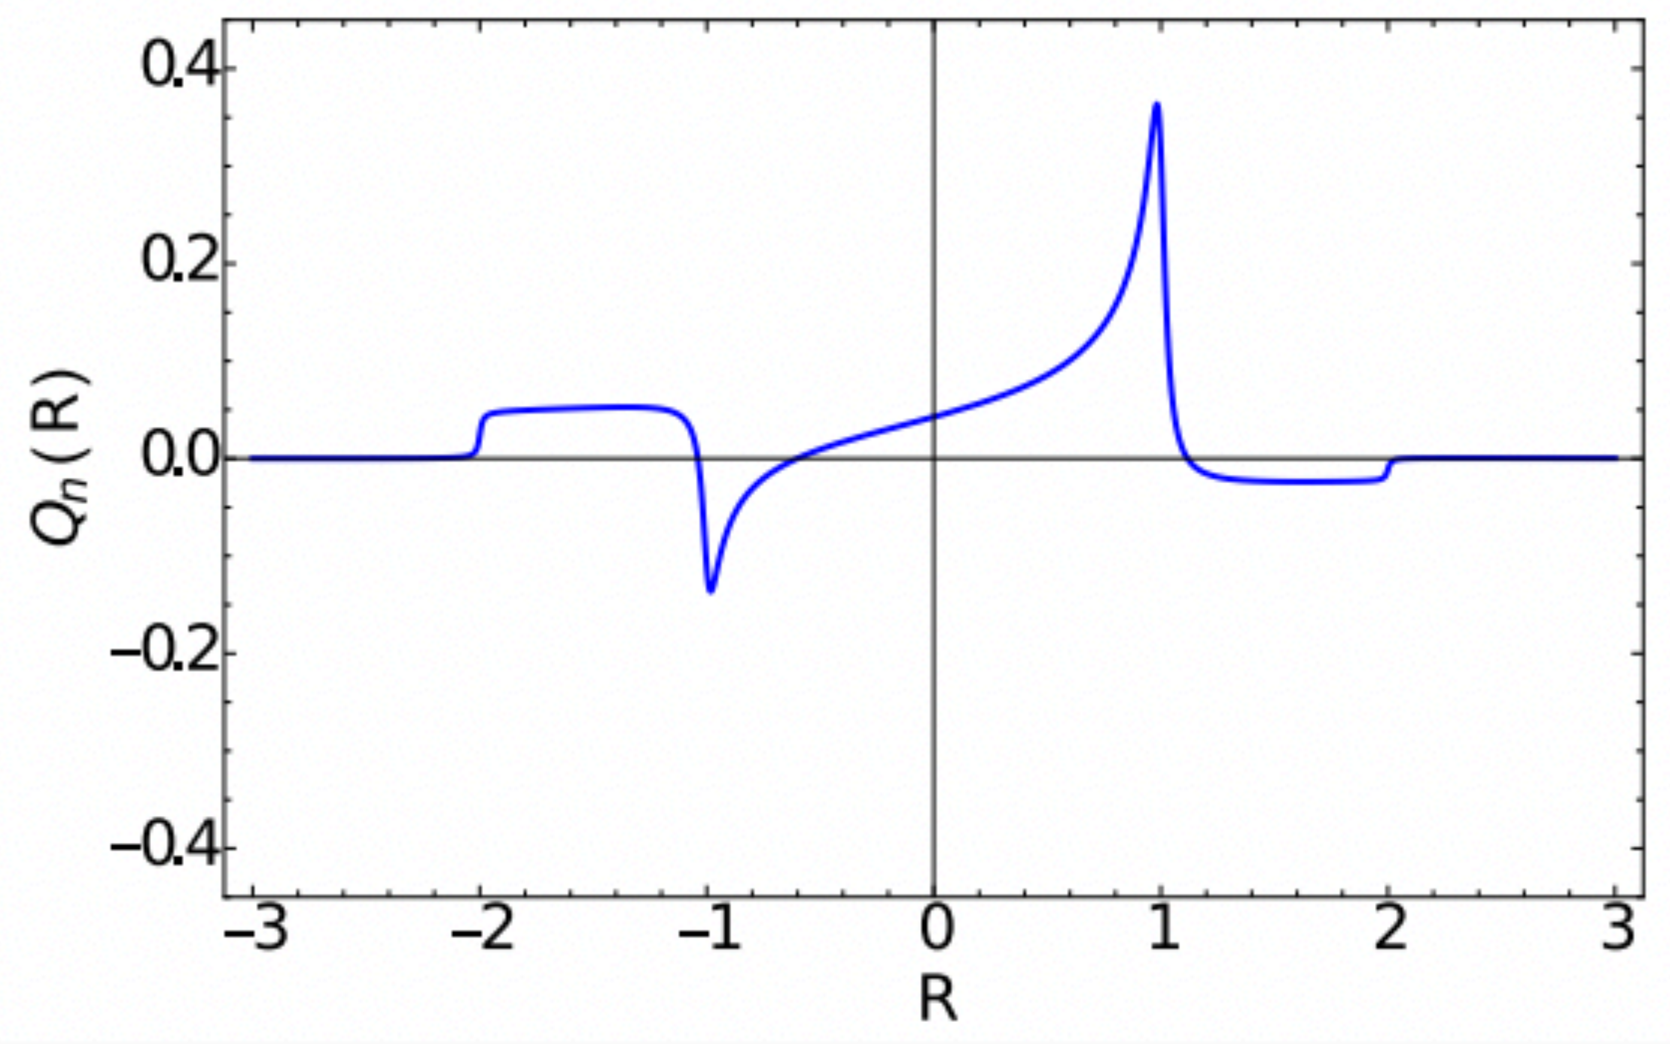

Differential Binning

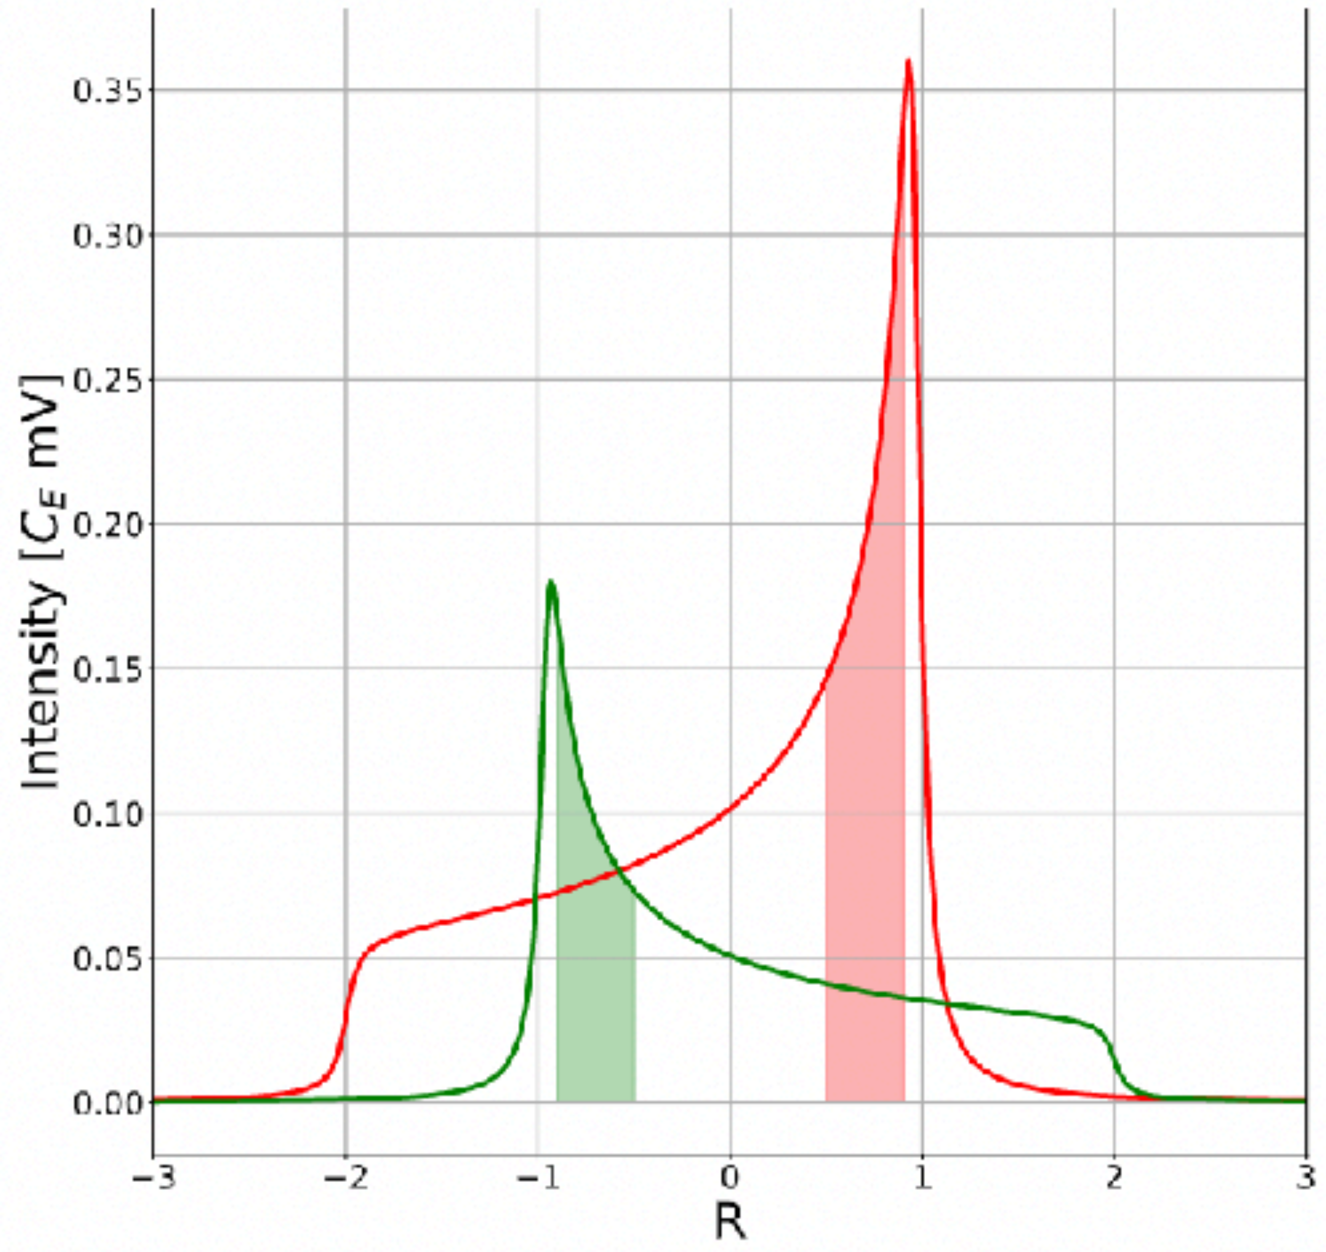

$$P(R) = C(I_+(R) + I\_(R))$$
$$P(\theta) = C(I_+(\theta) + I\_(\theta))$$

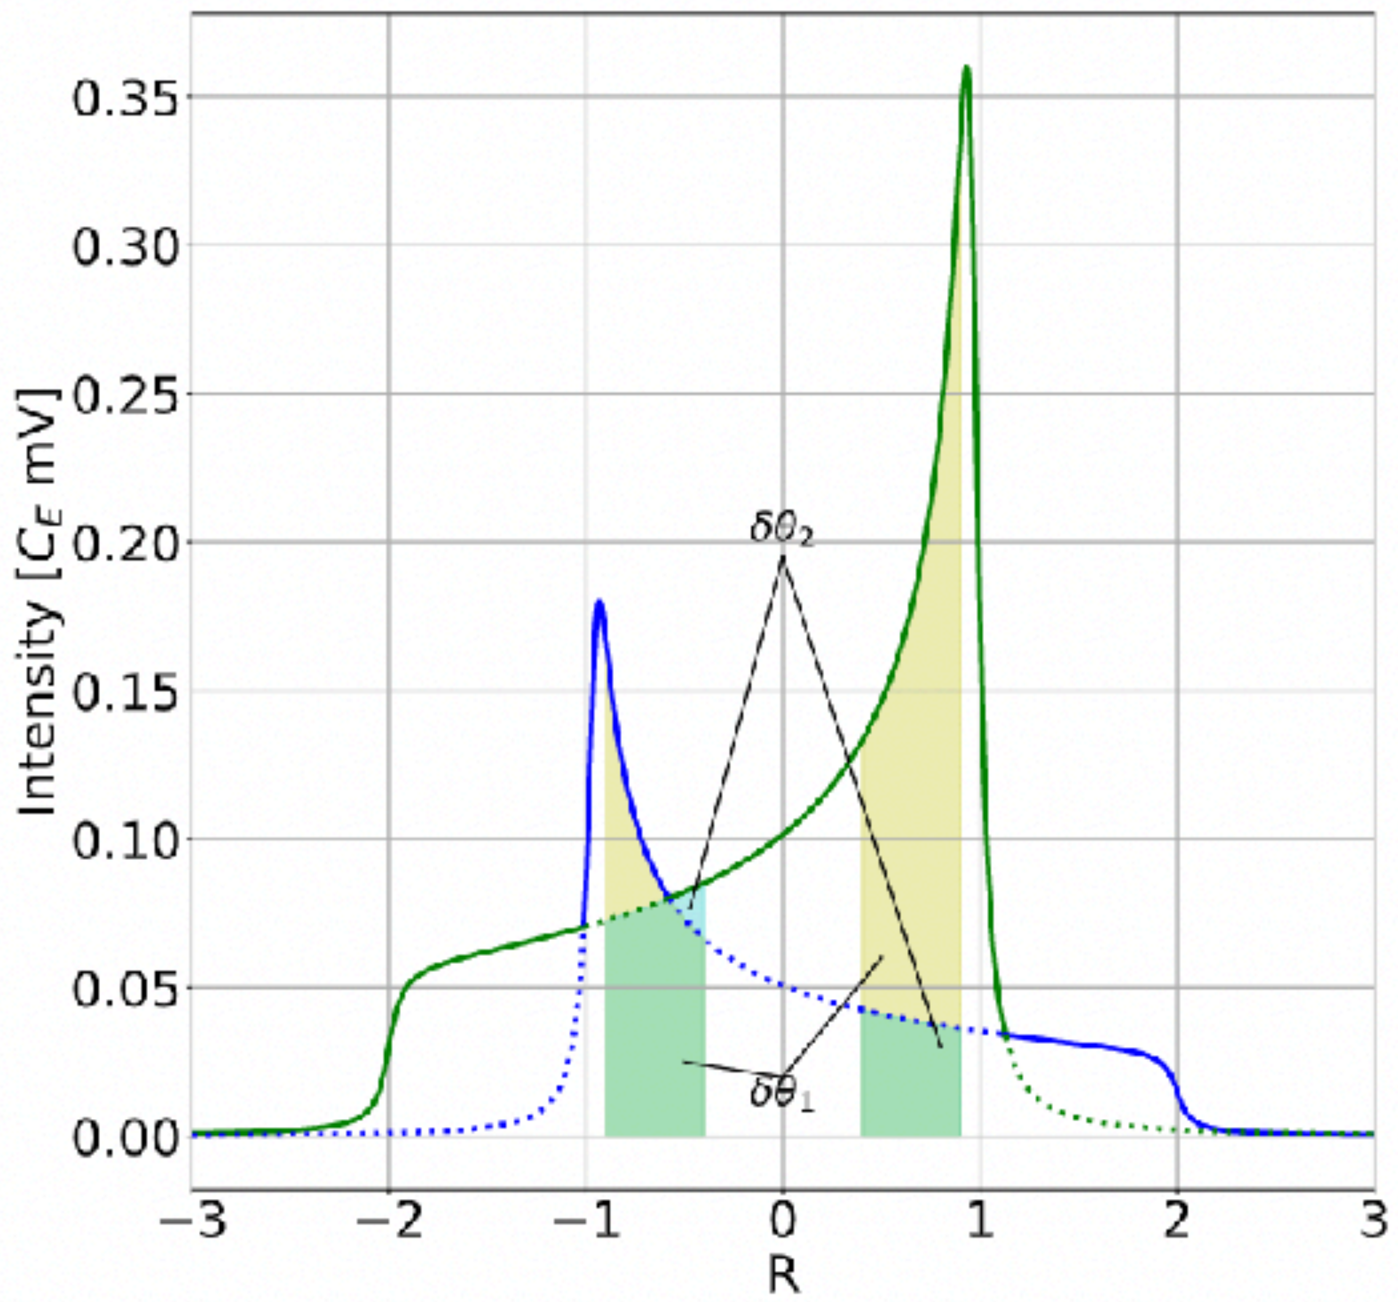

$$Q(R) = C(I_+(R) - I\_(R))$$

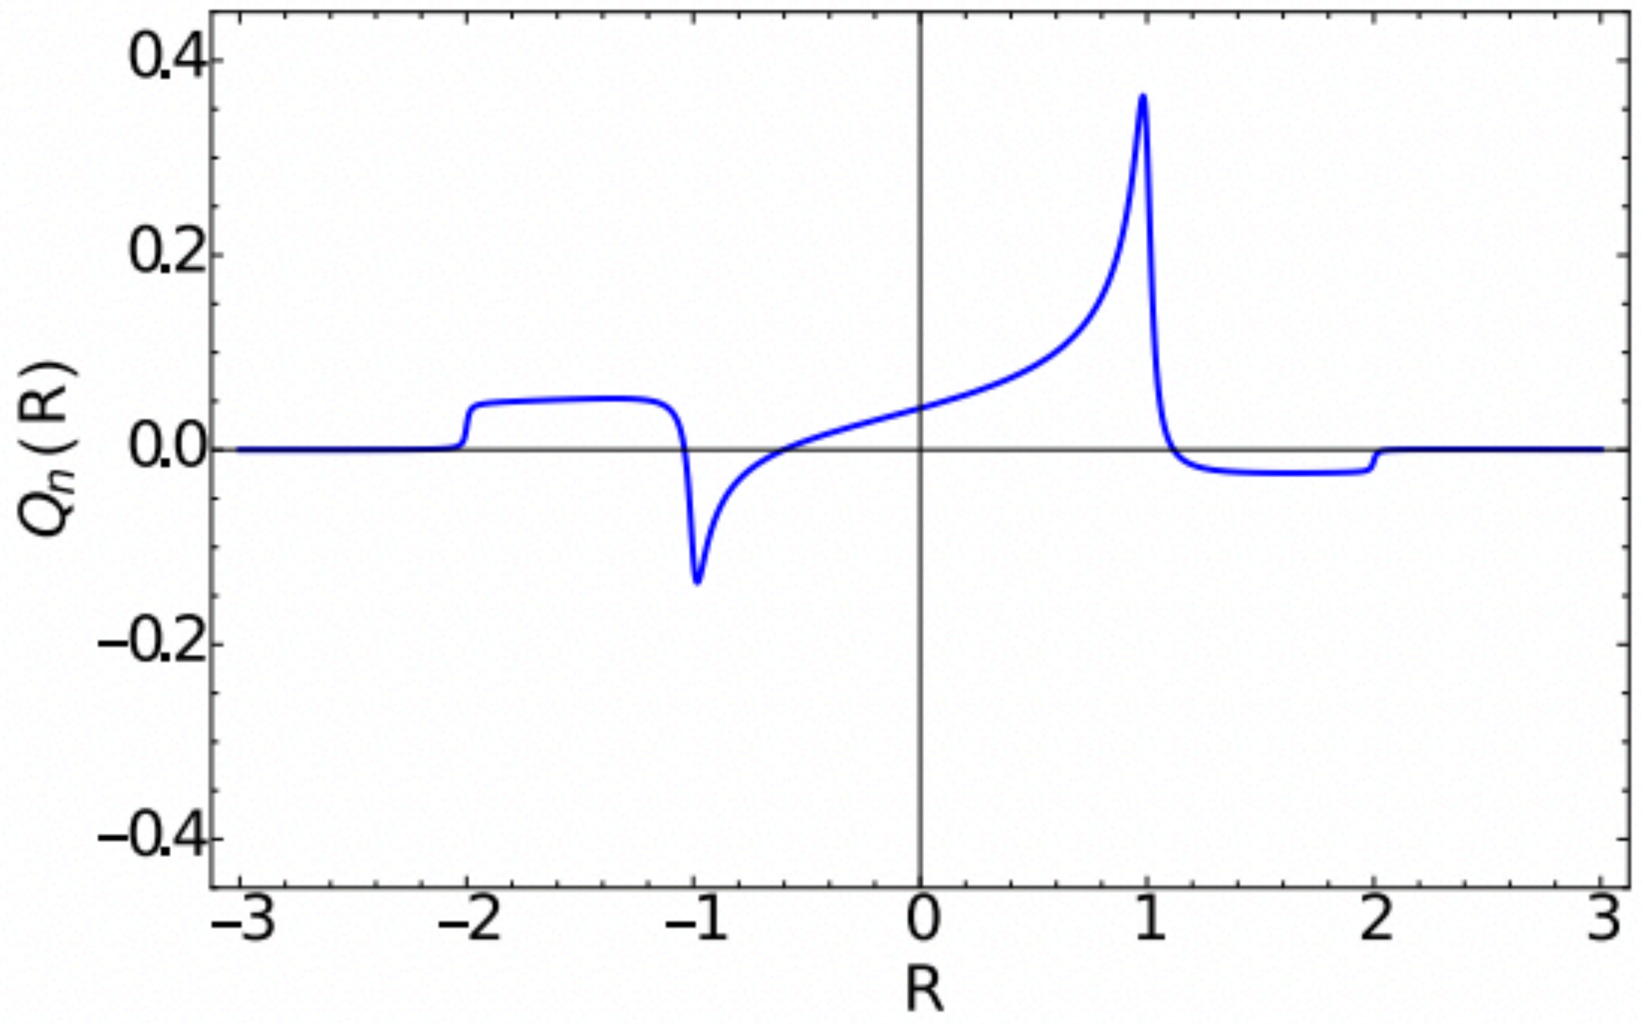

Depolarized back to zero

Spin Temperature Consistency

or strategic depolarization

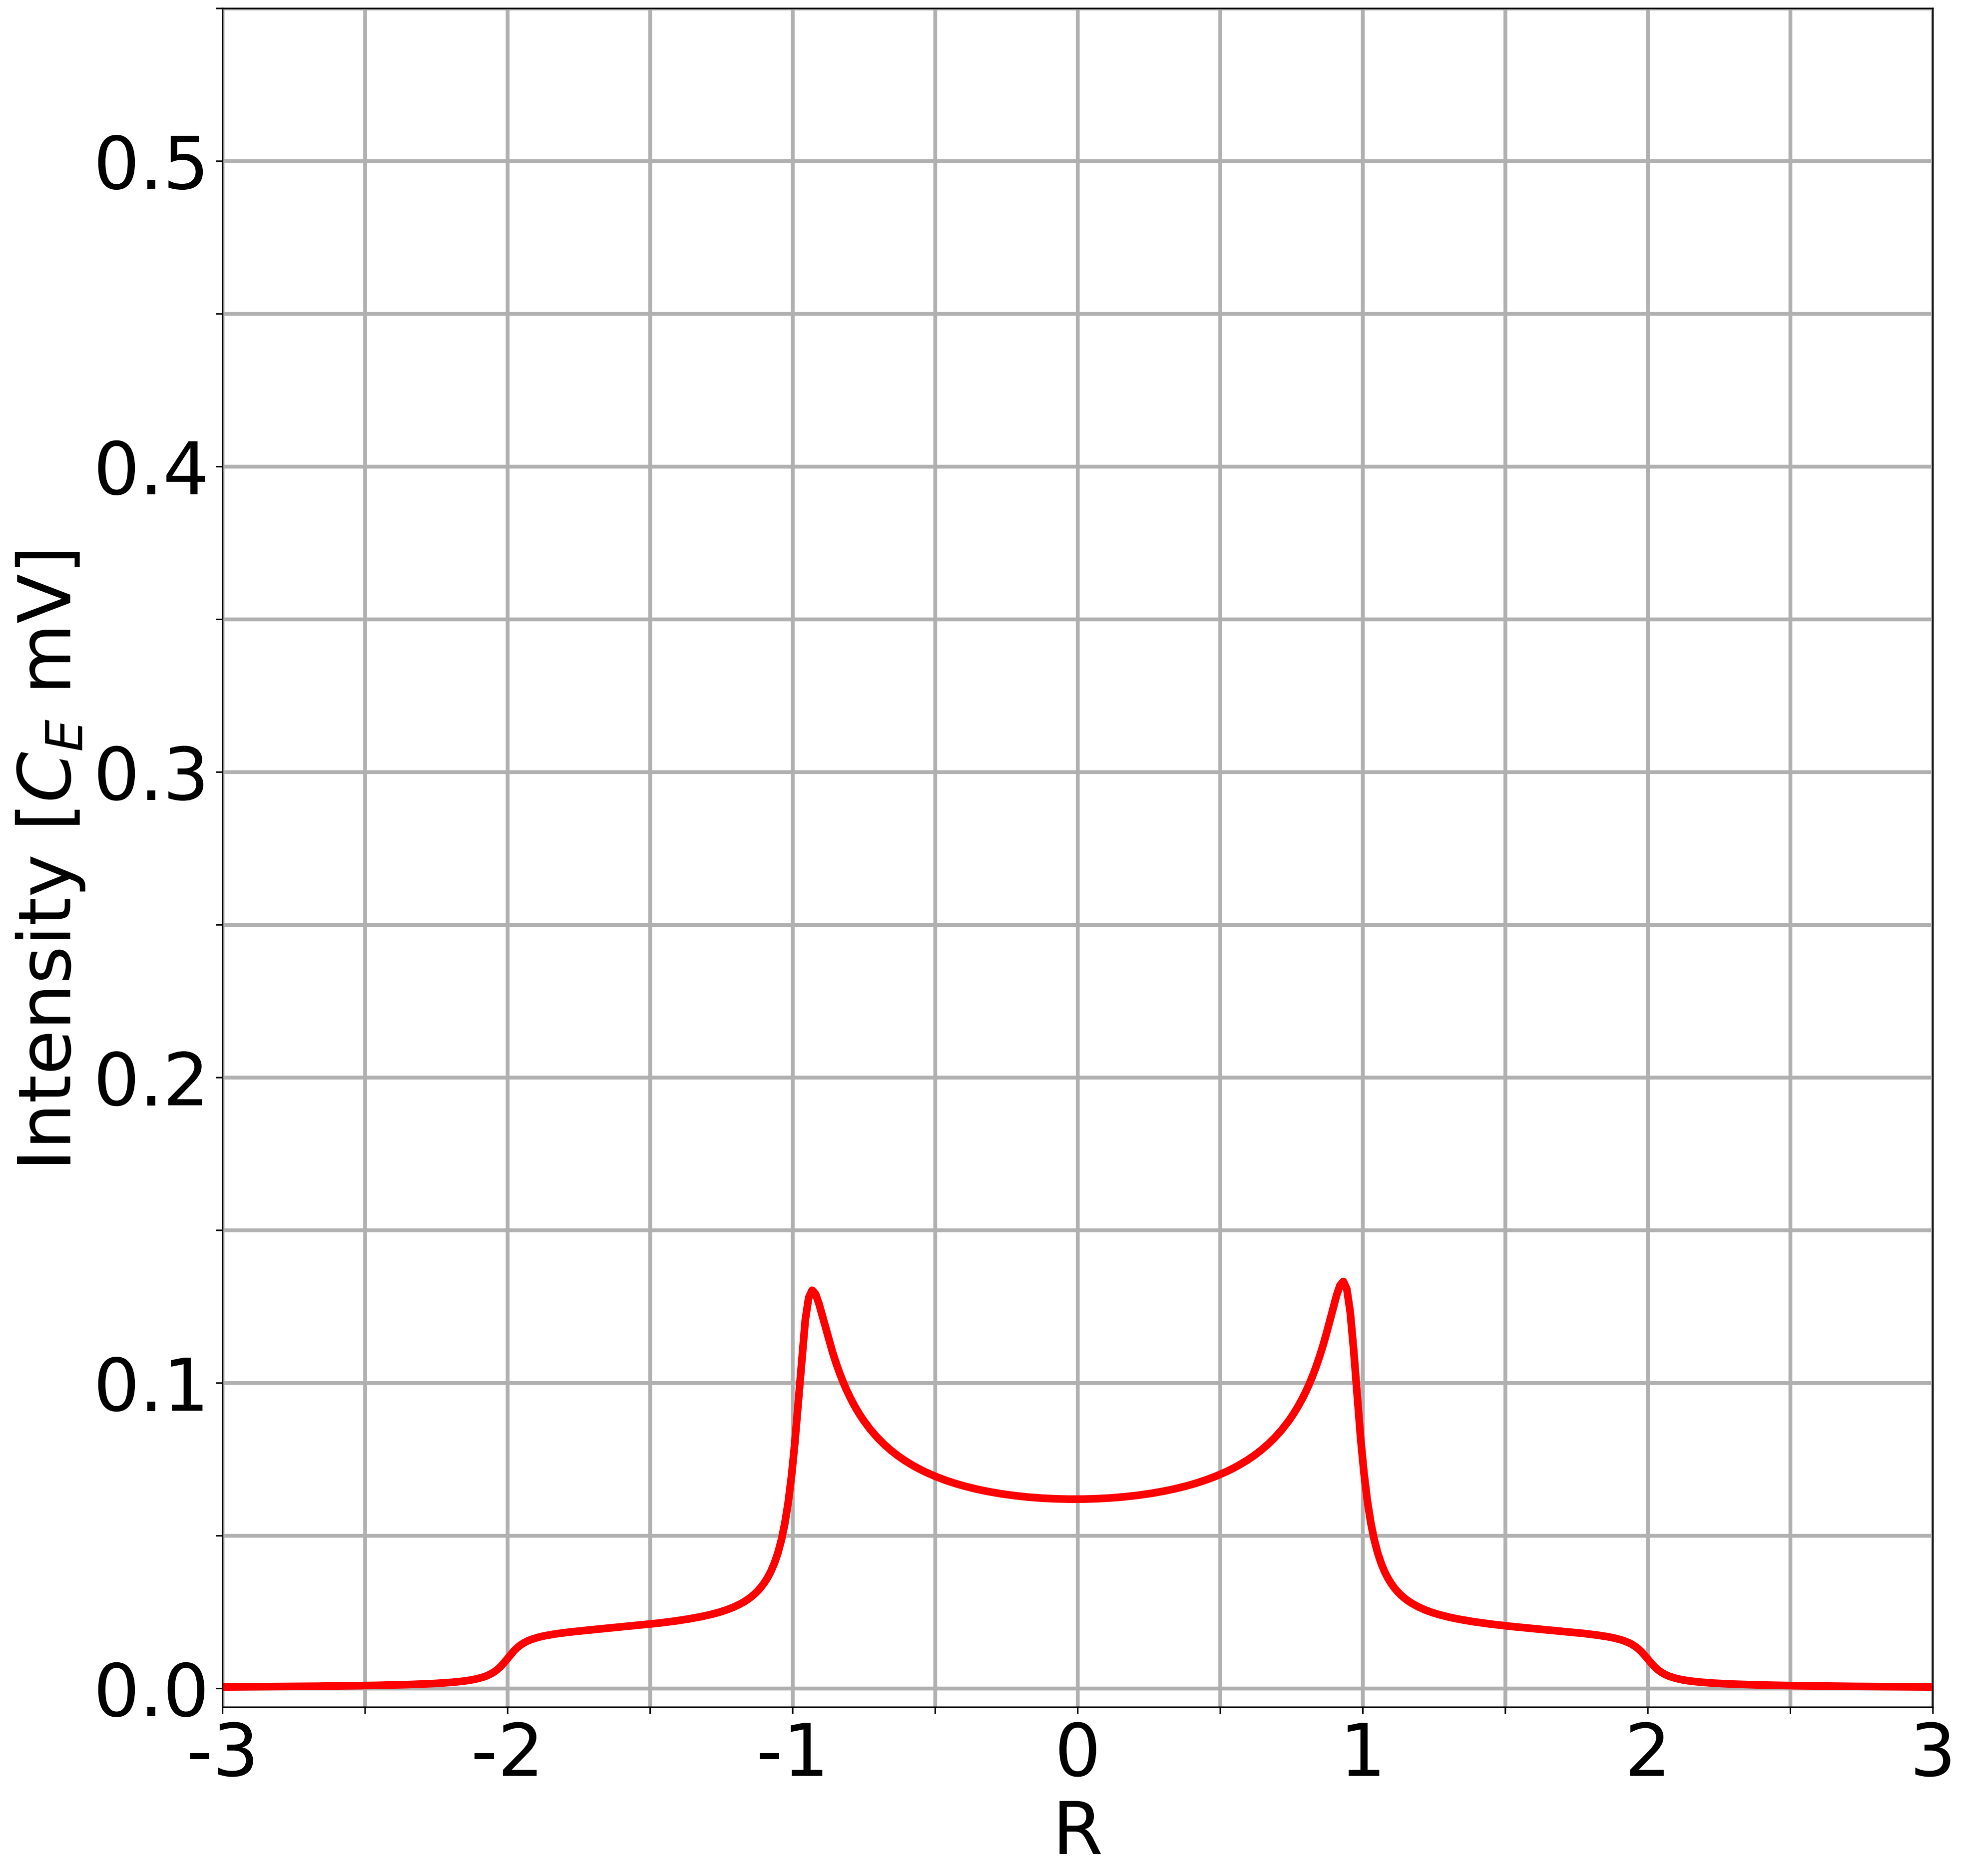

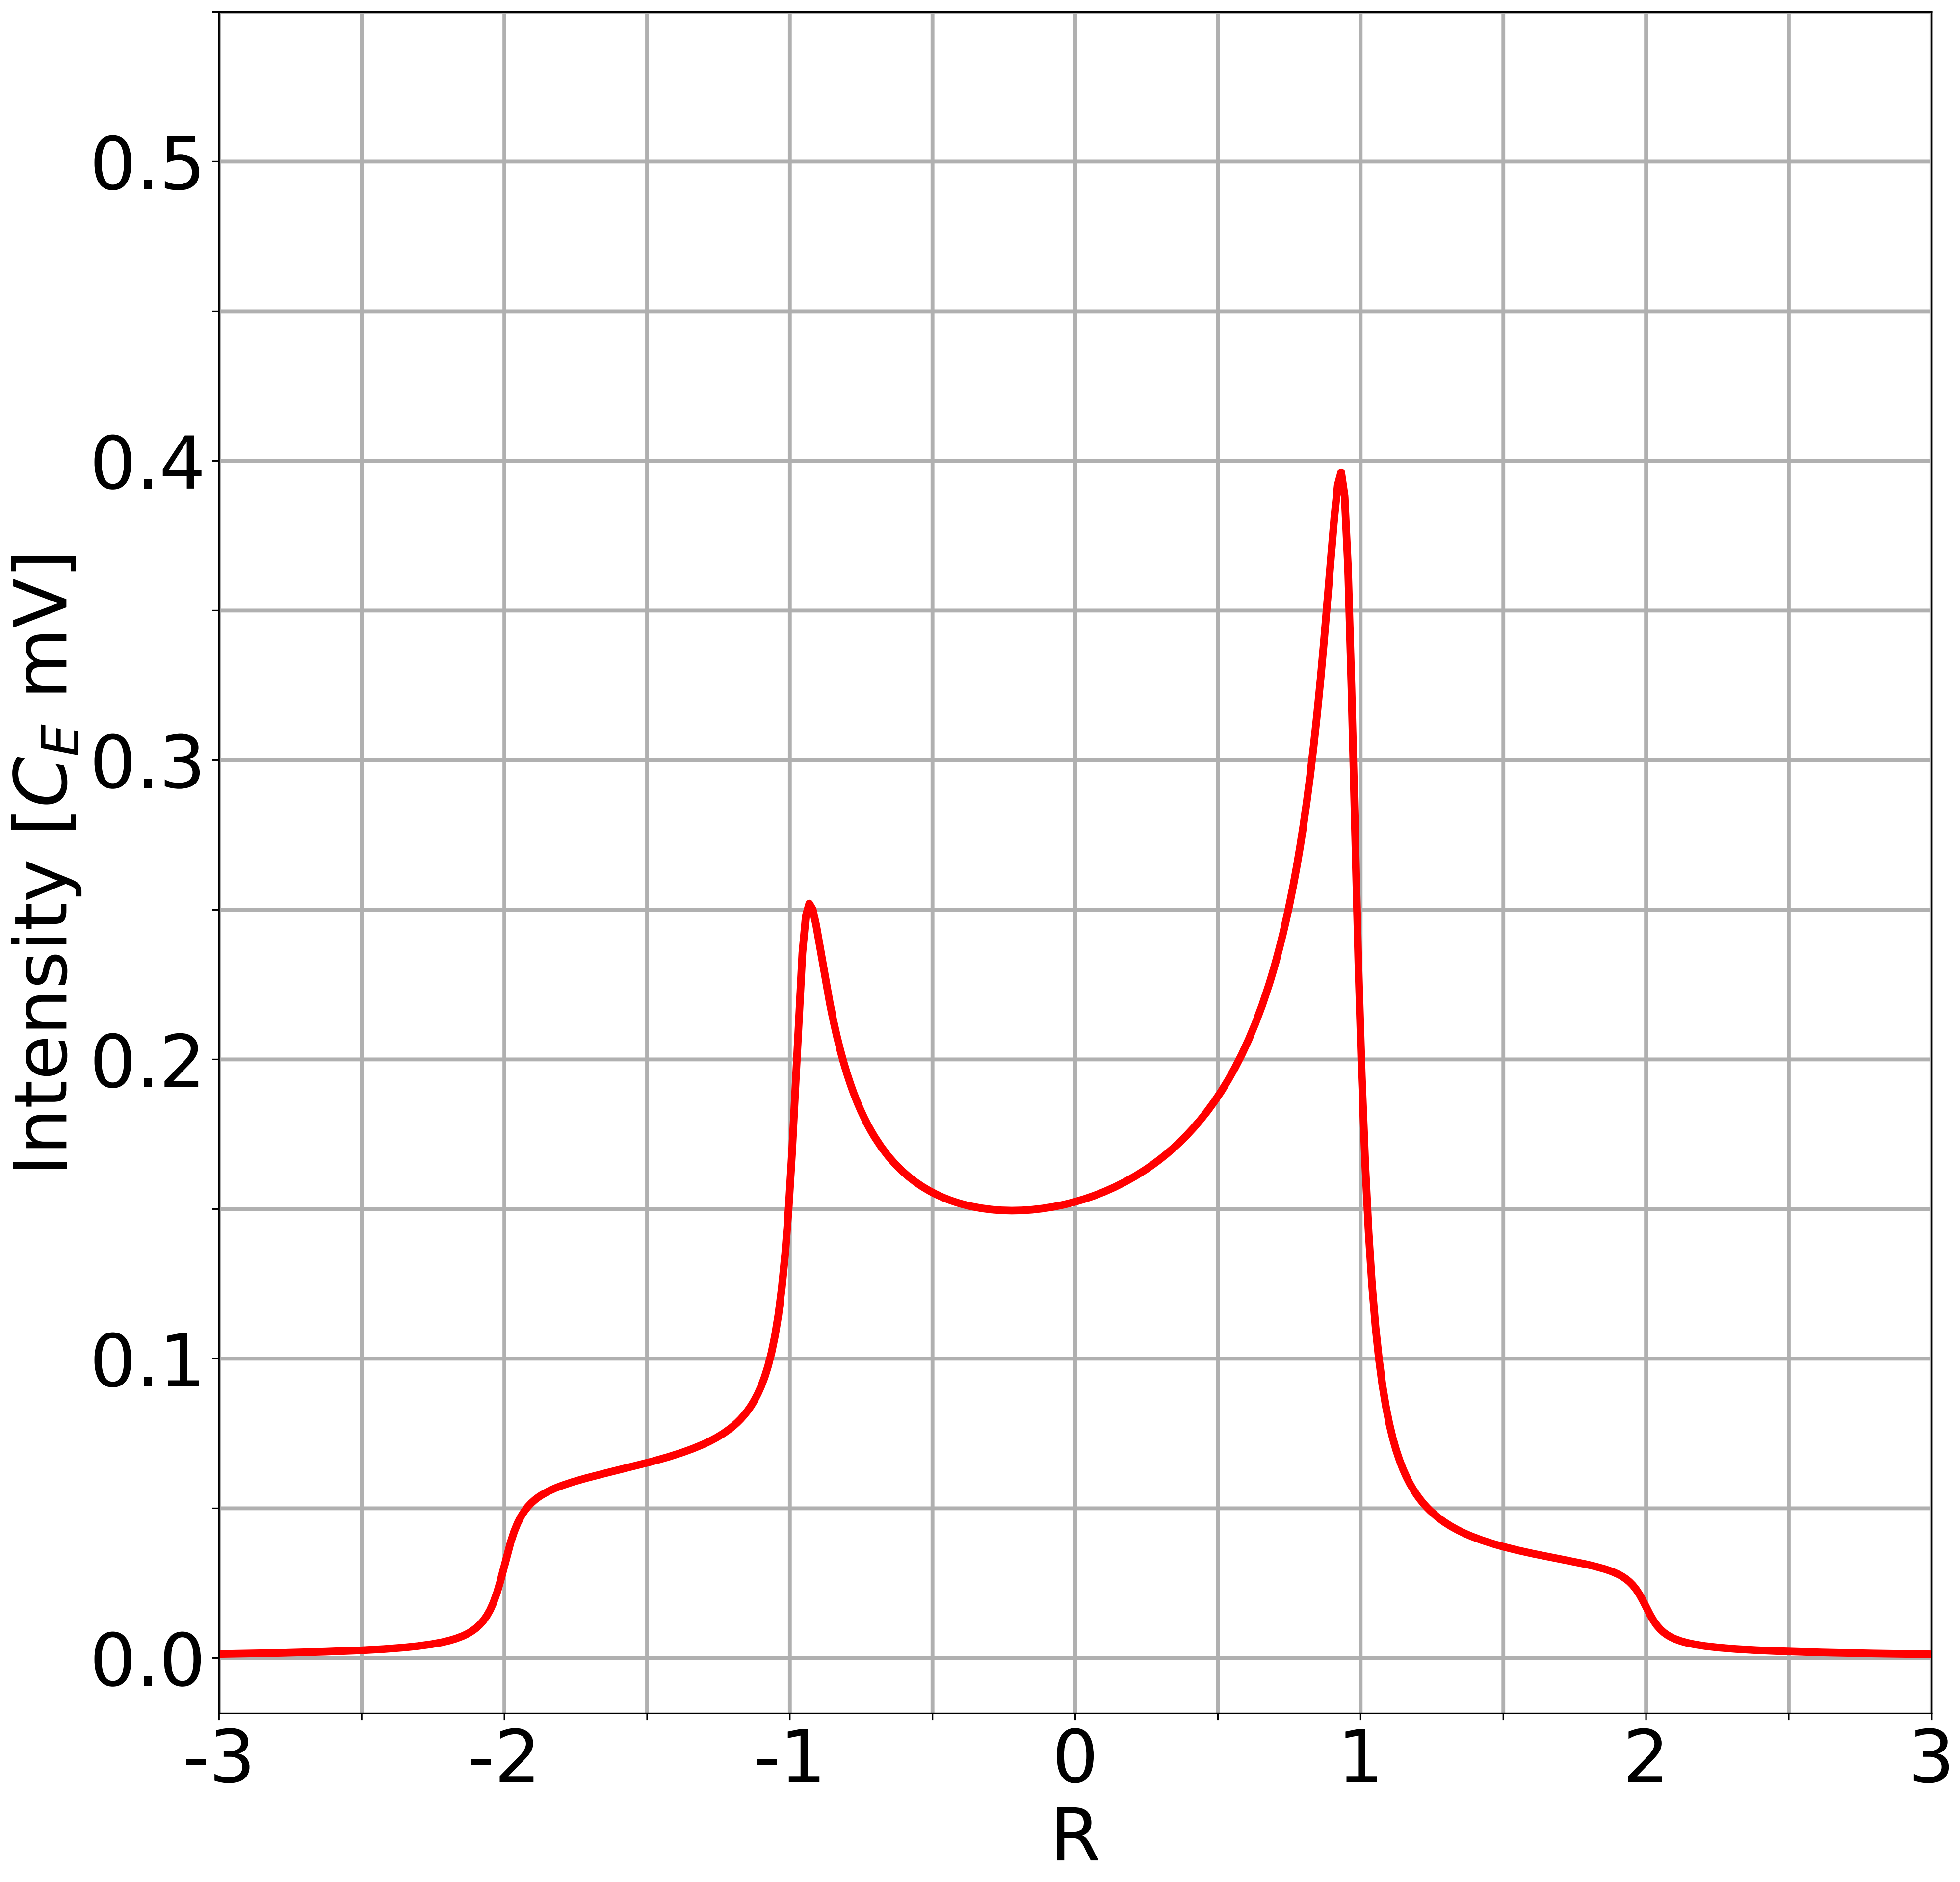

Spin Temperature Consistency

or strategic depolarization

Applying saturating or semi-saturating RF
depolarize across the selected domain

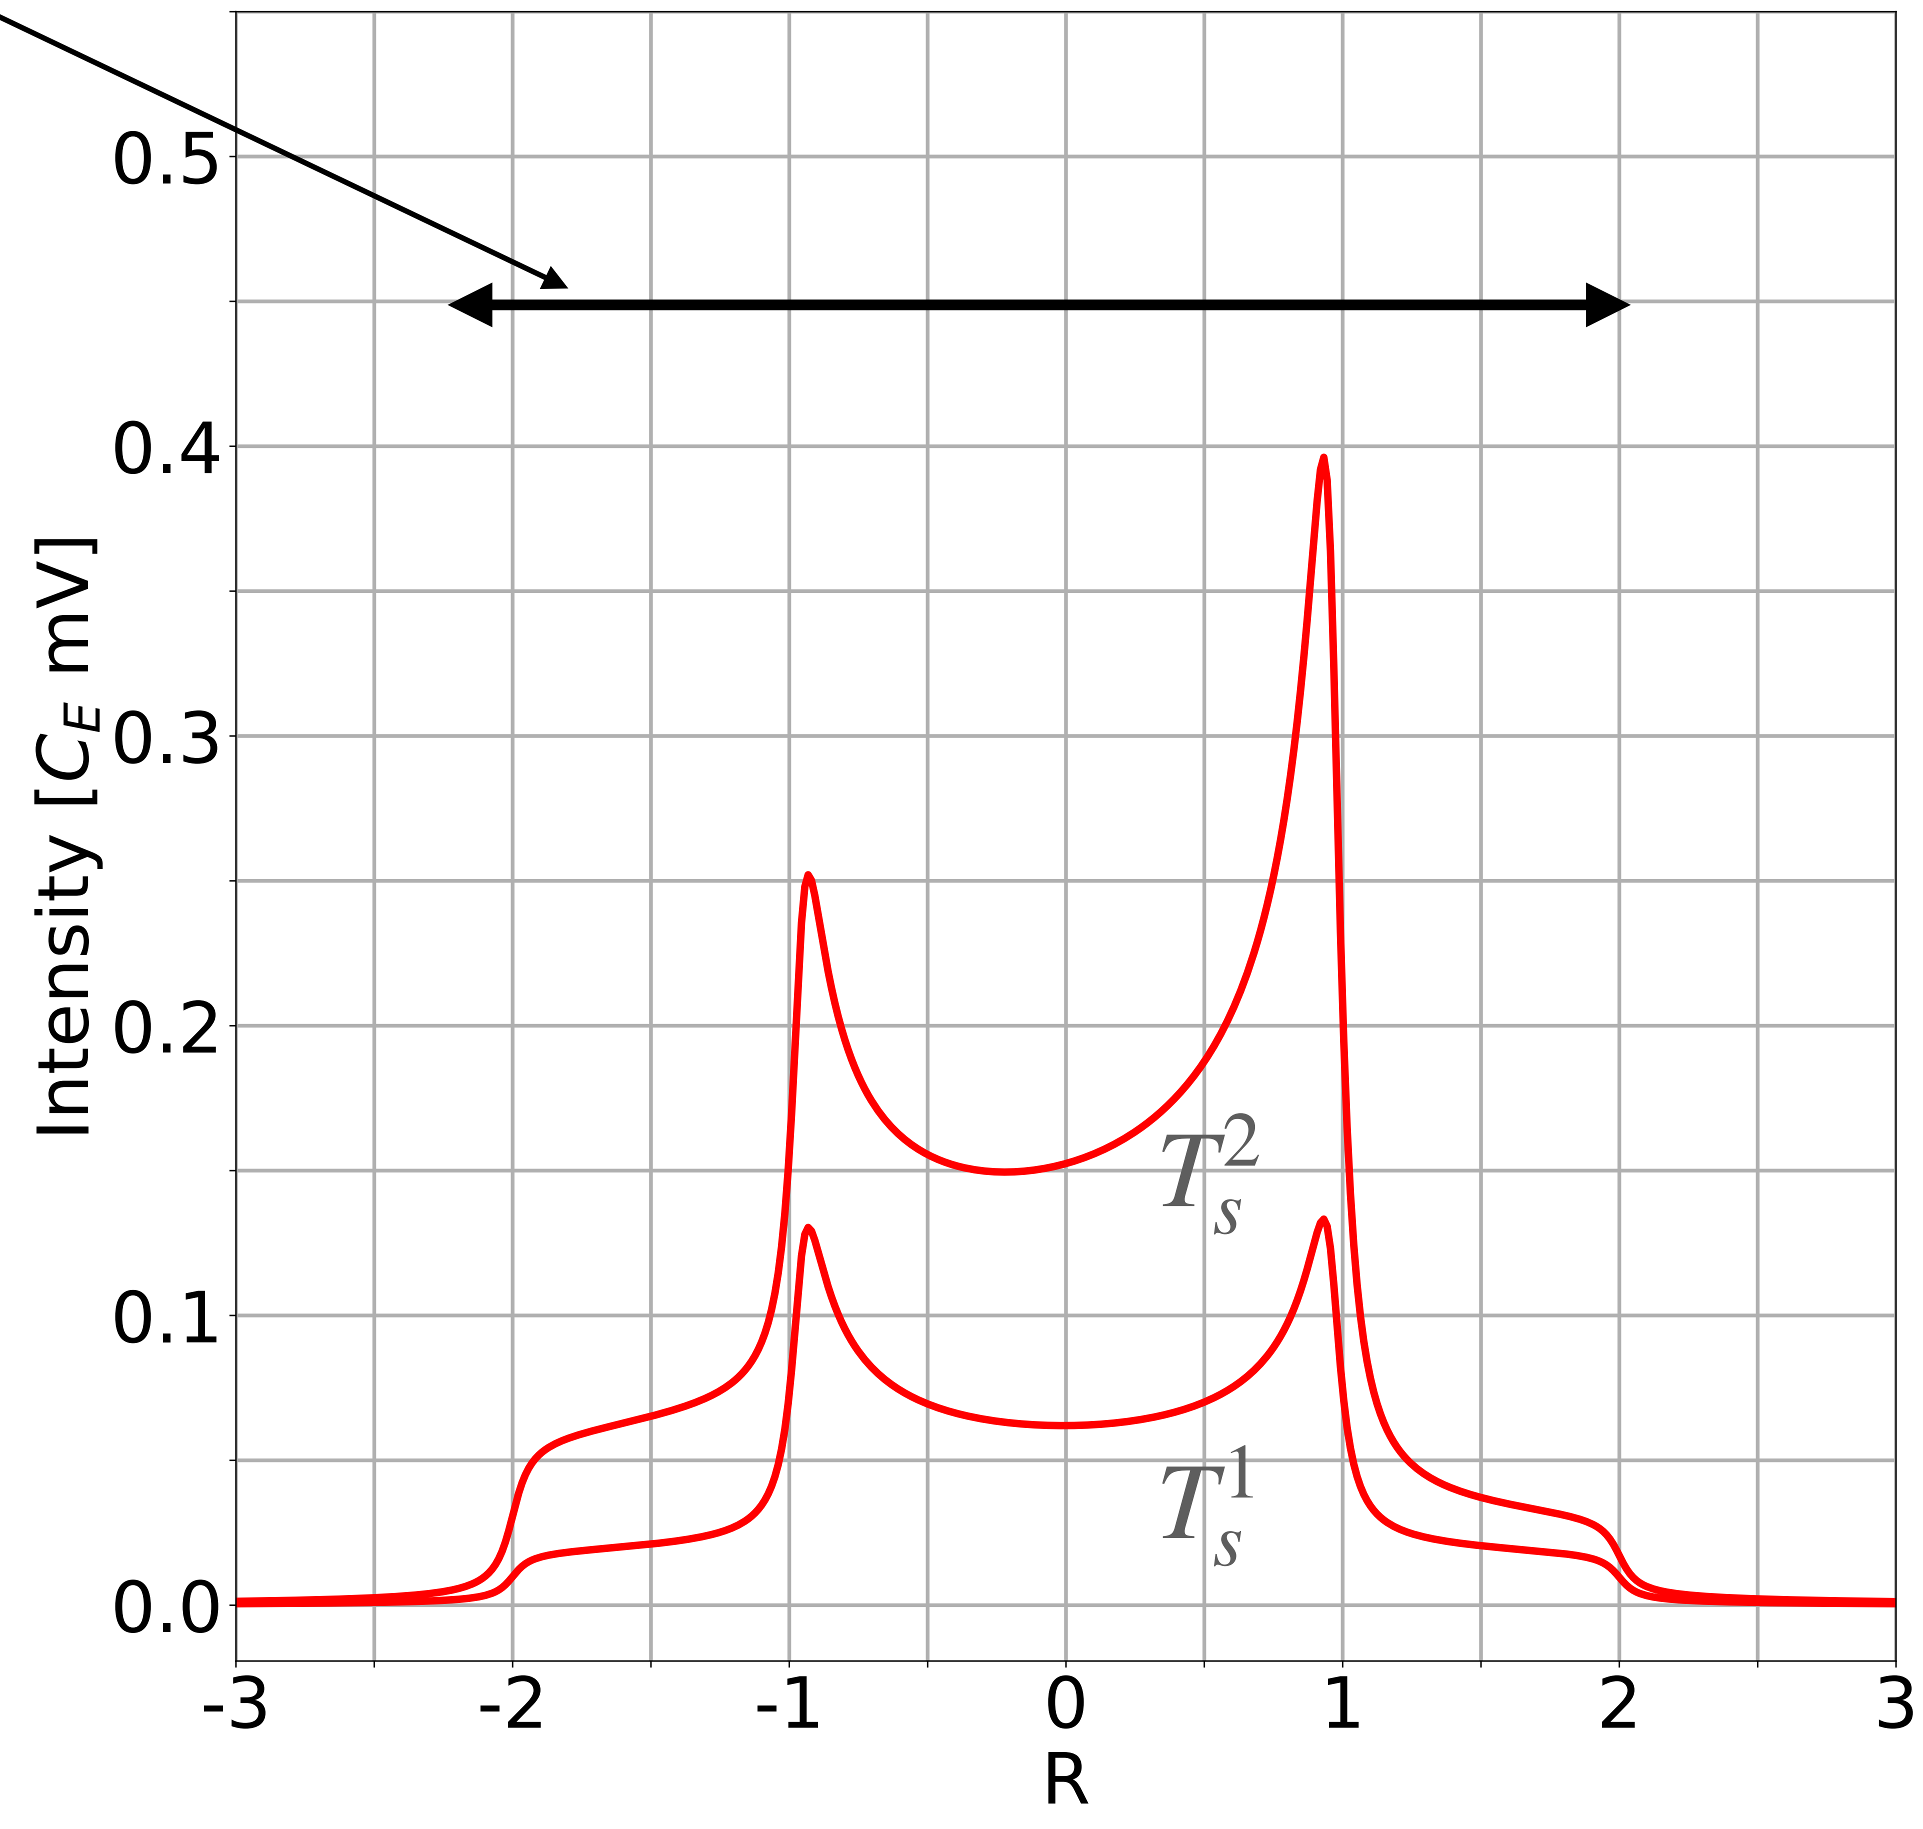

$$\omega_m \gg 2\pi/T_1$$

Fast enough that we treat this
As the application of a homogeneous field

Spin Temperature Consistency

or strategic depolarization

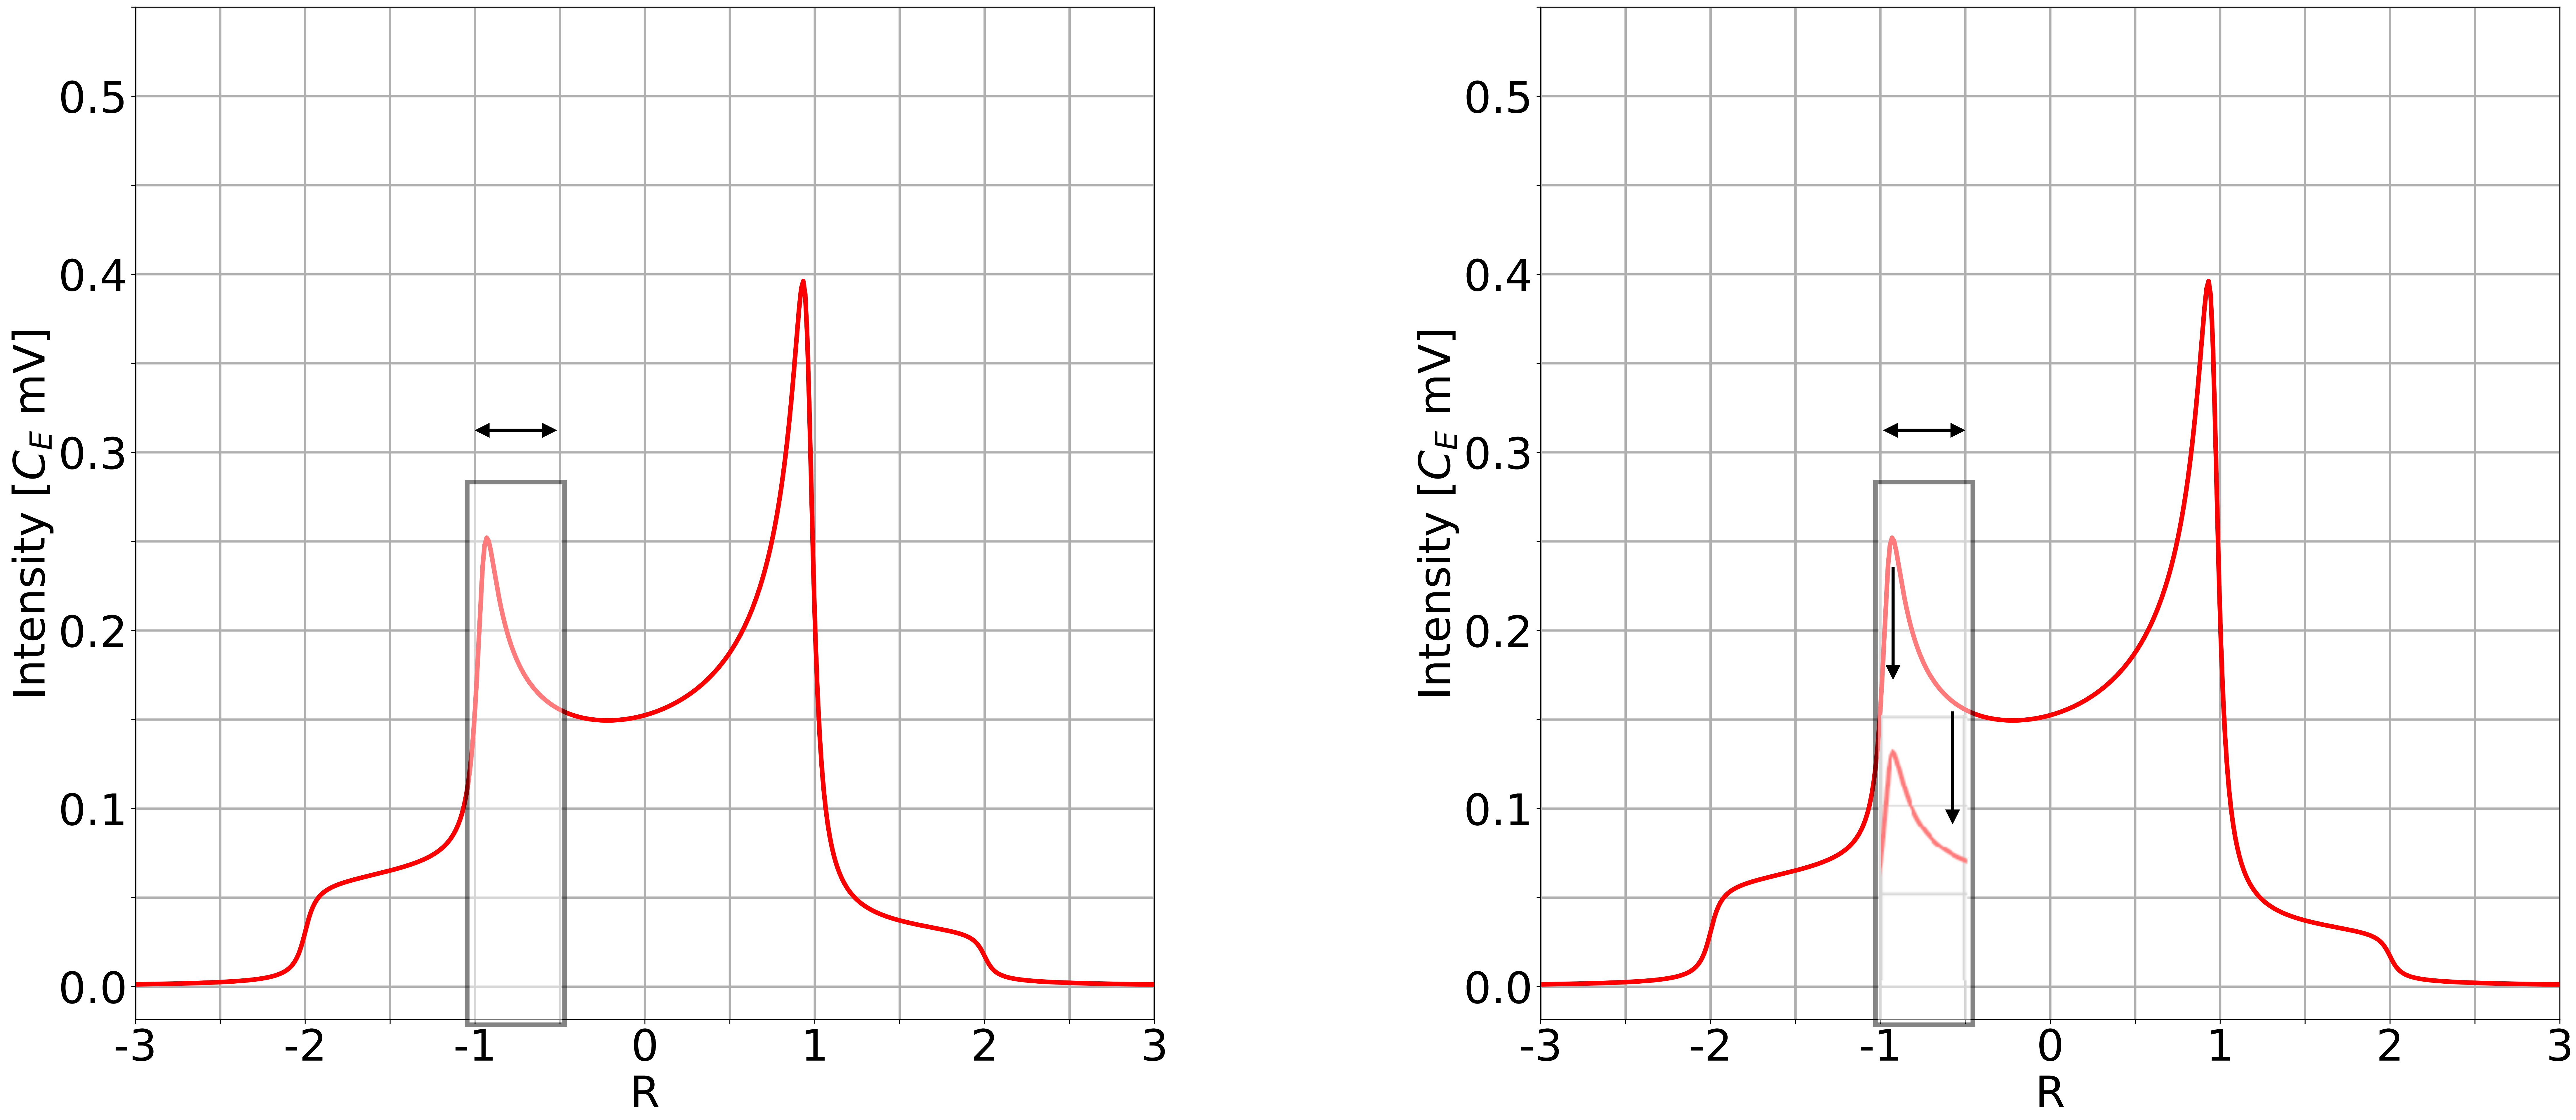

Spin Temperature Consistency

or strategic depolarization

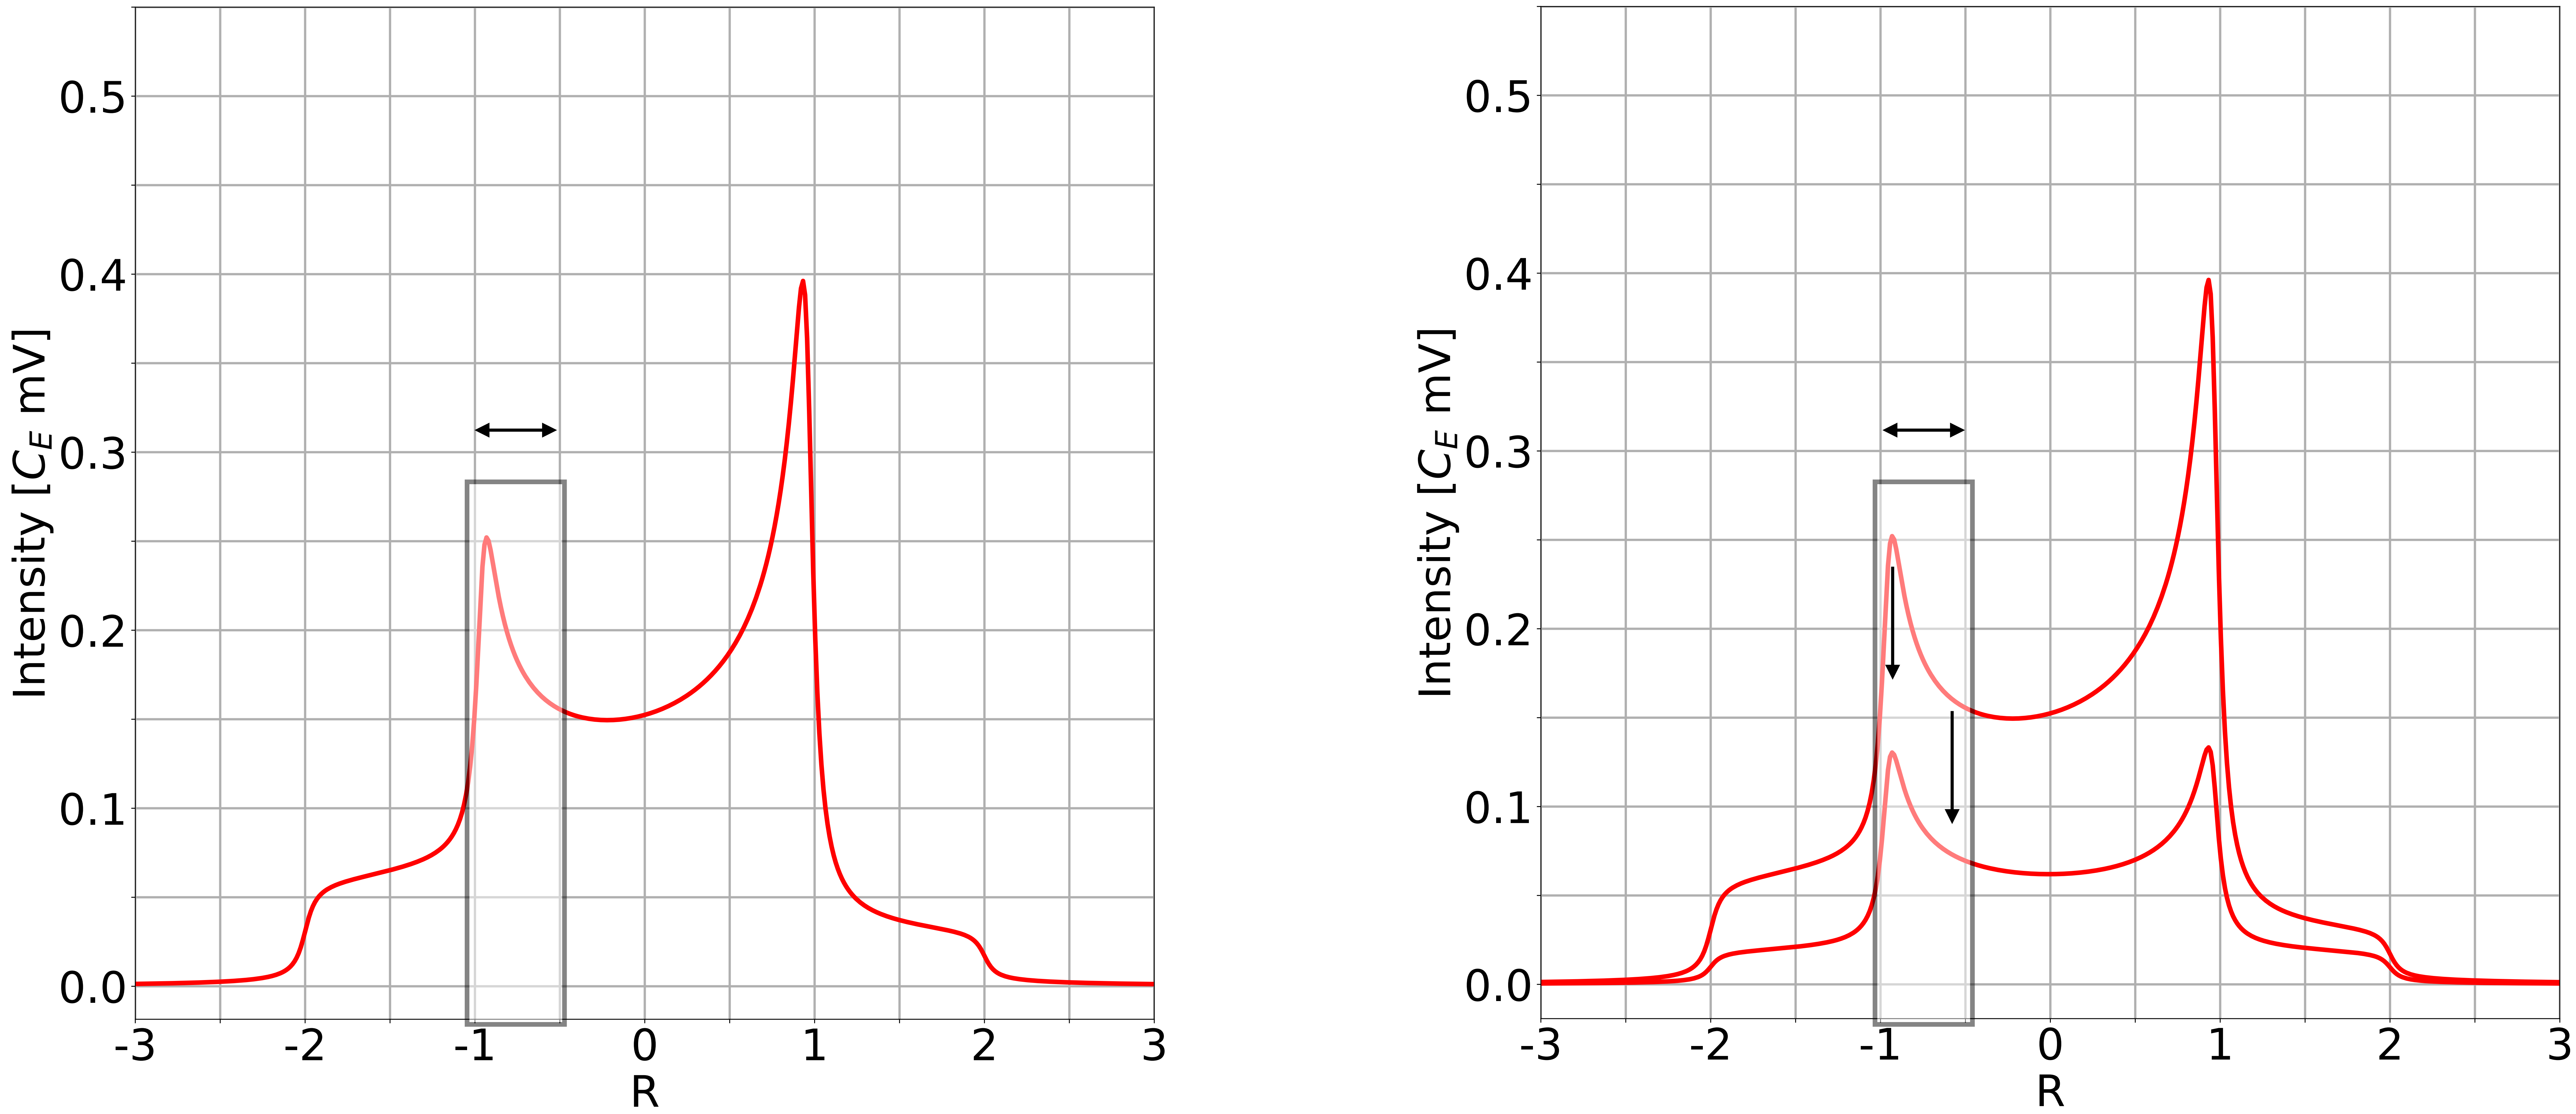

Spin Temperature Consistency

or strategic depolarization

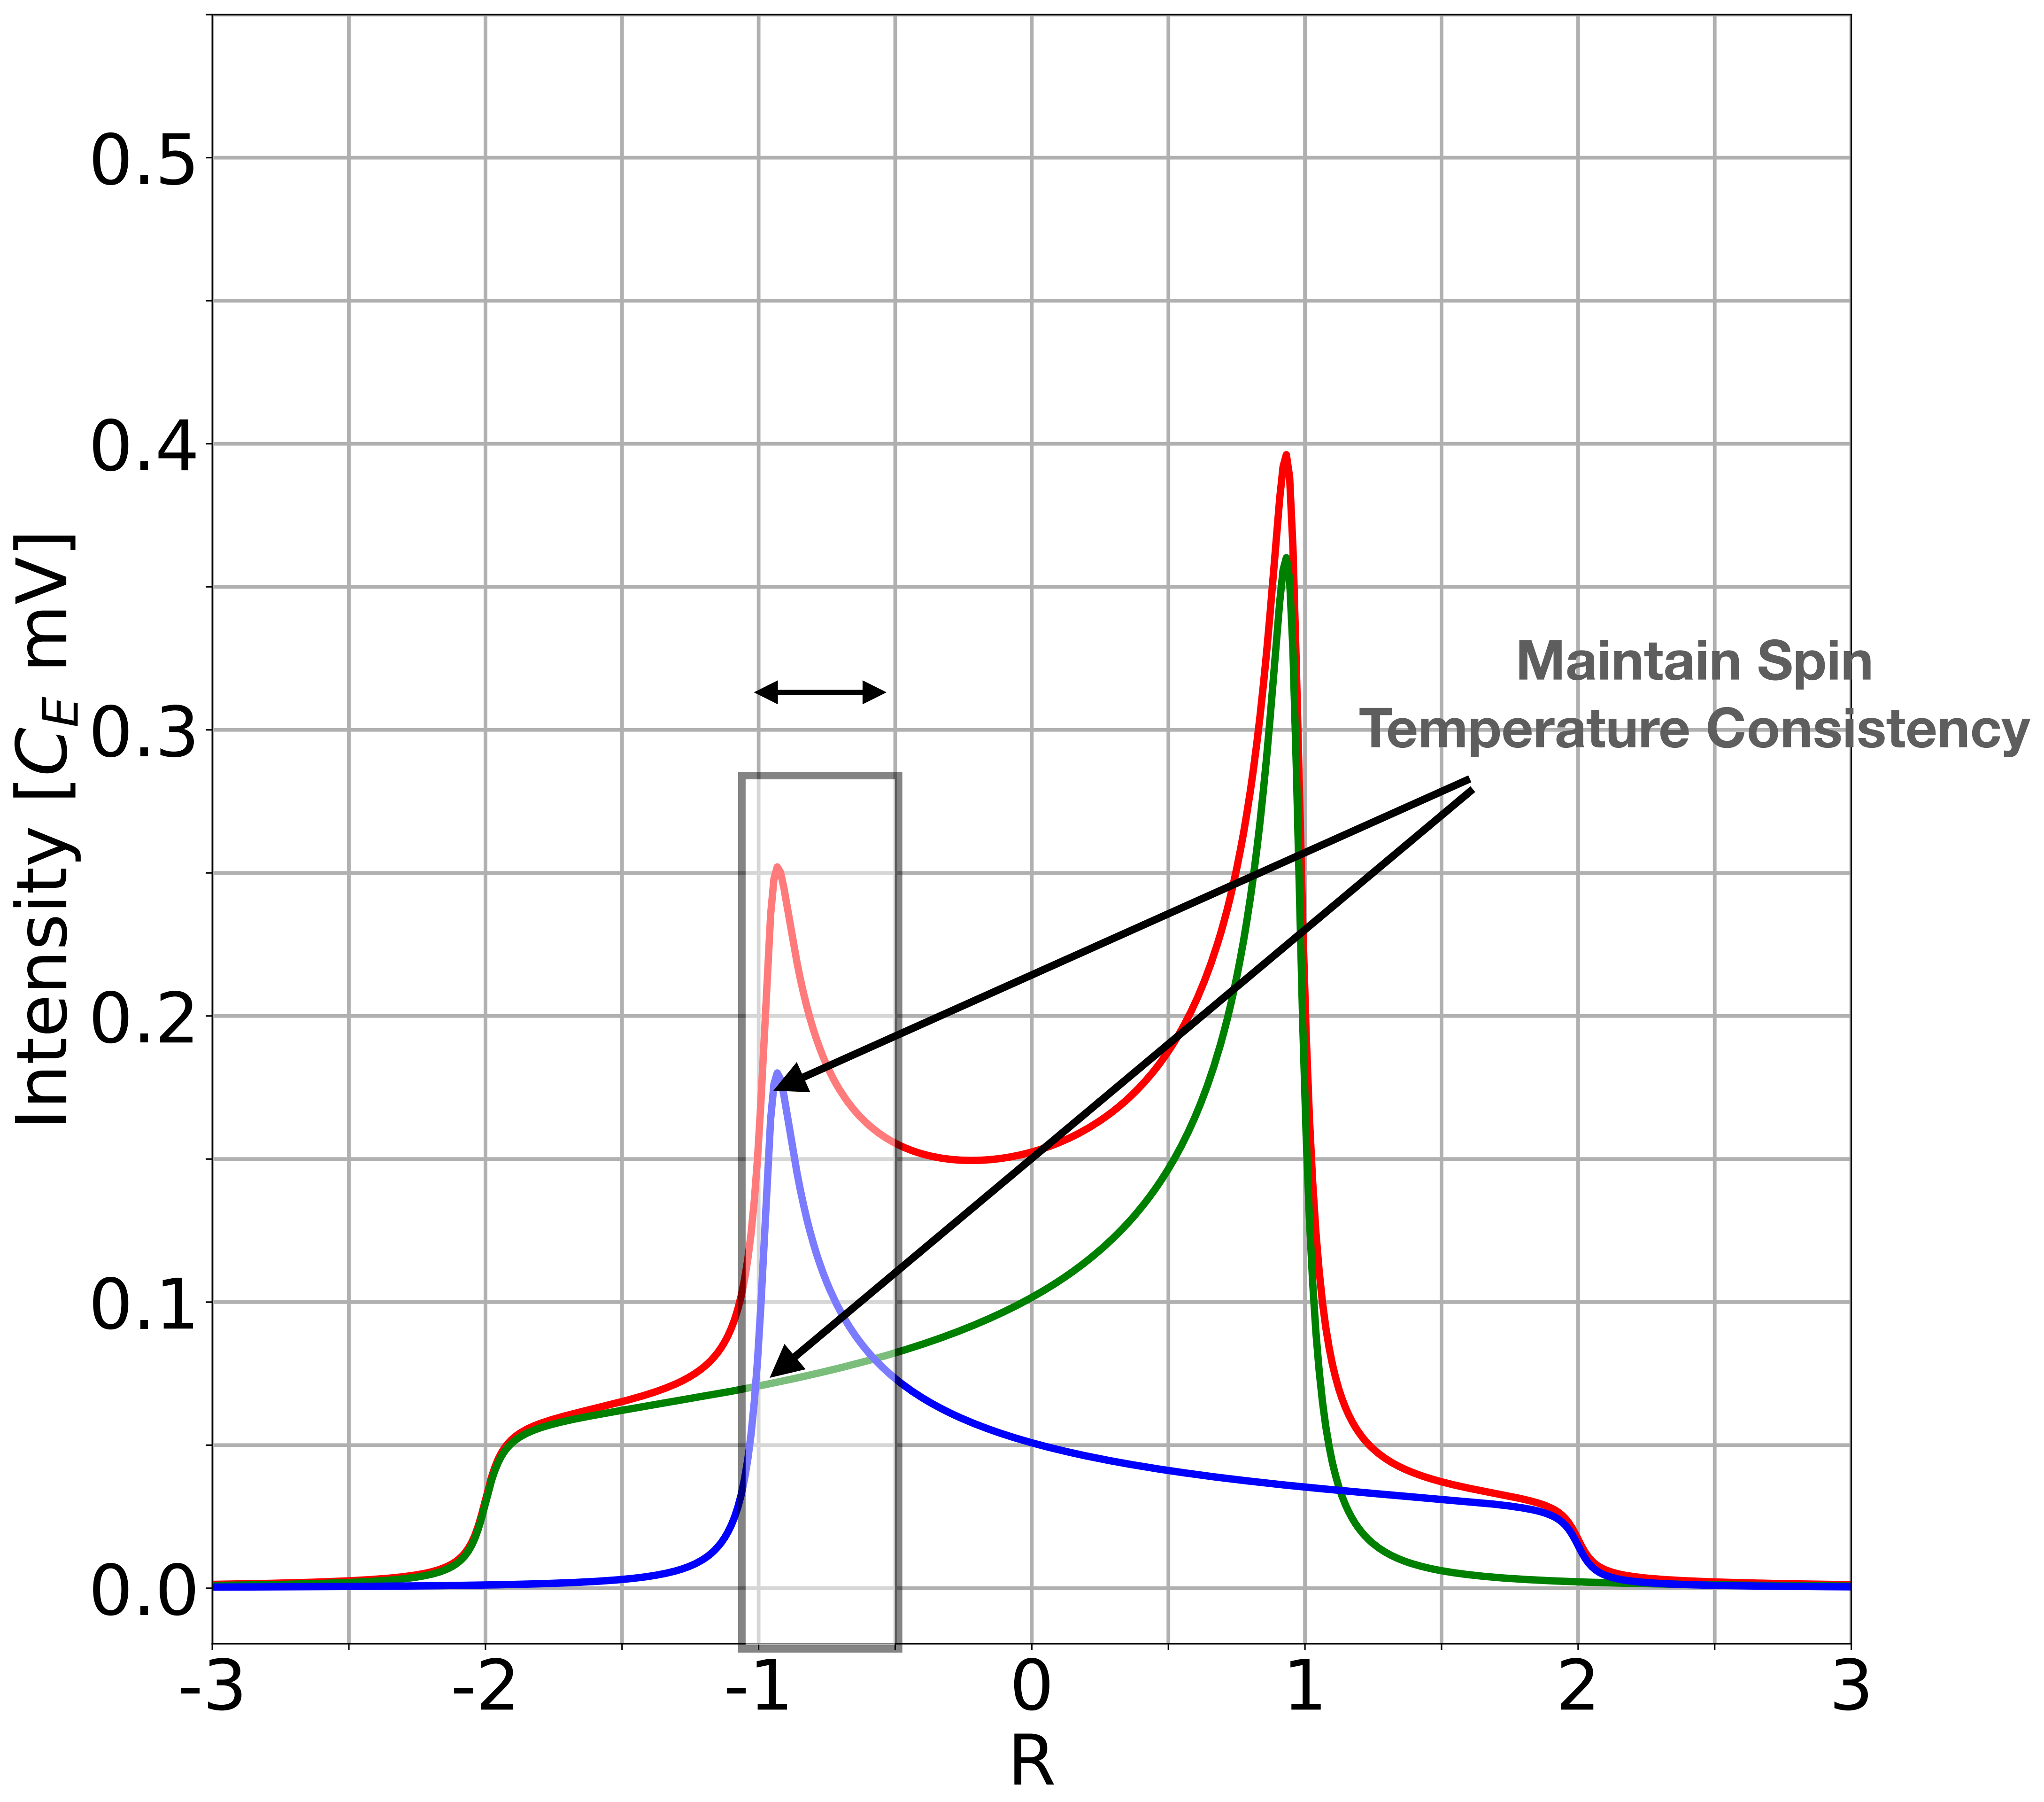

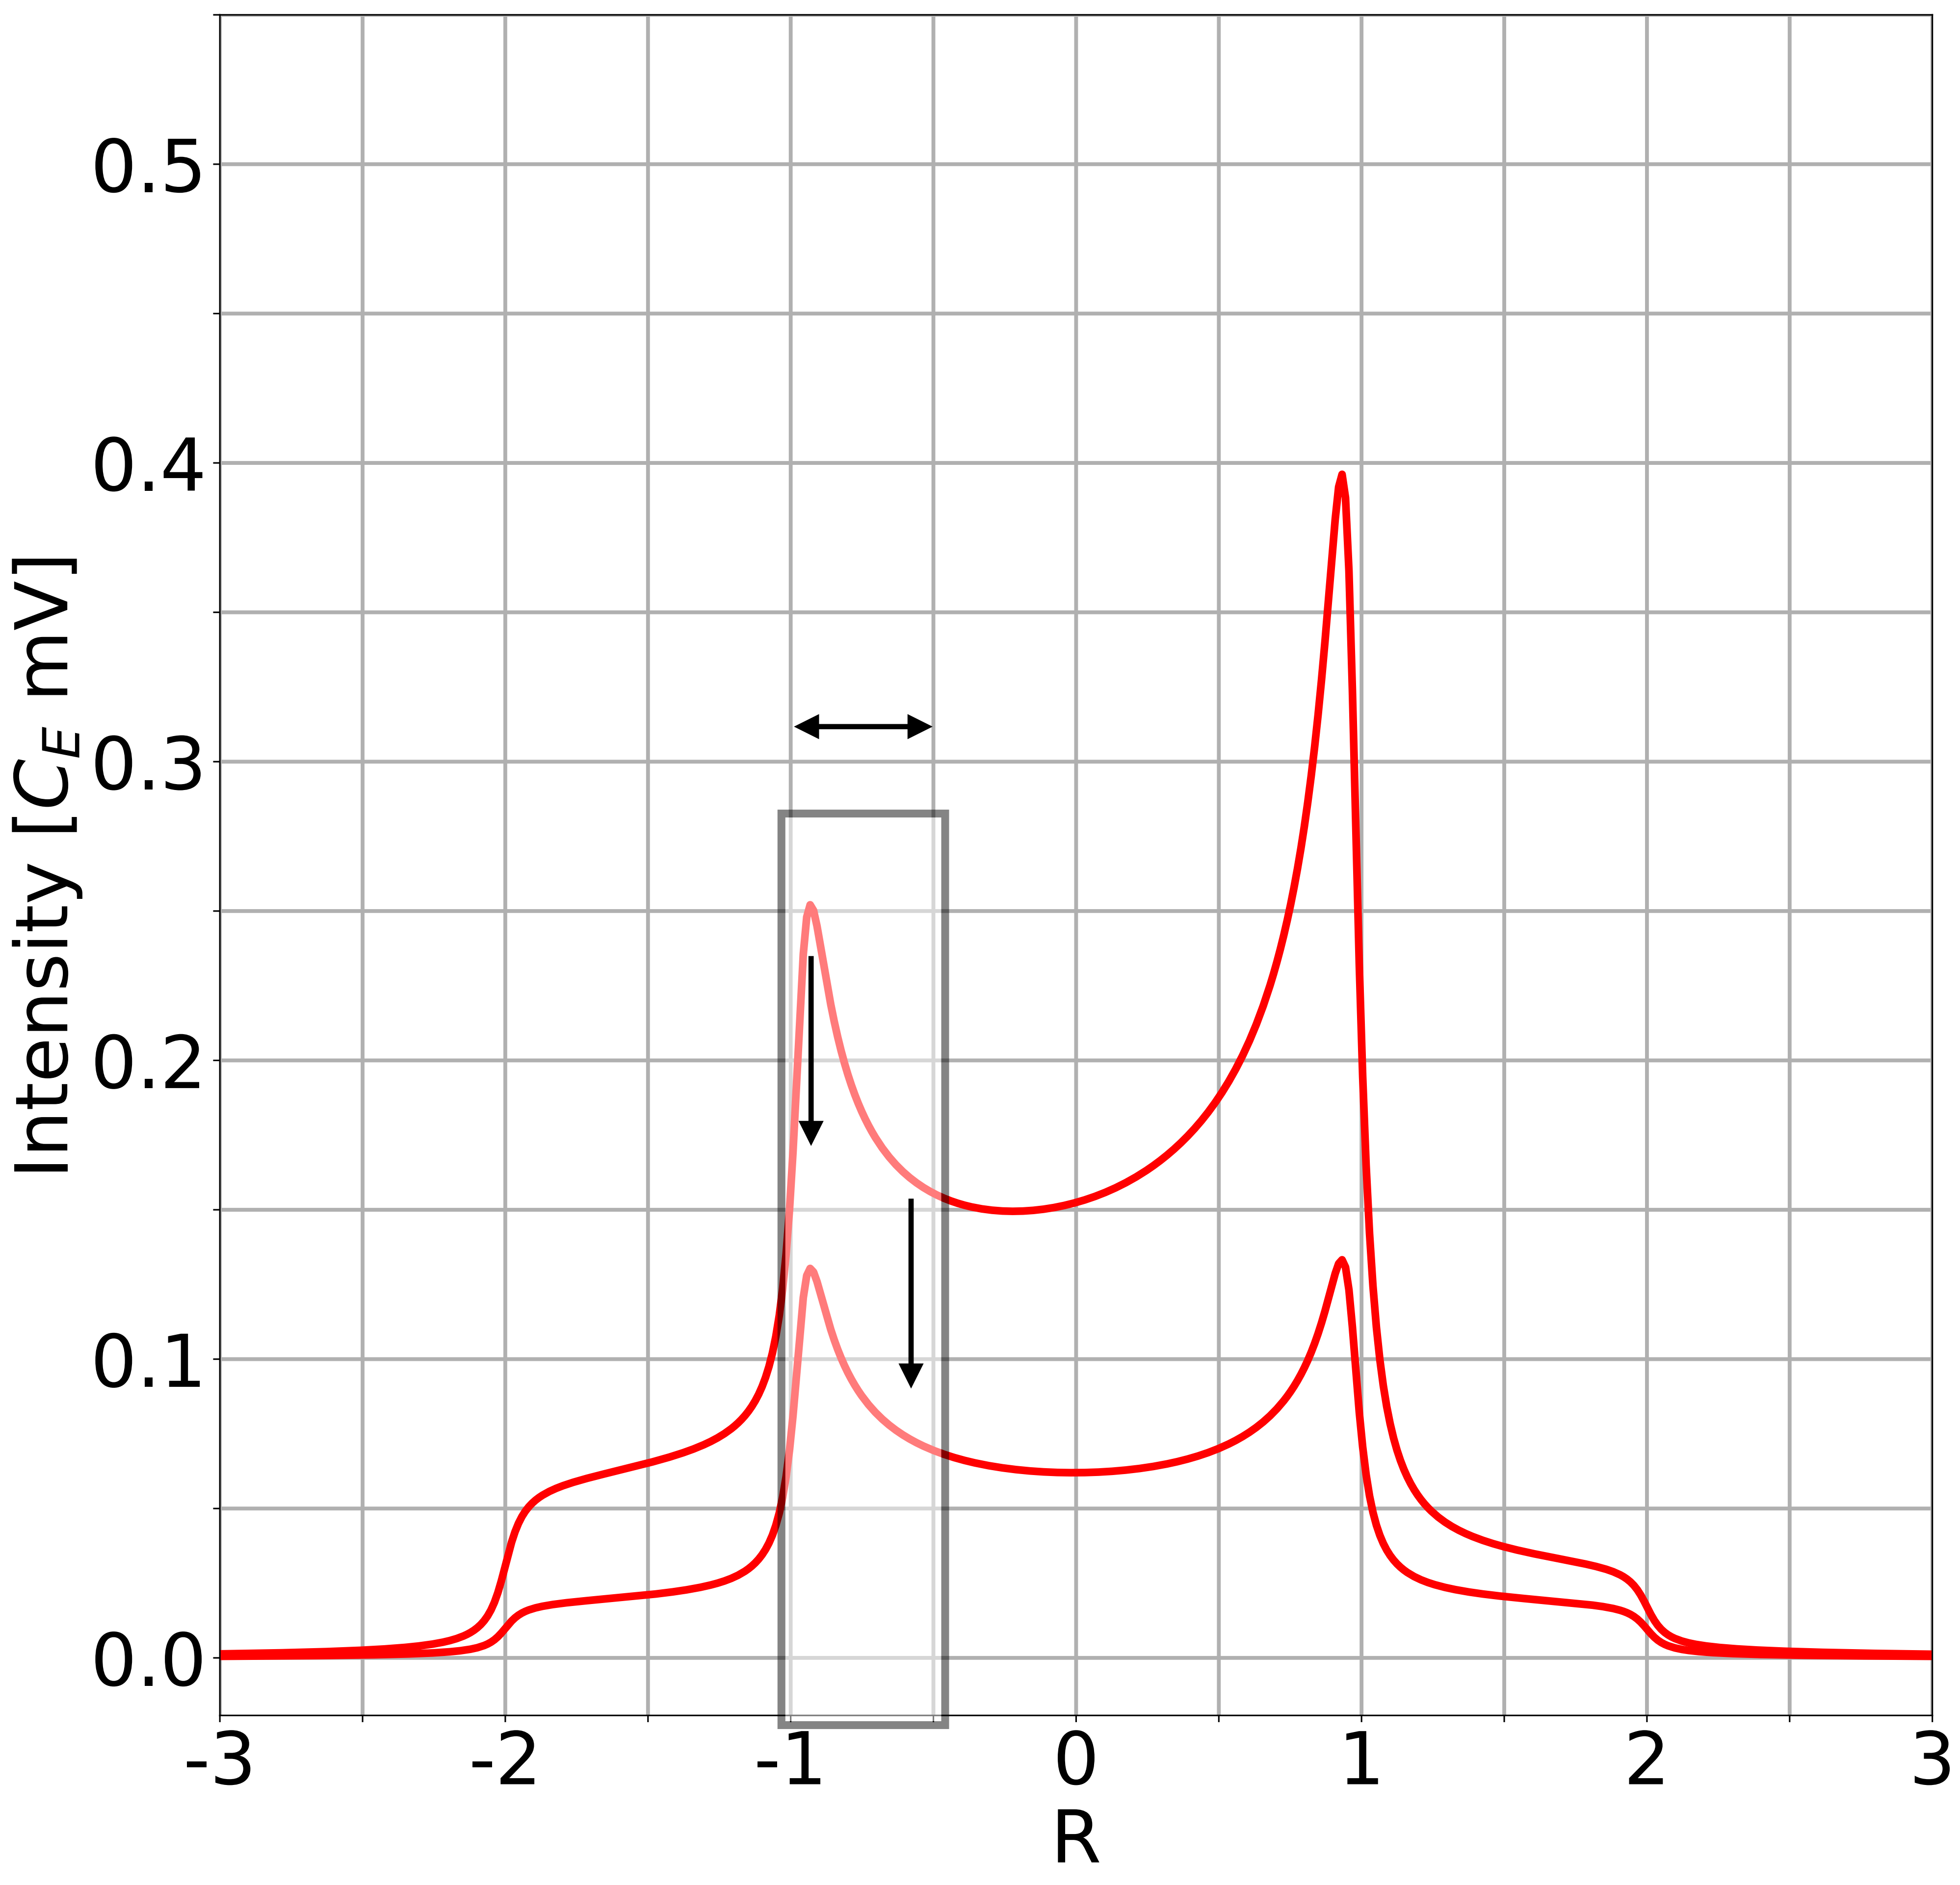

Principle: If you know the sum you know the difference

Rate Response

Or why you can get rid of models

$$\dot{i}_-(-\mathscr{R}) = -2C\xi\rho_0$$
$$\dot{i}_+(\mathscr{R}) = C\xi\rho_0$$

$$\dot{i}_+(-\mathscr{R}) = -2C\xi\rho_+$$
$$\dot{i}_-(\mathscr{R}) = C\xi\rho_+$$

$$\dot{i}_-(-\mathscr{R}) = -\frac{1}{2}\dot{i}_+(\mathscr{R})$$

$$\dot{i}_+(-\mathscr{R}) = -\frac{1}{2}\dot{i}_-(\mathscr{R})$$

$$\boxed{A_{gained} = \frac{1}{2}A_{lost}}$$

https://doi.org/10.1016/j.nima.2023.168177

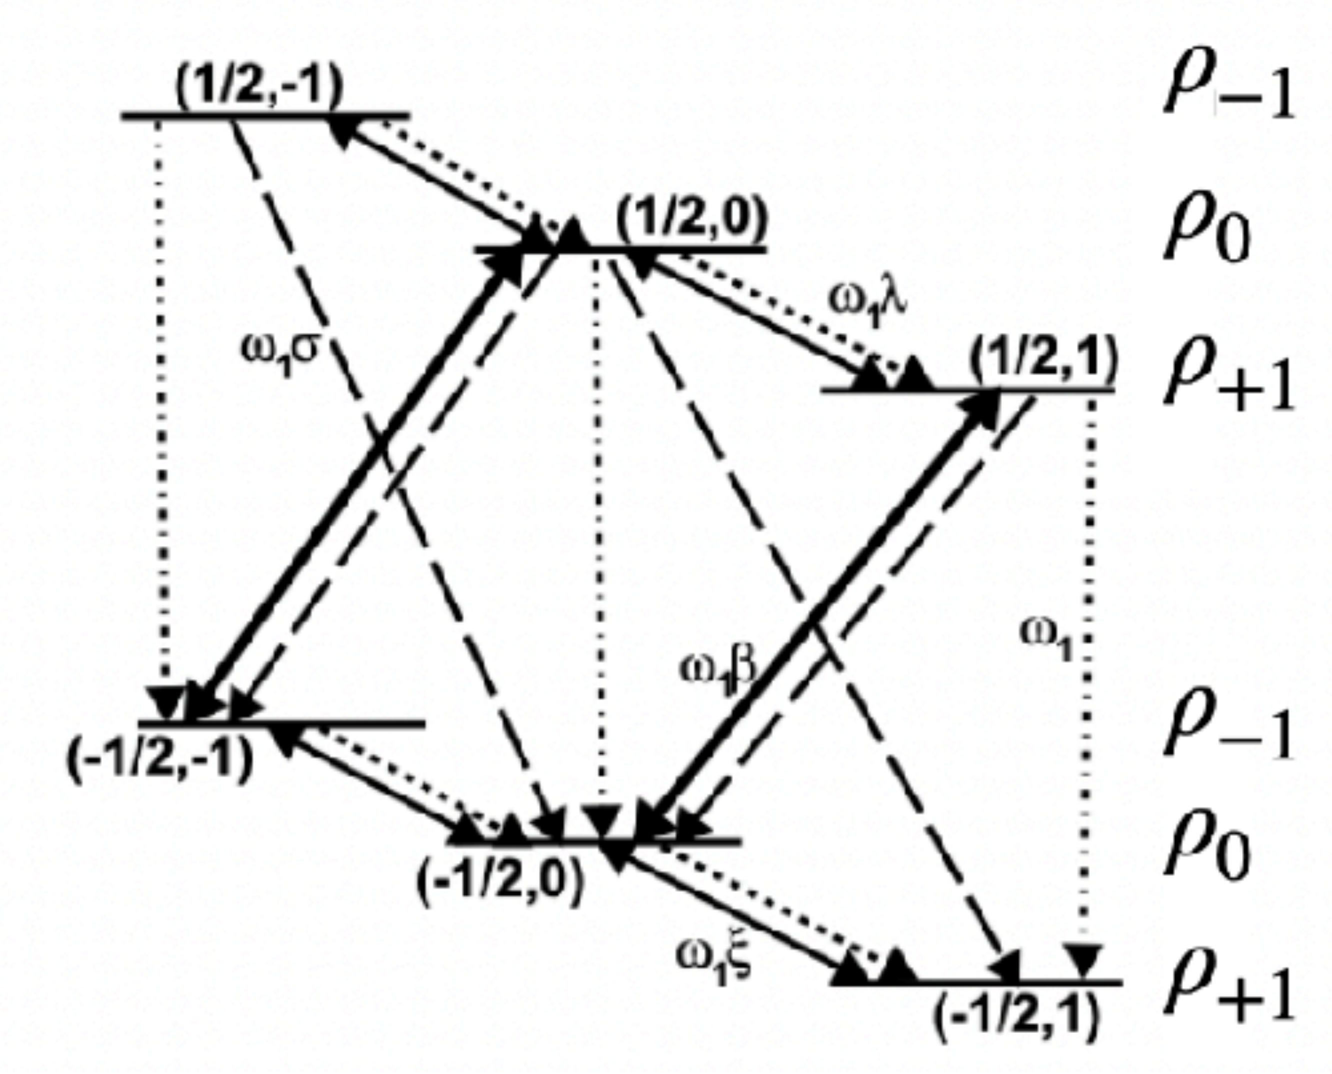

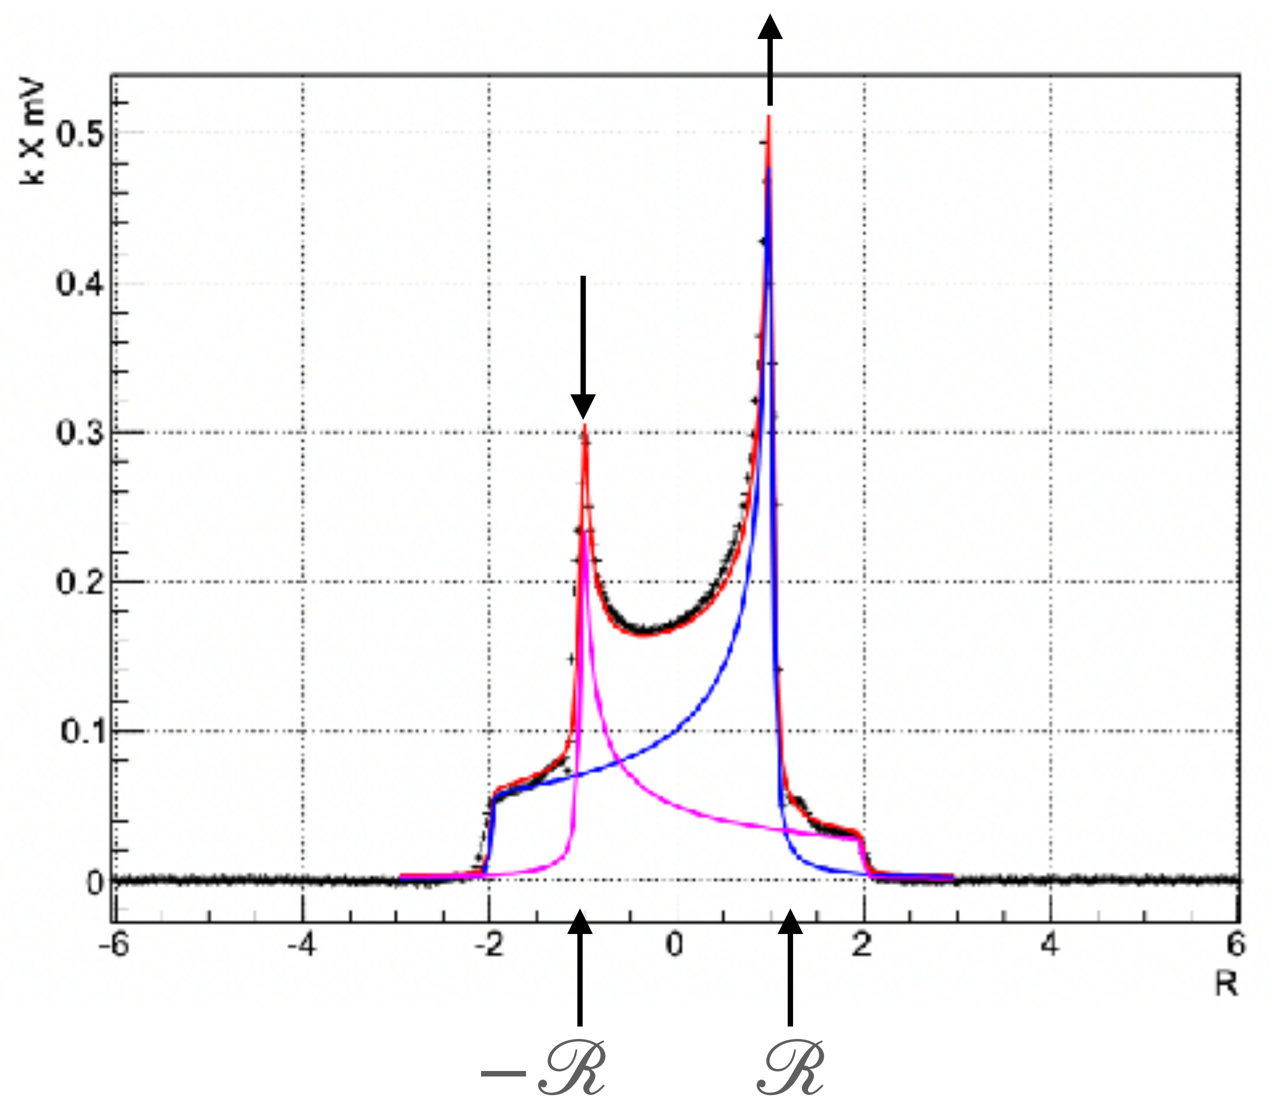

Putting These Conditions Together

Simple Measuring Tools

- The difference (Q) in intensities can by easily calculated using area dependent Boltzmann
- Apply $A_{gained} = \frac{1}{2} A_{lost}$
- Configure for any vector polarization and the particular RF region
- You're Done!!

Universally True
Any lineshape
Any material

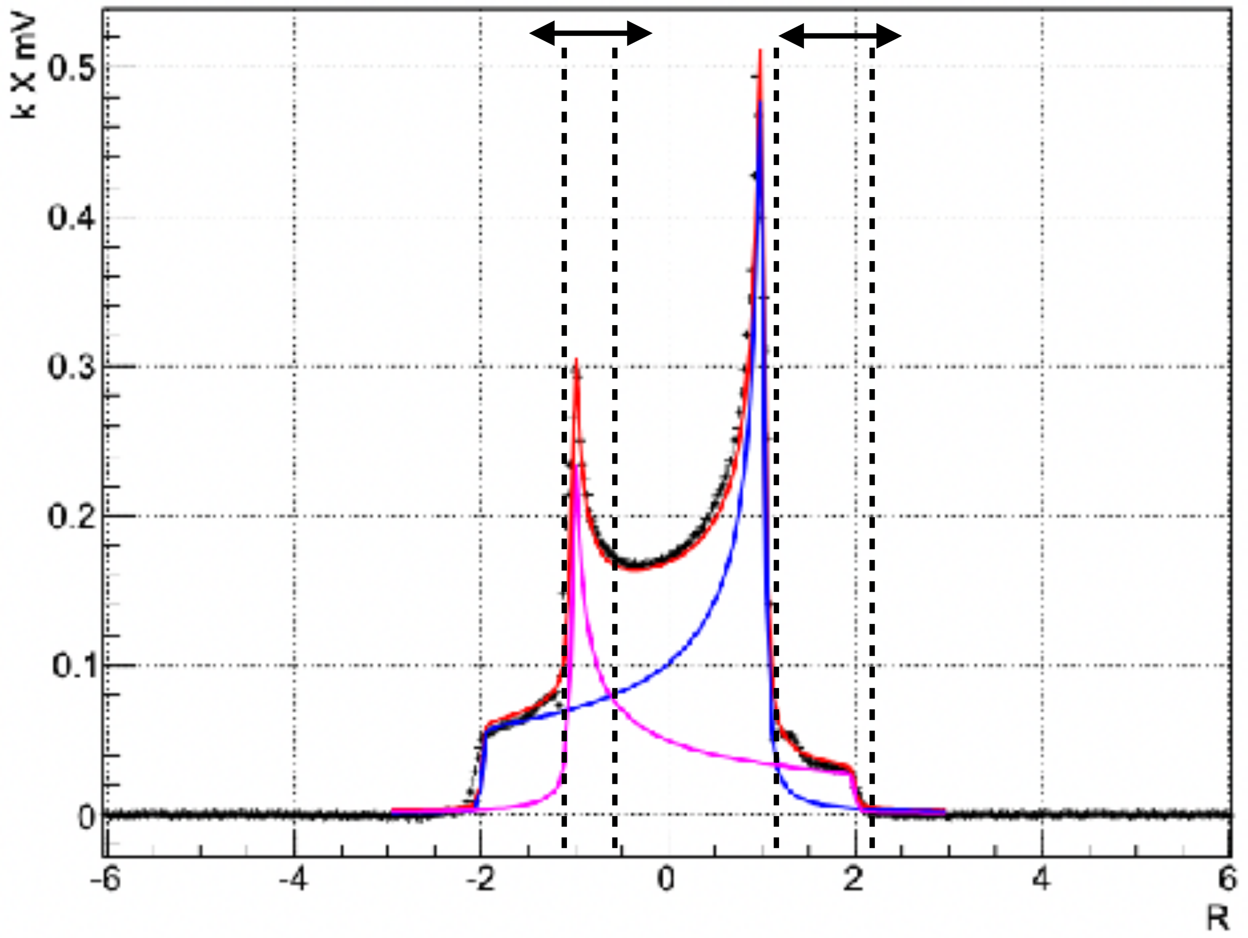

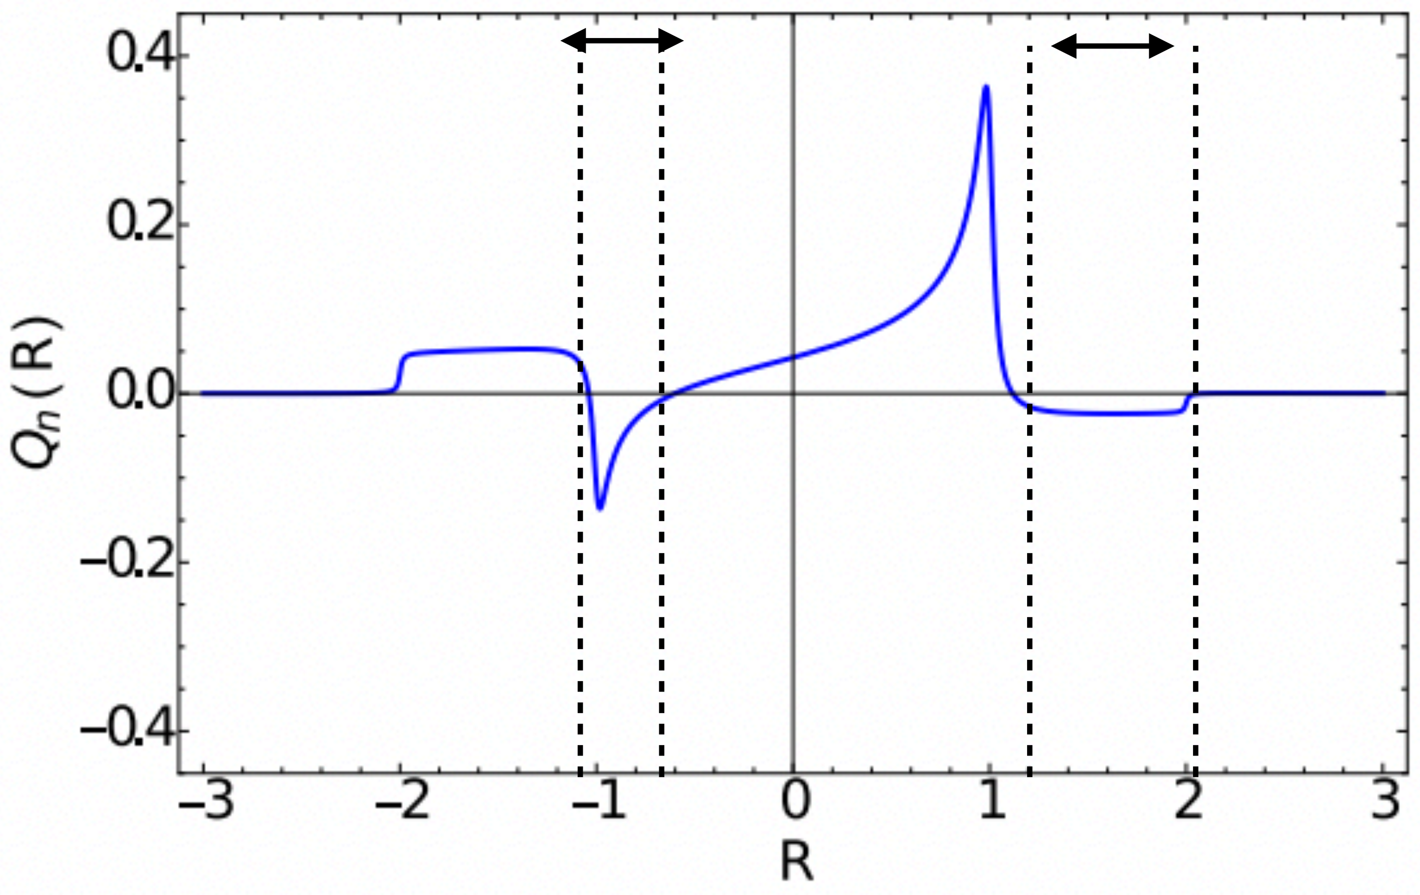

Caveats

What is exact and what is approximation

- Everything just laid out is exact for any polarization mechanism, any line shape, and any material and so **not model dependent**
- Everything just laid out is in reference to longitudinal (spin-lattice) relaxation pathways and not transverse (spin-spin, like spin diffusion)
  - To take into account the transverse relaxation pathways one needs to fit the hole using a Voigt (convolution of Gaussian and Lorentzian)
  - These fits are sensitive to Q-factor of coil, degree of tuning and matching of RF circuit, amplification parameters, and transverse relaxation of material

Addressing Transverse Pathways

Voigt Profile

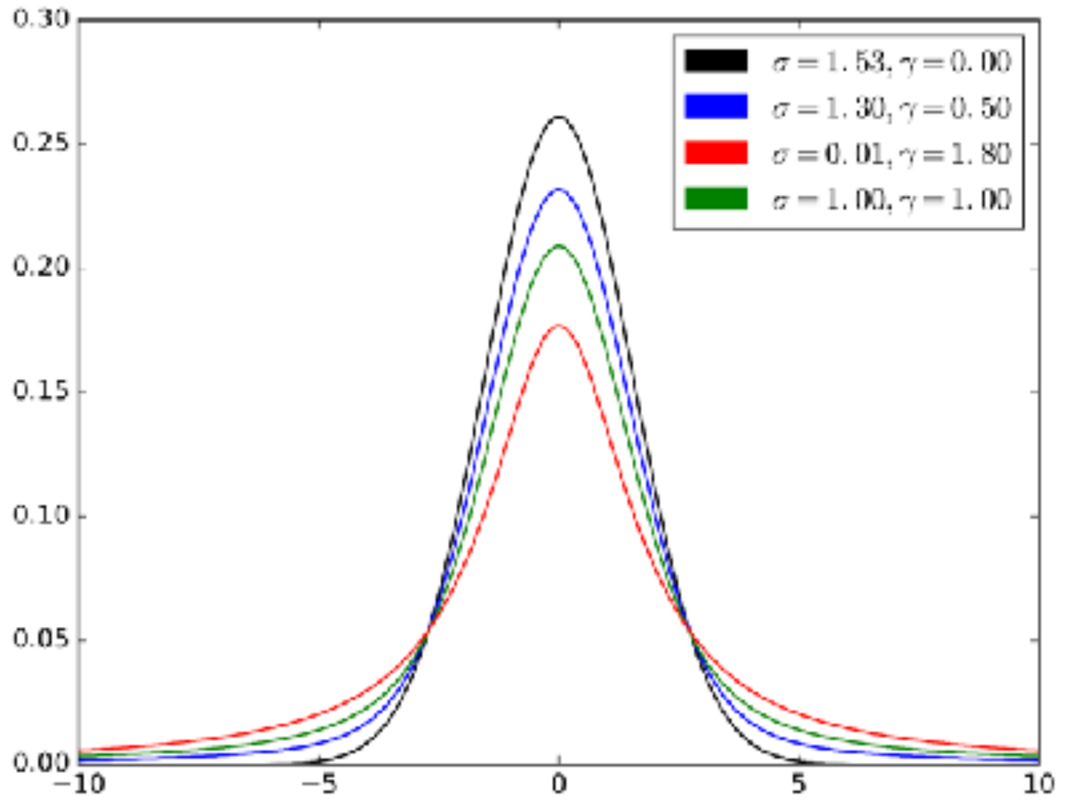

$$V(x;\sigma,\gamma) \equiv \int_{-\infty}^{\infty} G(x';\sigma)L(x-x';\gamma)\,dx'$$

$$G(x;\sigma) \equiv \frac{e^{-x^2/(2\sigma^2)}}{\sigma\sqrt{2\pi}}$$

$$L(x;\gamma) \equiv \frac{\gamma}{\pi(x^2+\gamma^2)}$$

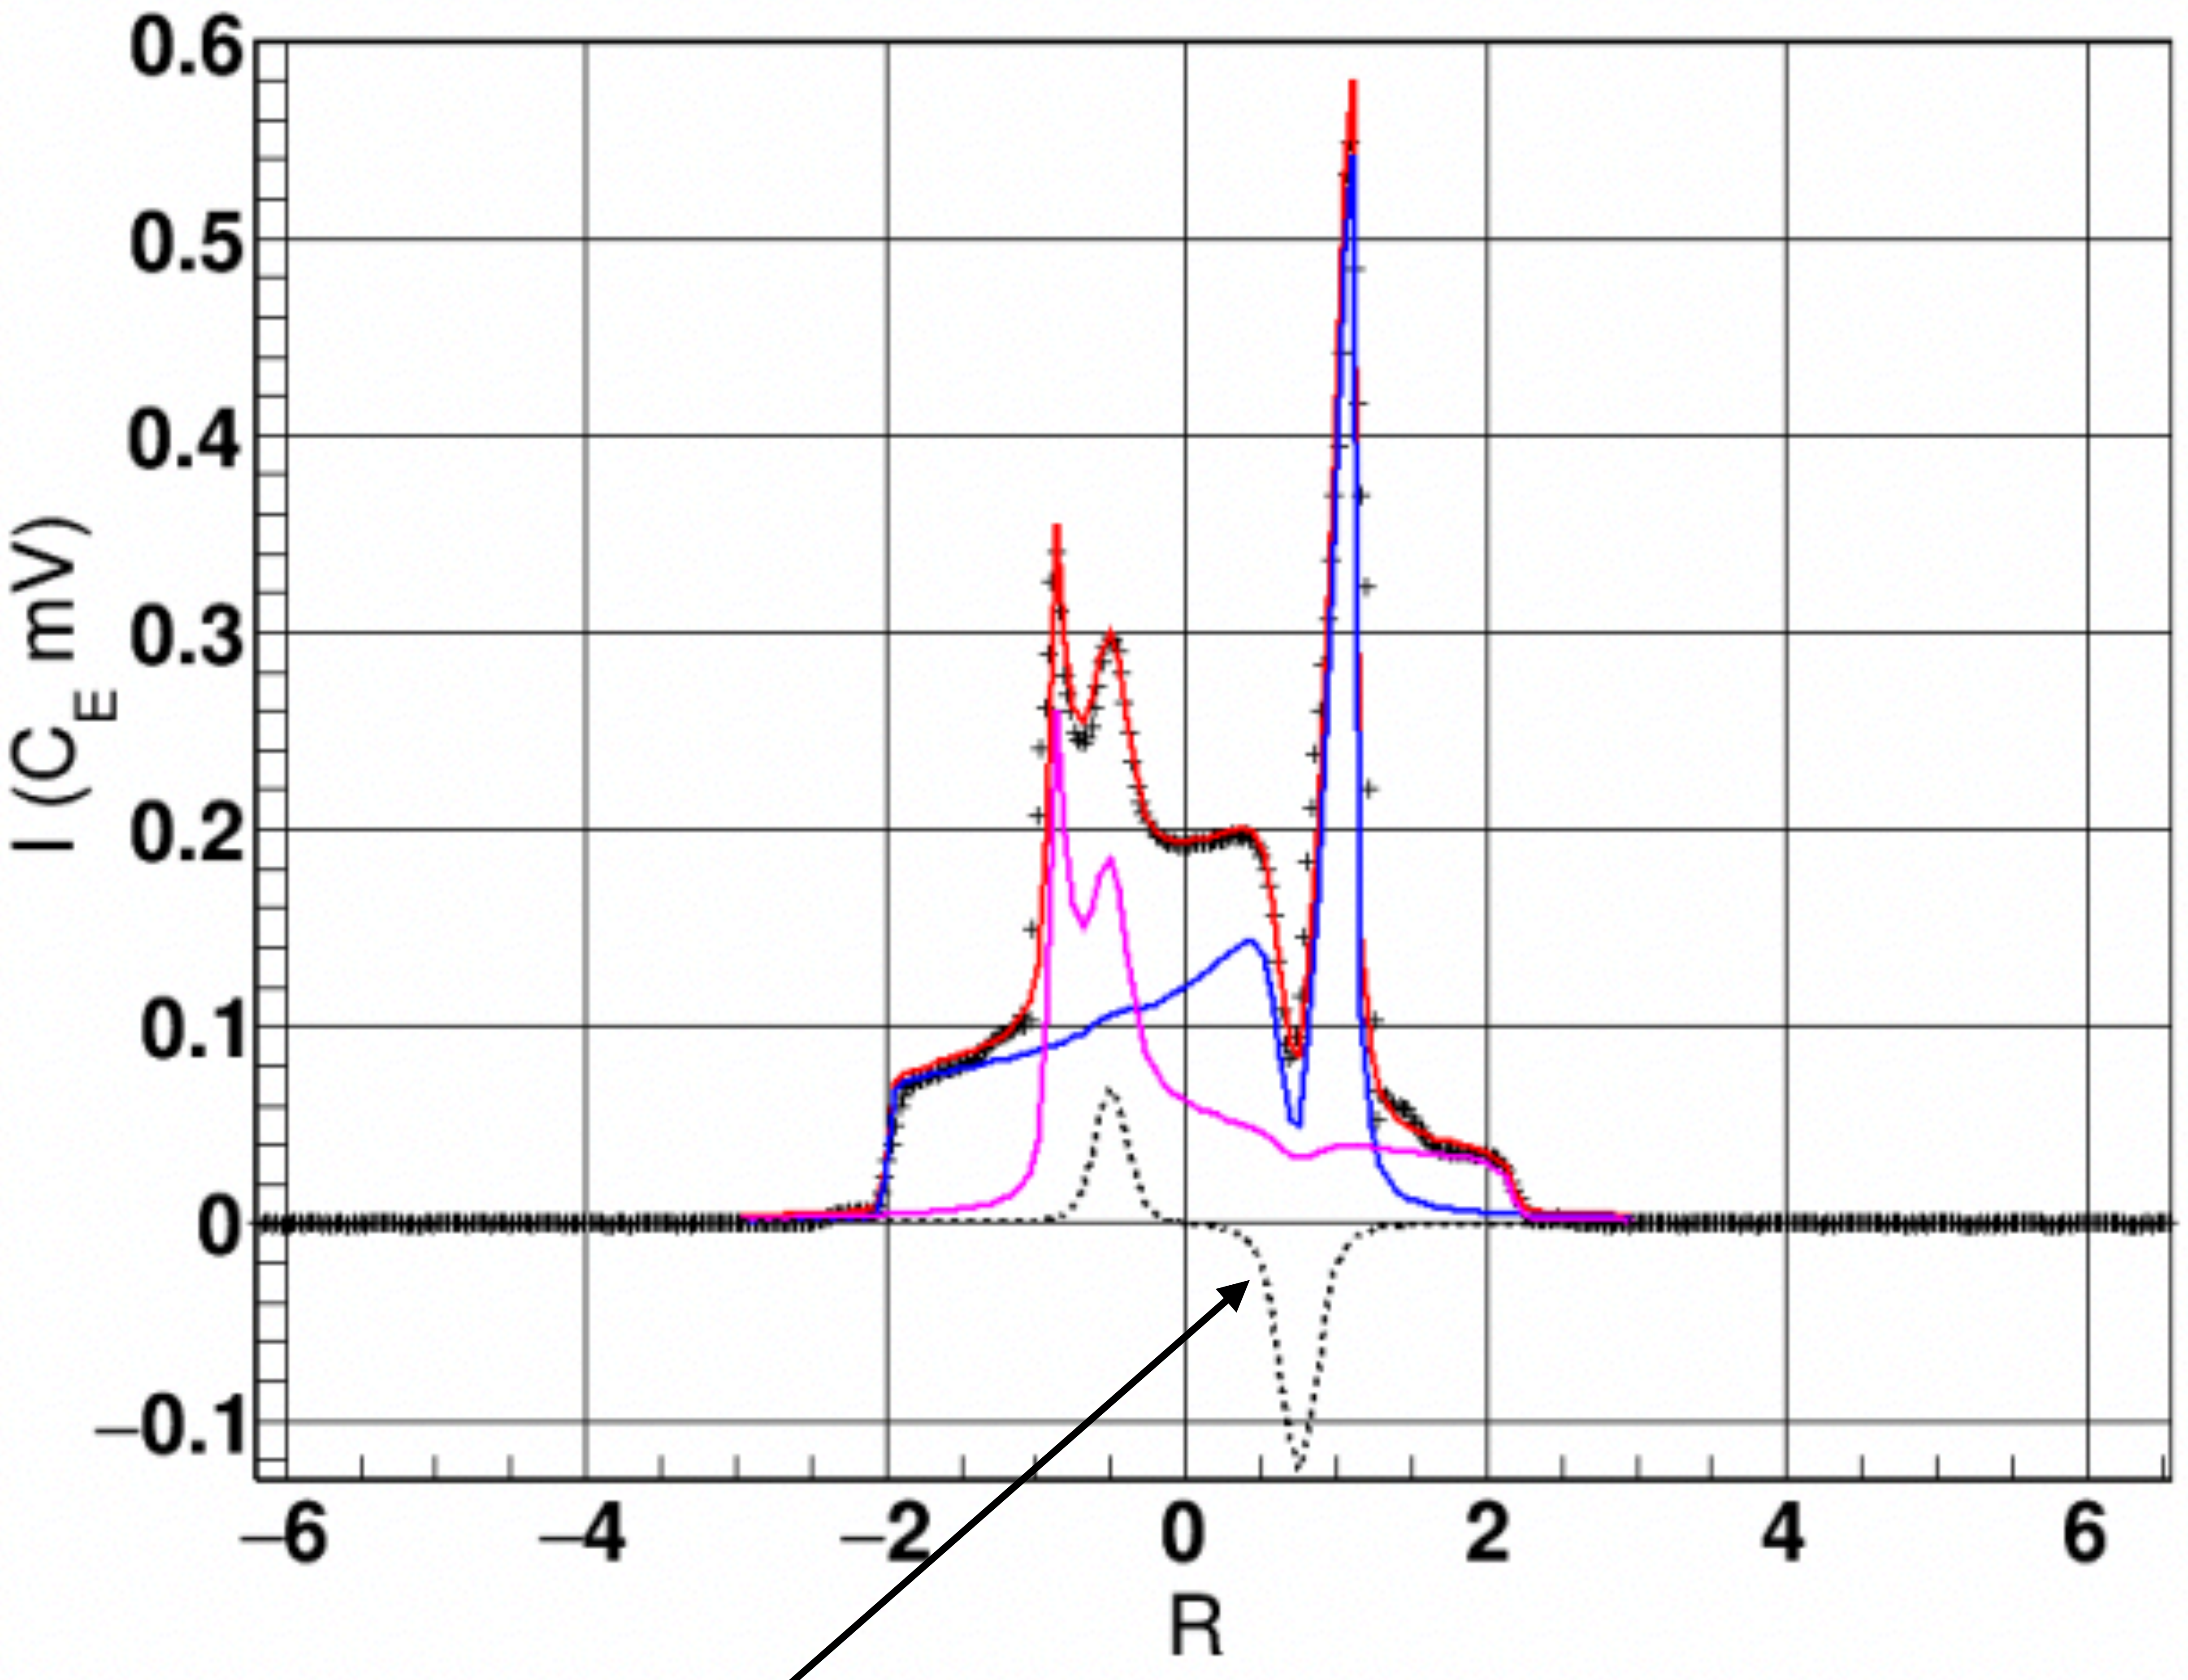

Dotted line is signal before ss-RF minus signal after ss-RF

Practical Aspects of Doing This with NMR

ss-RF Q-meter based prototype

- Set the single sweep
- RF switch to turn off Q-meter and run ss-RF
- Then run and change ss-RF power/duration
  - But measurements and response must be fast or errors will overwhelm

Possible without a new style of NMR

Yes!!

General Idea

For ss-RF NMR

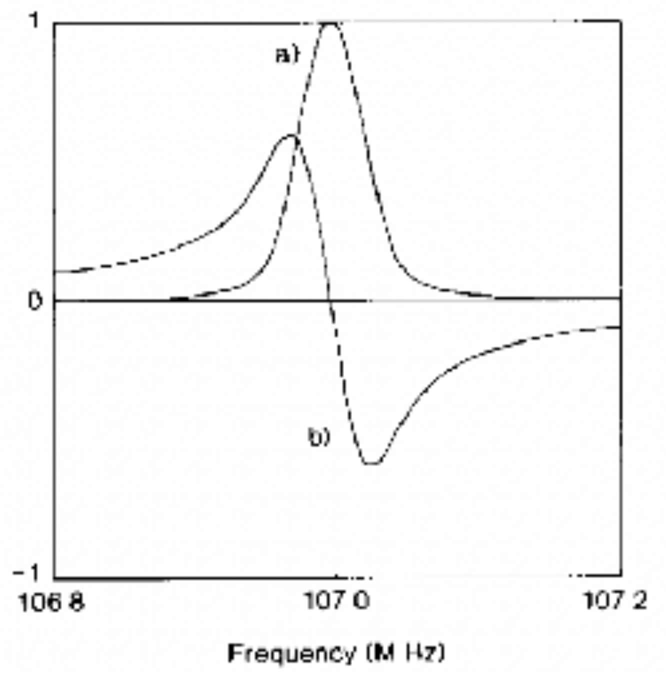

Design Needed

- Absorption/Dispersion
- Fast Phase Sampling
- ss-RF power profile
- Separation of RF range

A. Measure A/D 1000 phases

B. AI inference (streaming covariance)

C. ss-RF Enhance Q per sweep

D. goto A.

Standard Error in
Experiments: 4%
Theoretical Limit: 0.8%

Original Liverpool design

- Phase Sensitive detector
- Constant Current
- Non-destructive (low amplitude)
- Full spectral range

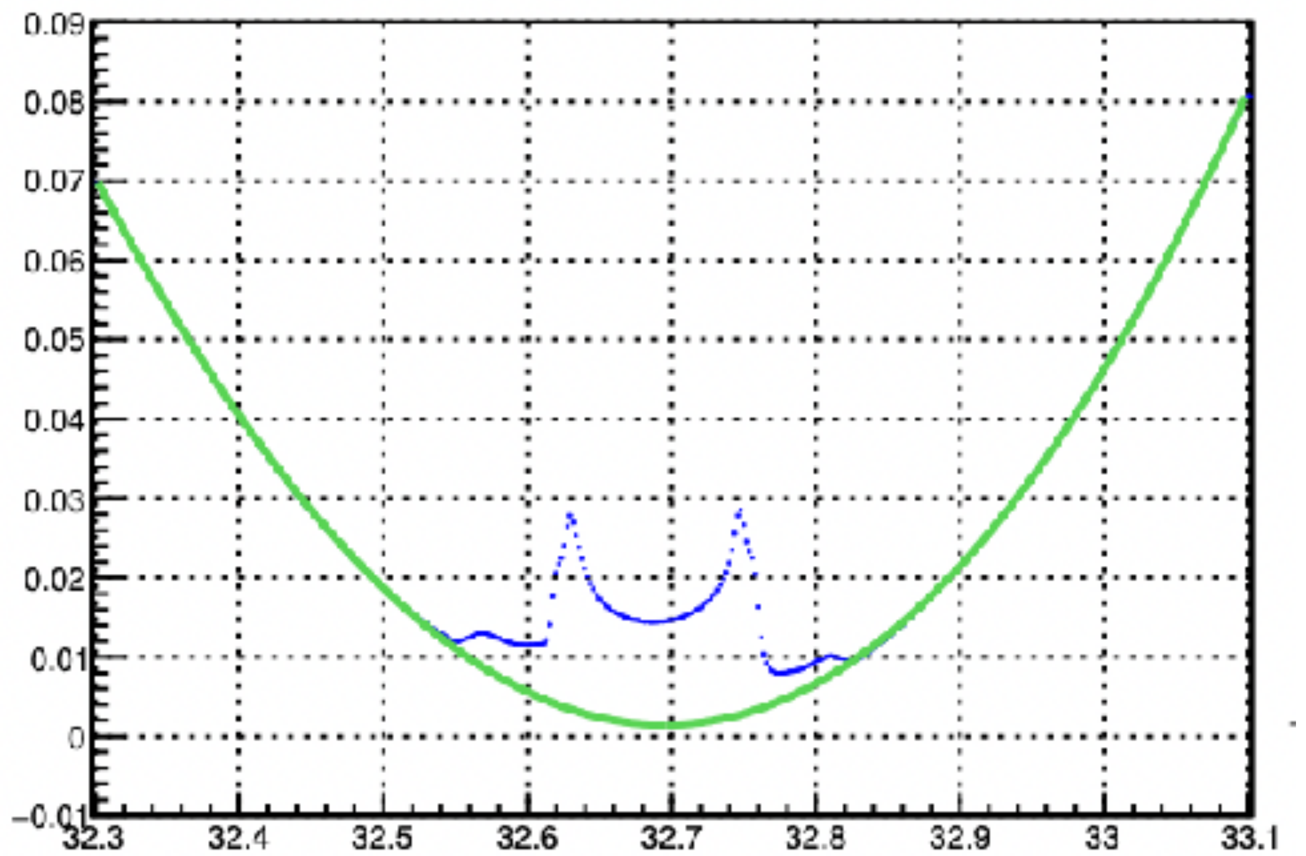

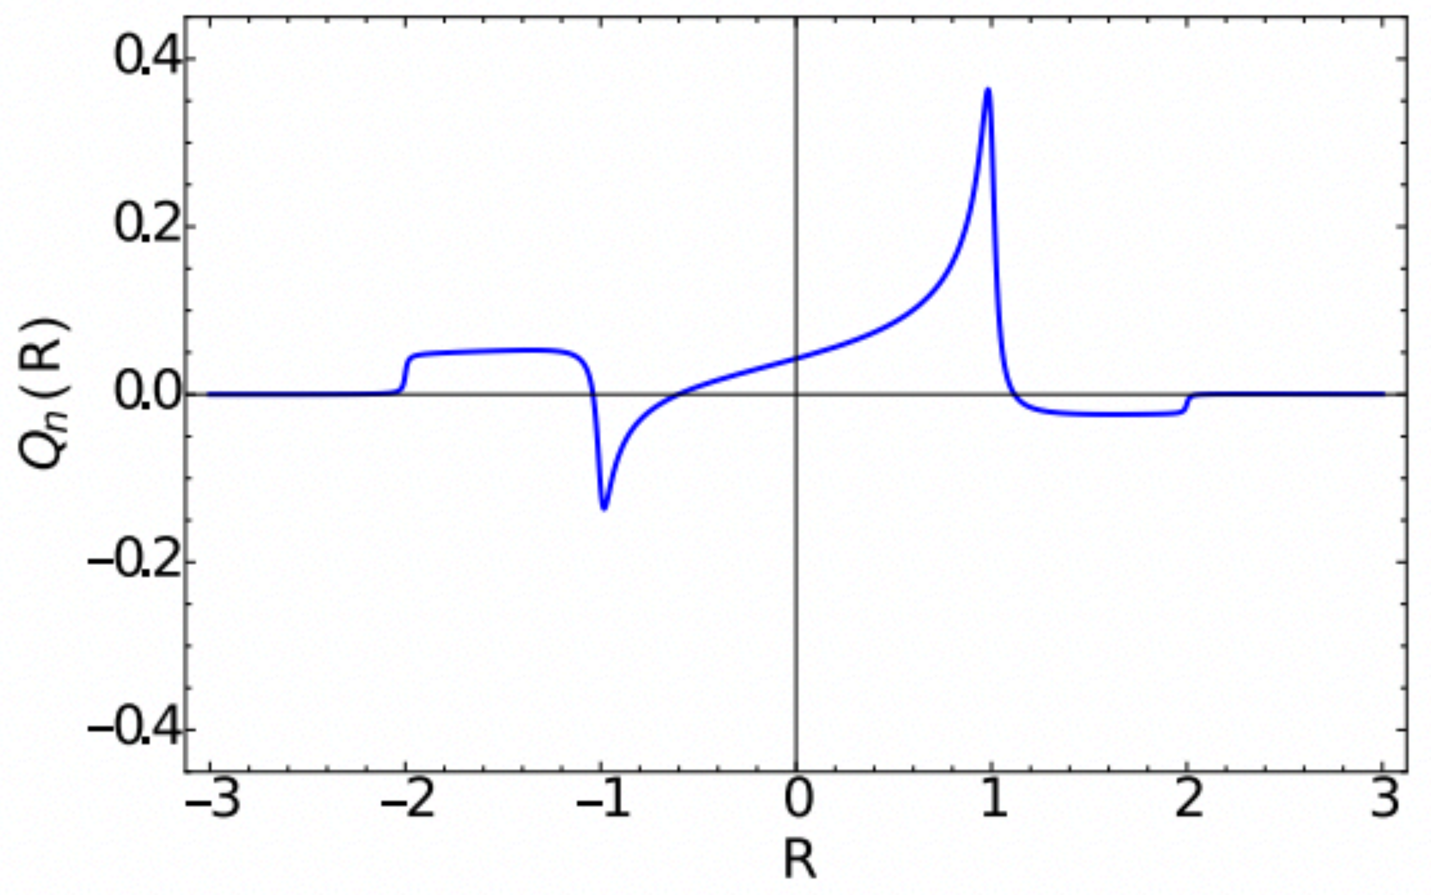

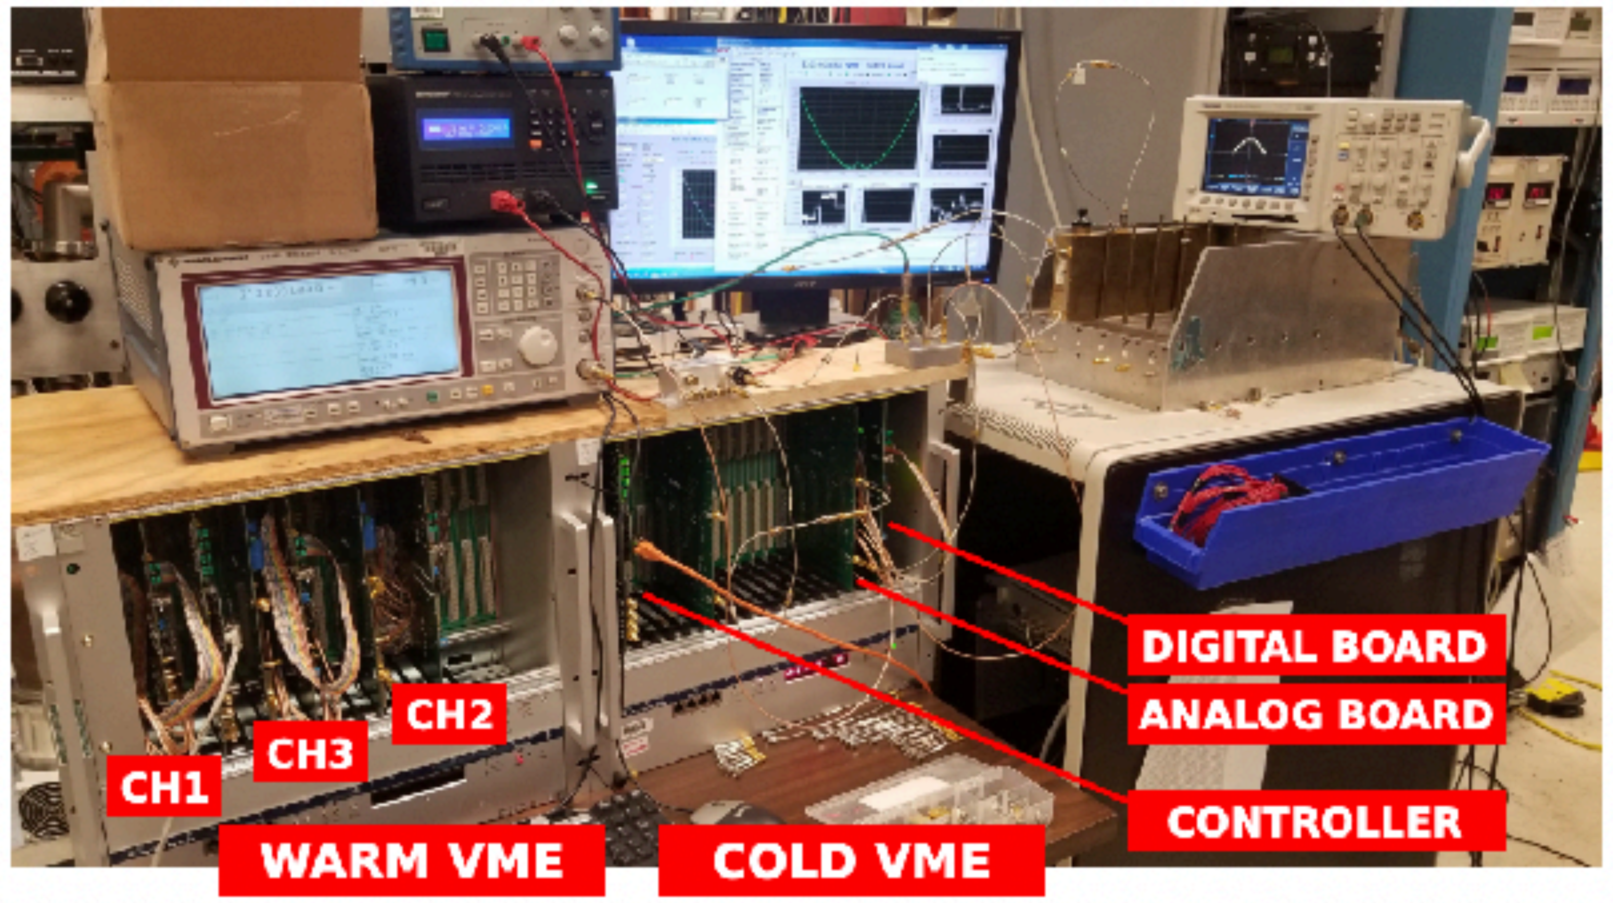

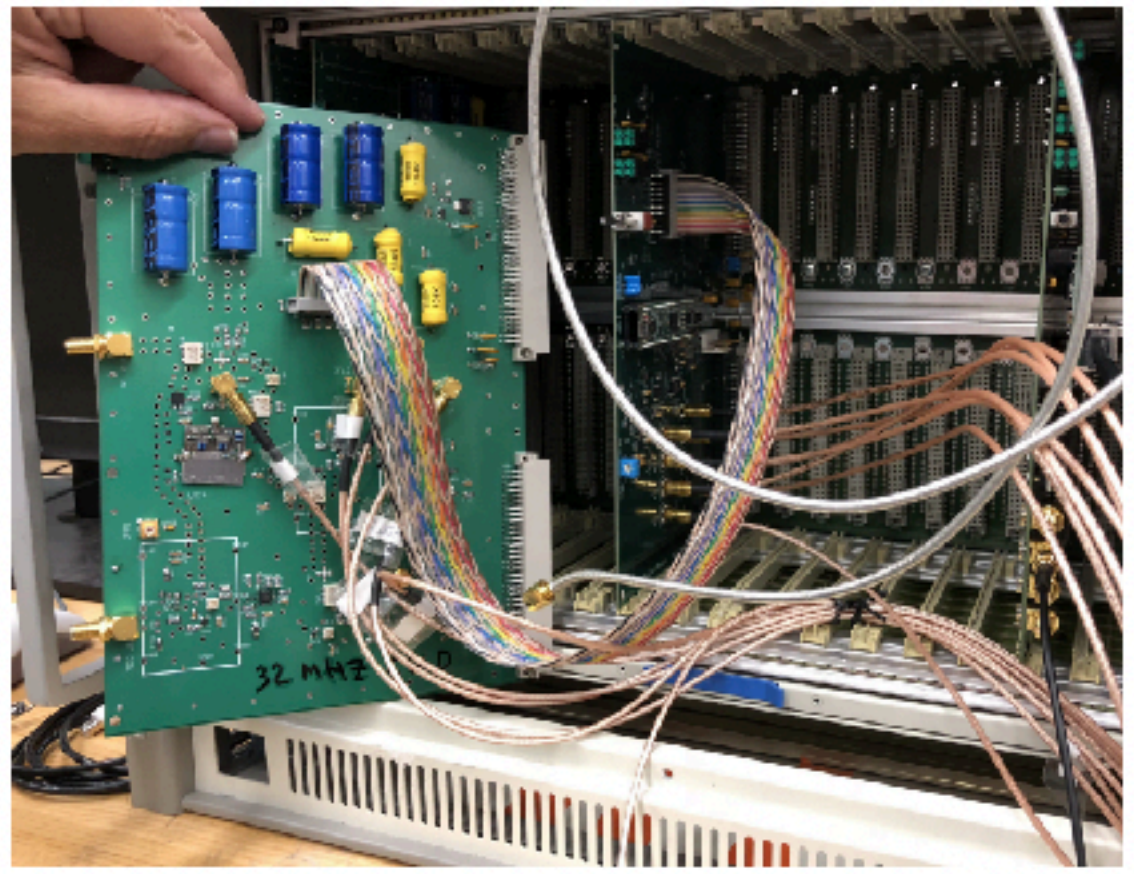

Status of New ss-RF based NMR

As of now

- Initial prototype used in cooldowns (large errors and no AI)
- Generic first round LabView based software/simulation
- New AI-GPU based extraction (Devin Seay)
- High-level design complete
- Applied for NSF funds to build
- Acquired funding
- Acquired hardware
- Constructed
- Commissioned

Measurements of Tensor Enhancement

Experimental results (all with irradiated d-Butanol)

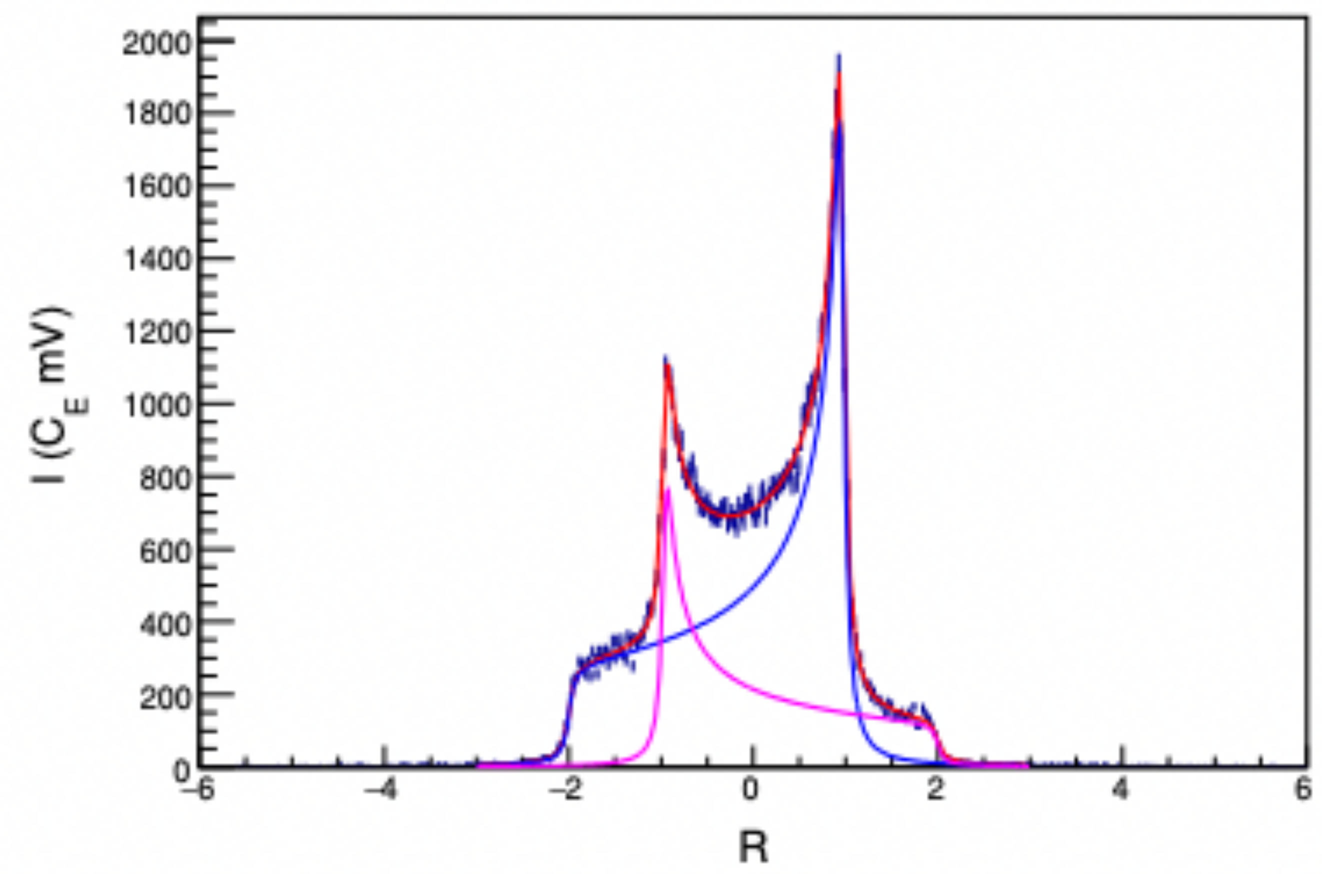

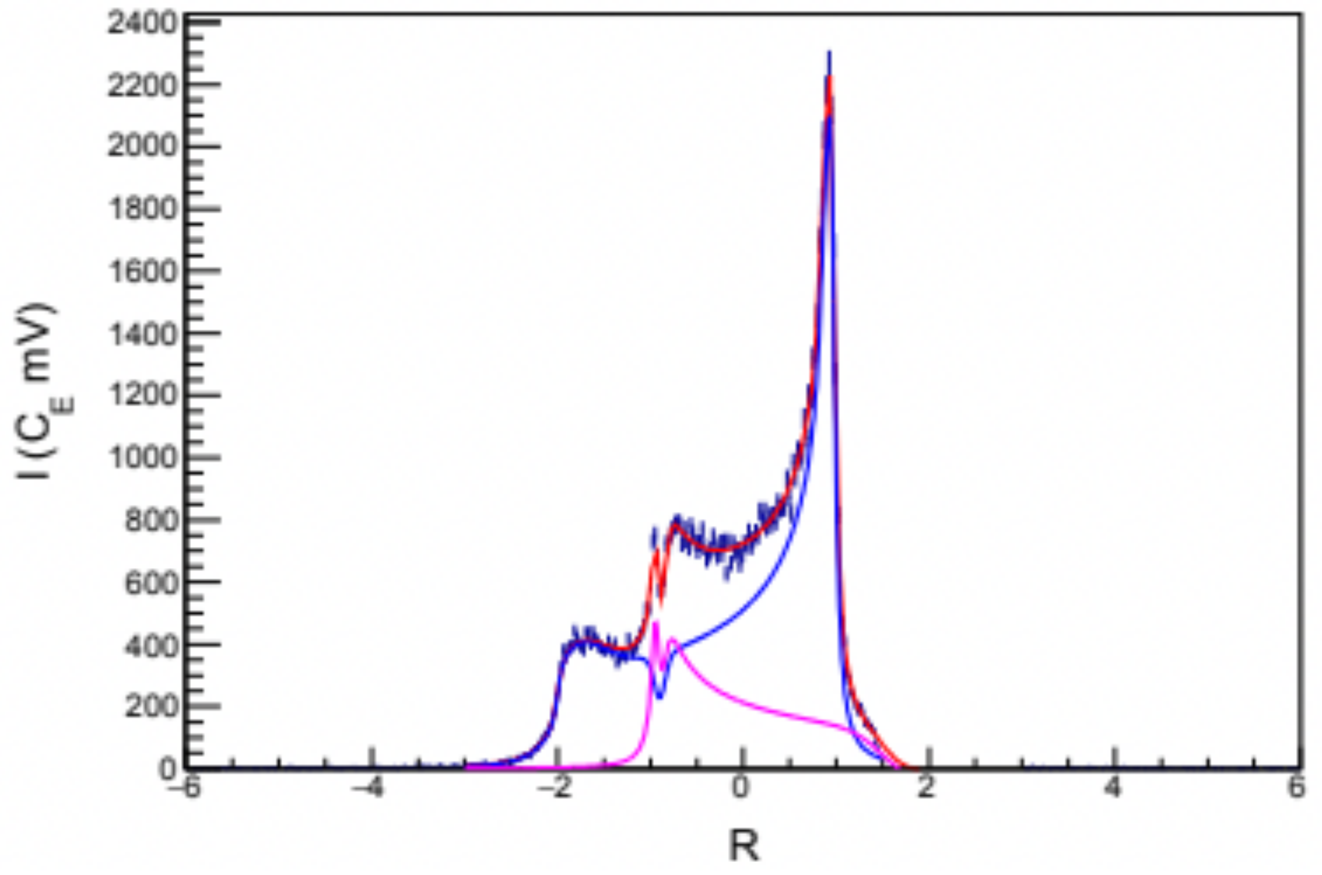

ss-RF Enhanced Measurements					
Peak (MHz)	Amp (mV)	Pedestal (MHz)	Amp (mV)	P_{zz} (%)	Error (%)
32.62(0.000)	20	32.85(0.015)	70	26.7	5.4
32.63(0.015)	30	32.85(0.020)	40	28.8	5.7
32.64(0.015)	30	32.84(0.025)	40	29.4	7.2
32.64(0.015)	25	32.83(0.035)	20	26.5	6.8
32.64(0.015)	20	32.85(0.035)	70	30.3	7.8
32.64(0.020)	20	32.85(0.025)	40	27.5	4.7
32.64(0.015)	40	32.88(0.055)	50	31.1	8.5

NIM A 981, (2020), 164504

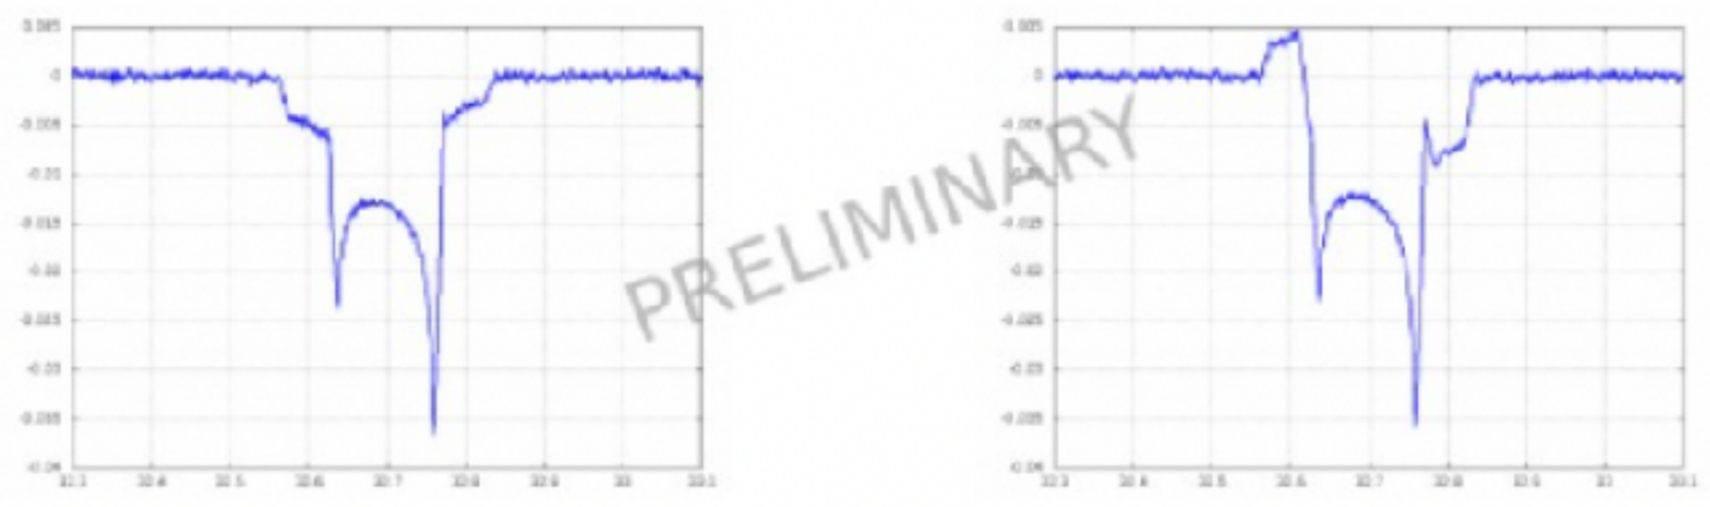

before

after

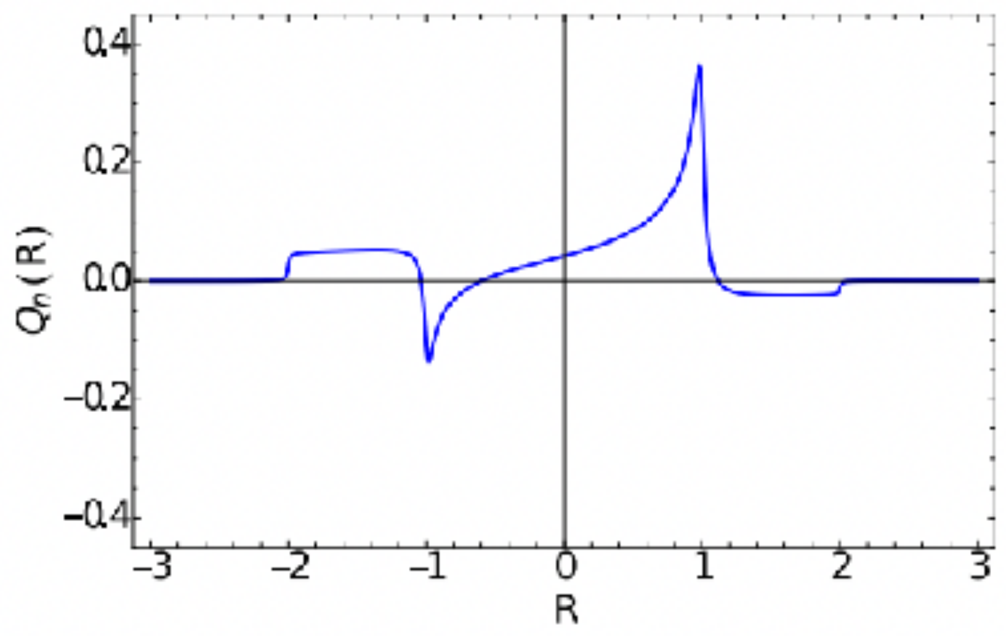

Rotating Targets (work still in progress)

And results (slow rotation)

- Rotate to TRY to burn one entire absorption line
- Spin Diffusion fights repopulates with rotation but changes for every angle

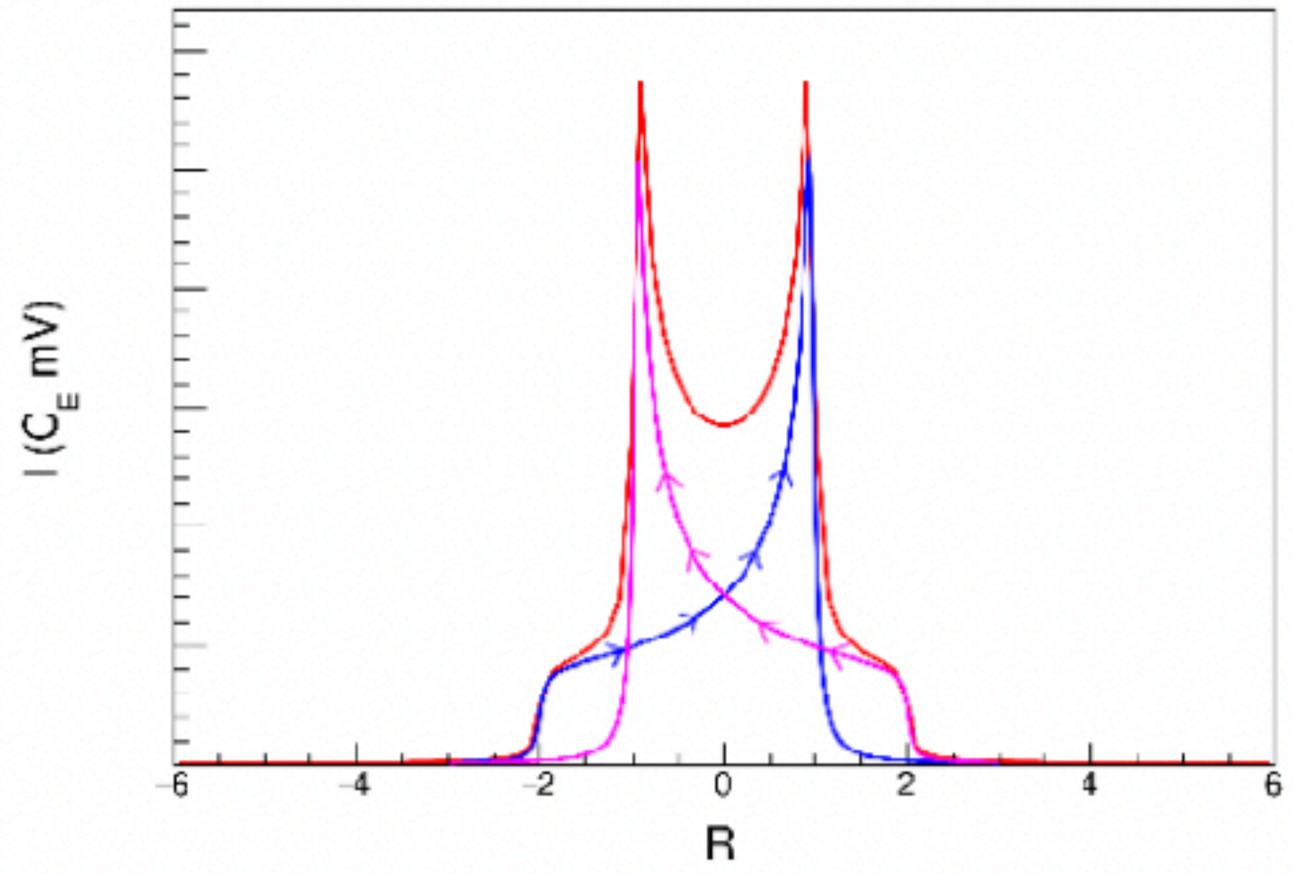

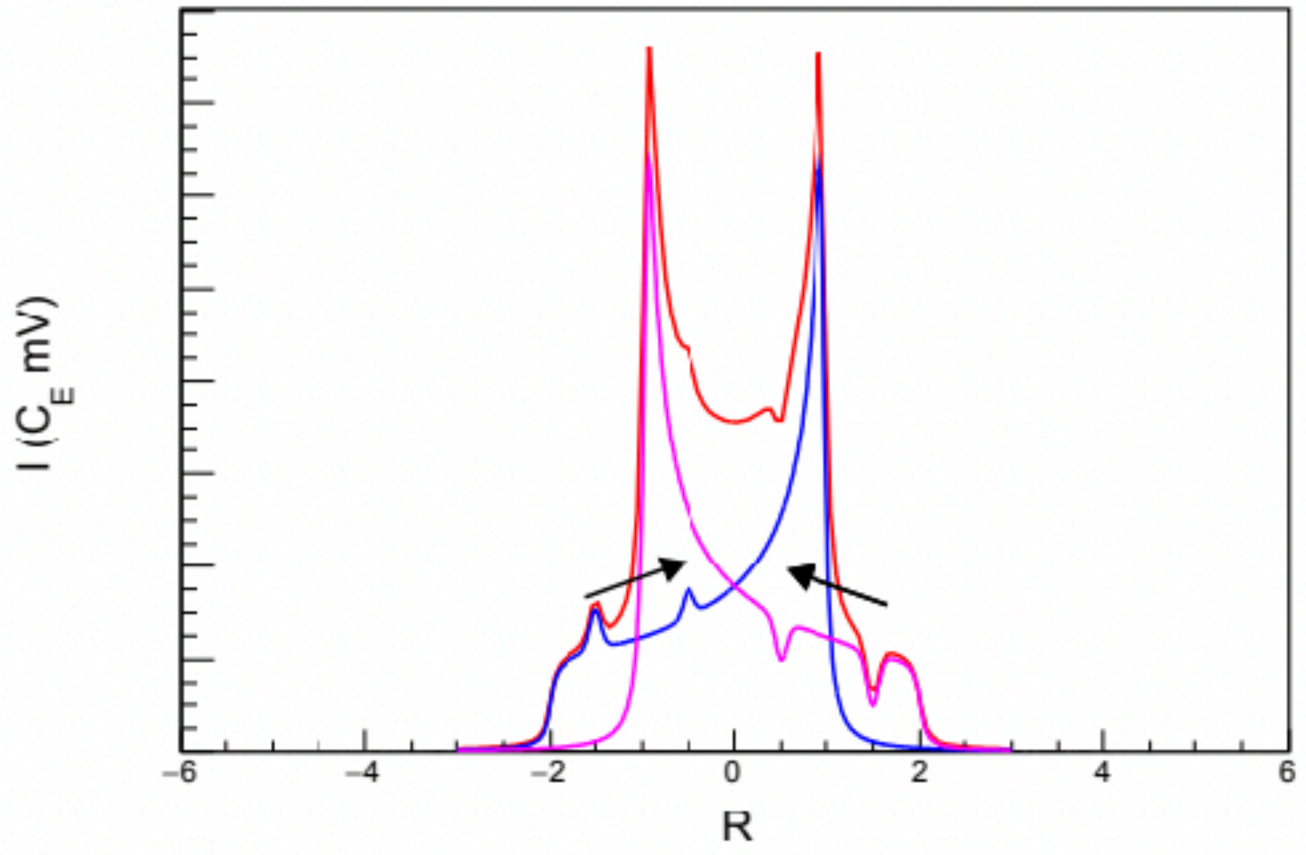

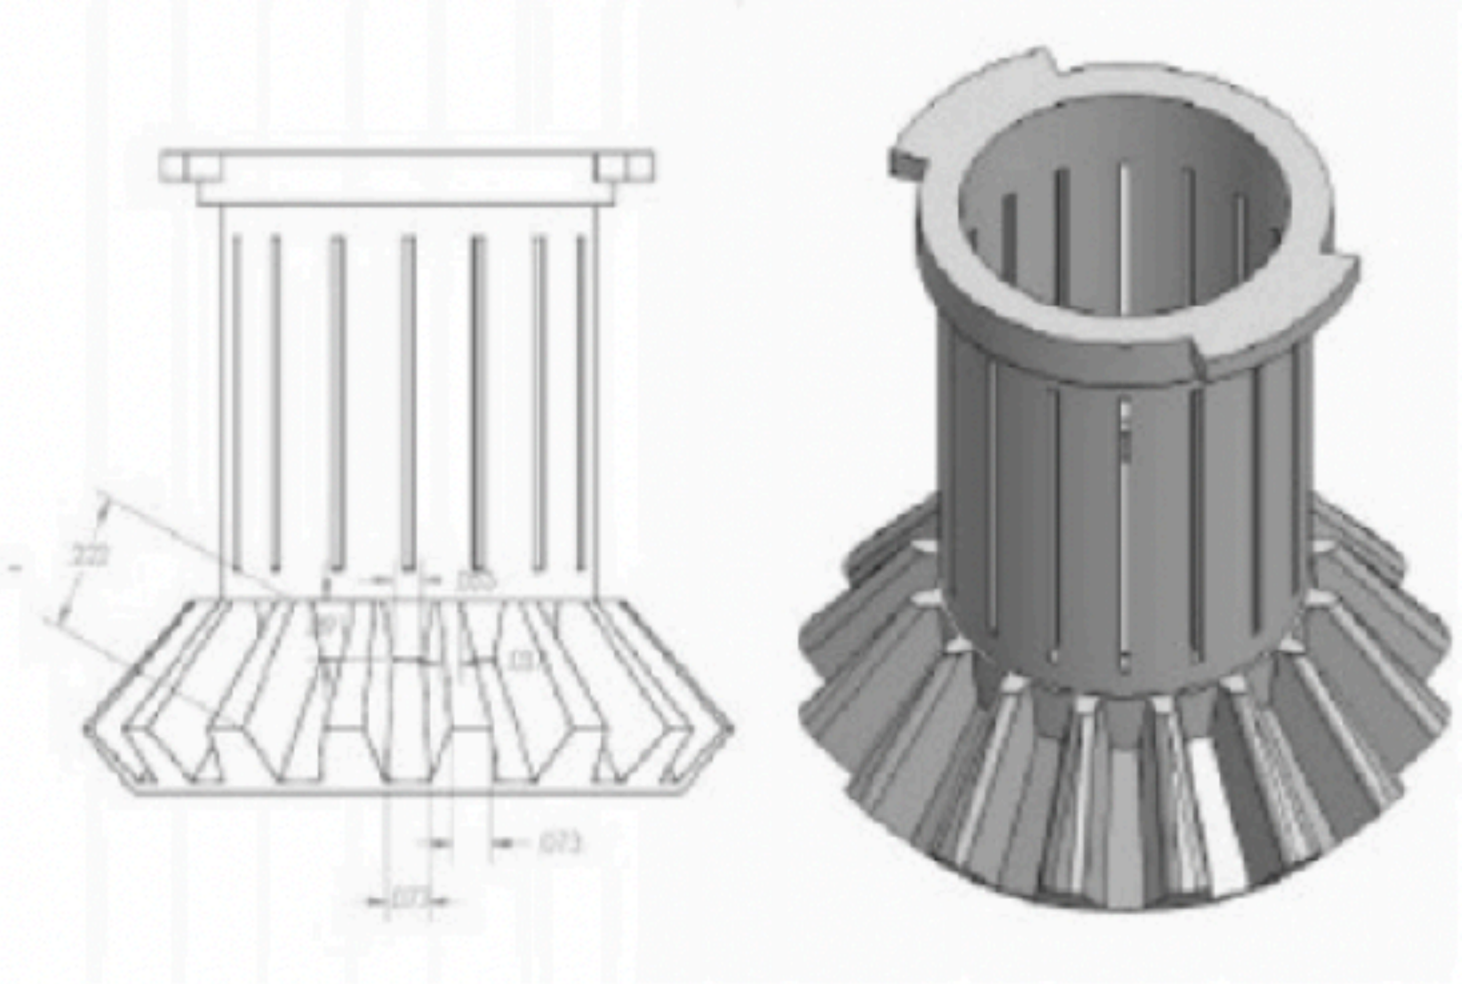

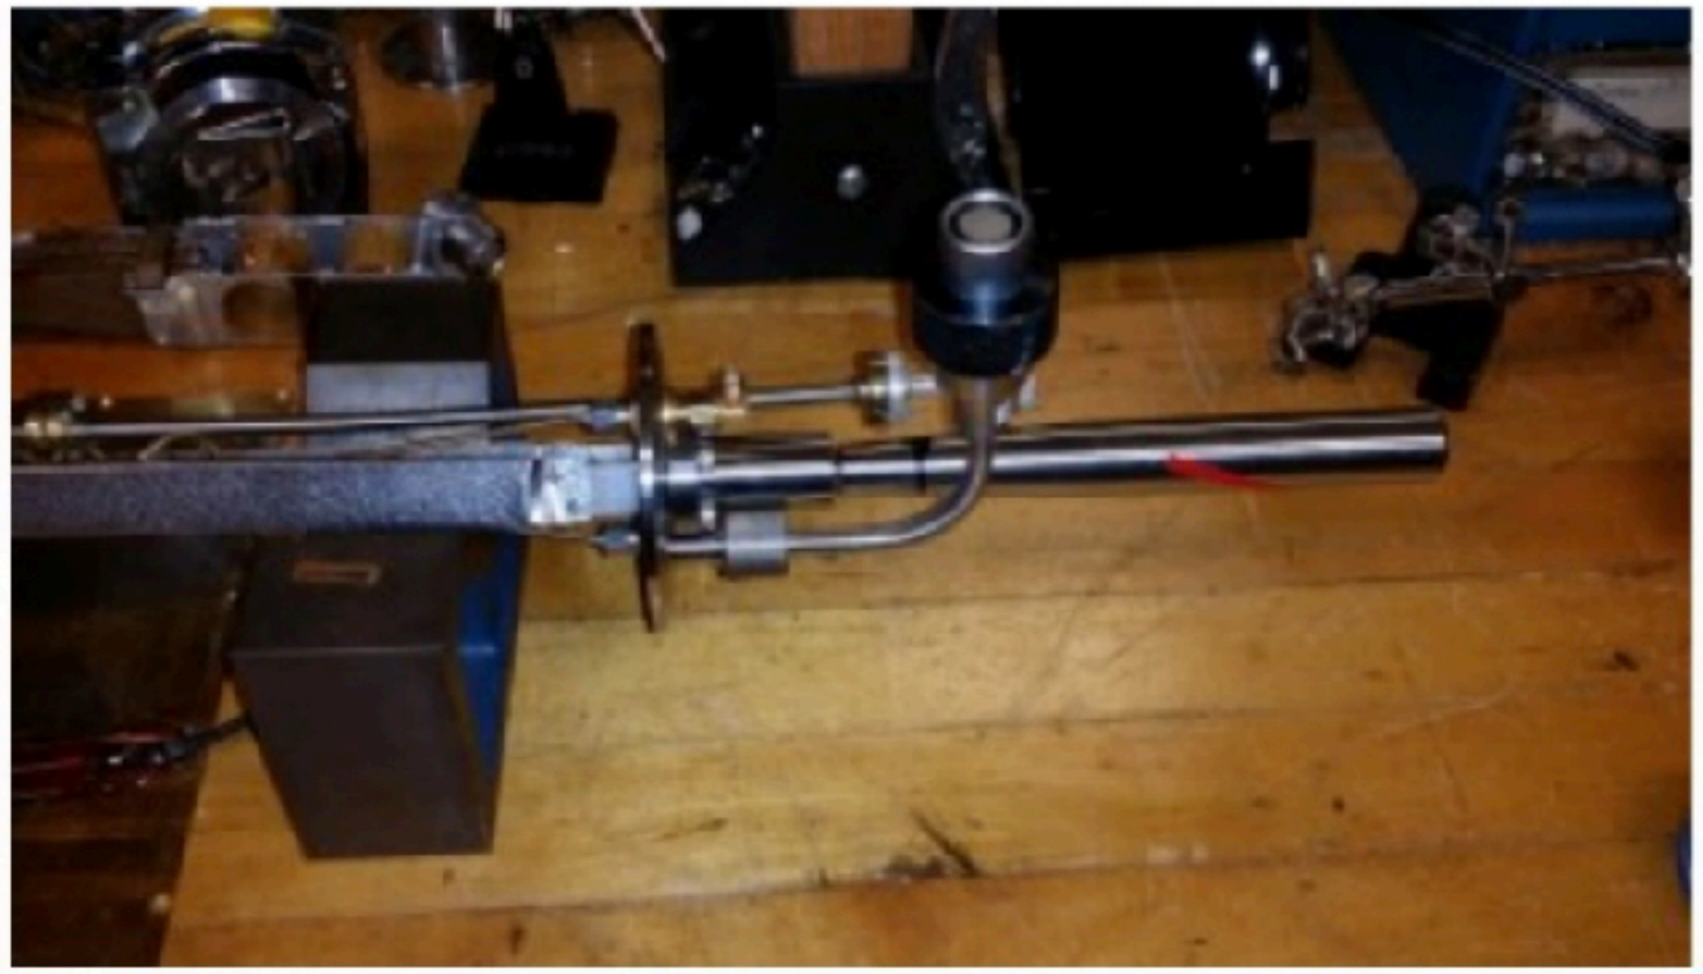

- Kel-F $(C_2ClF_3)_n$ cup and driving gear
- Motor outside cryostat
- NMR coil around cup
- Already used with several designs at UVA
- 1 Hz achieved with no problem
- Fixed beam spot

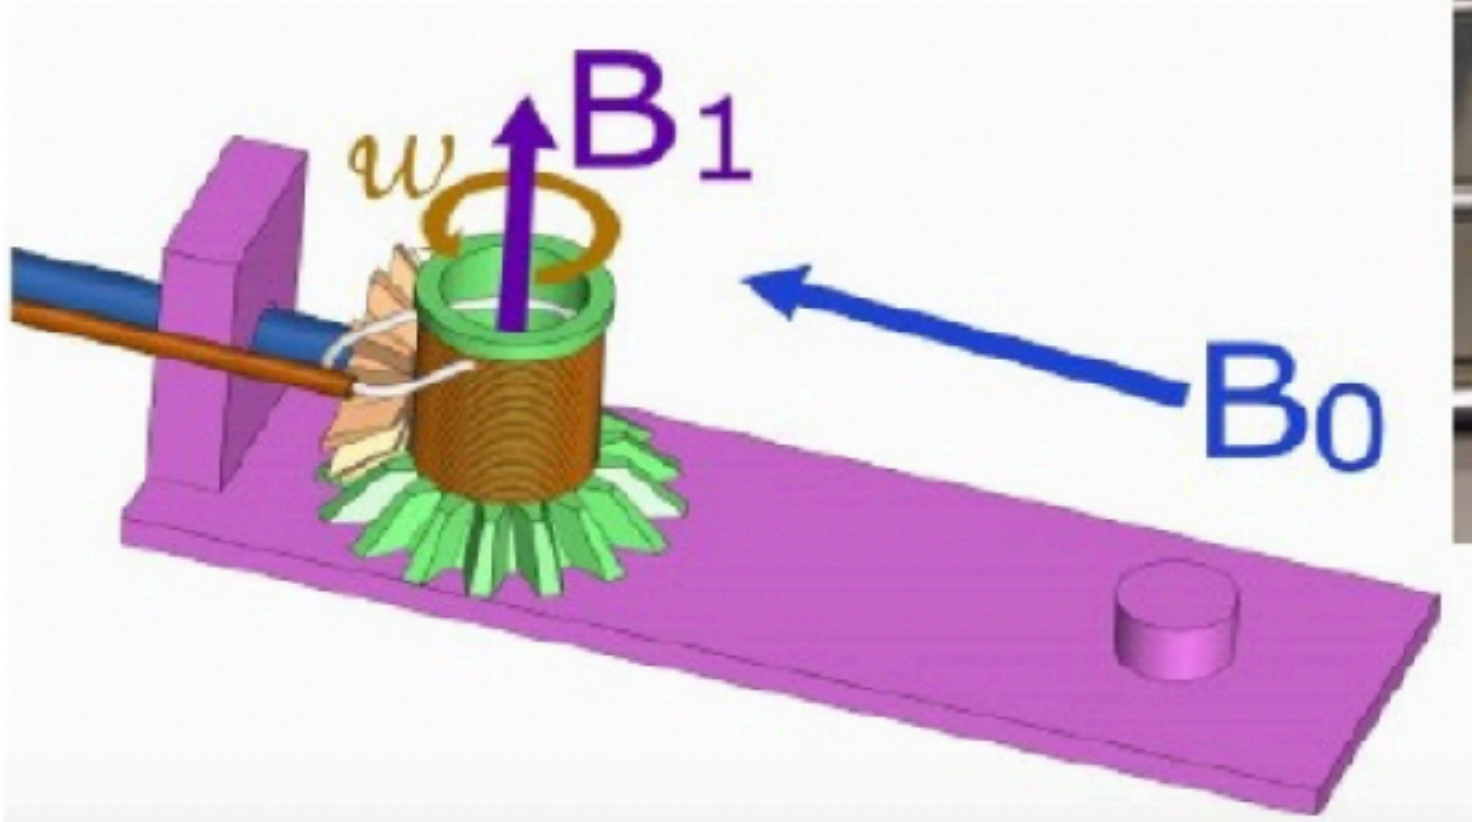

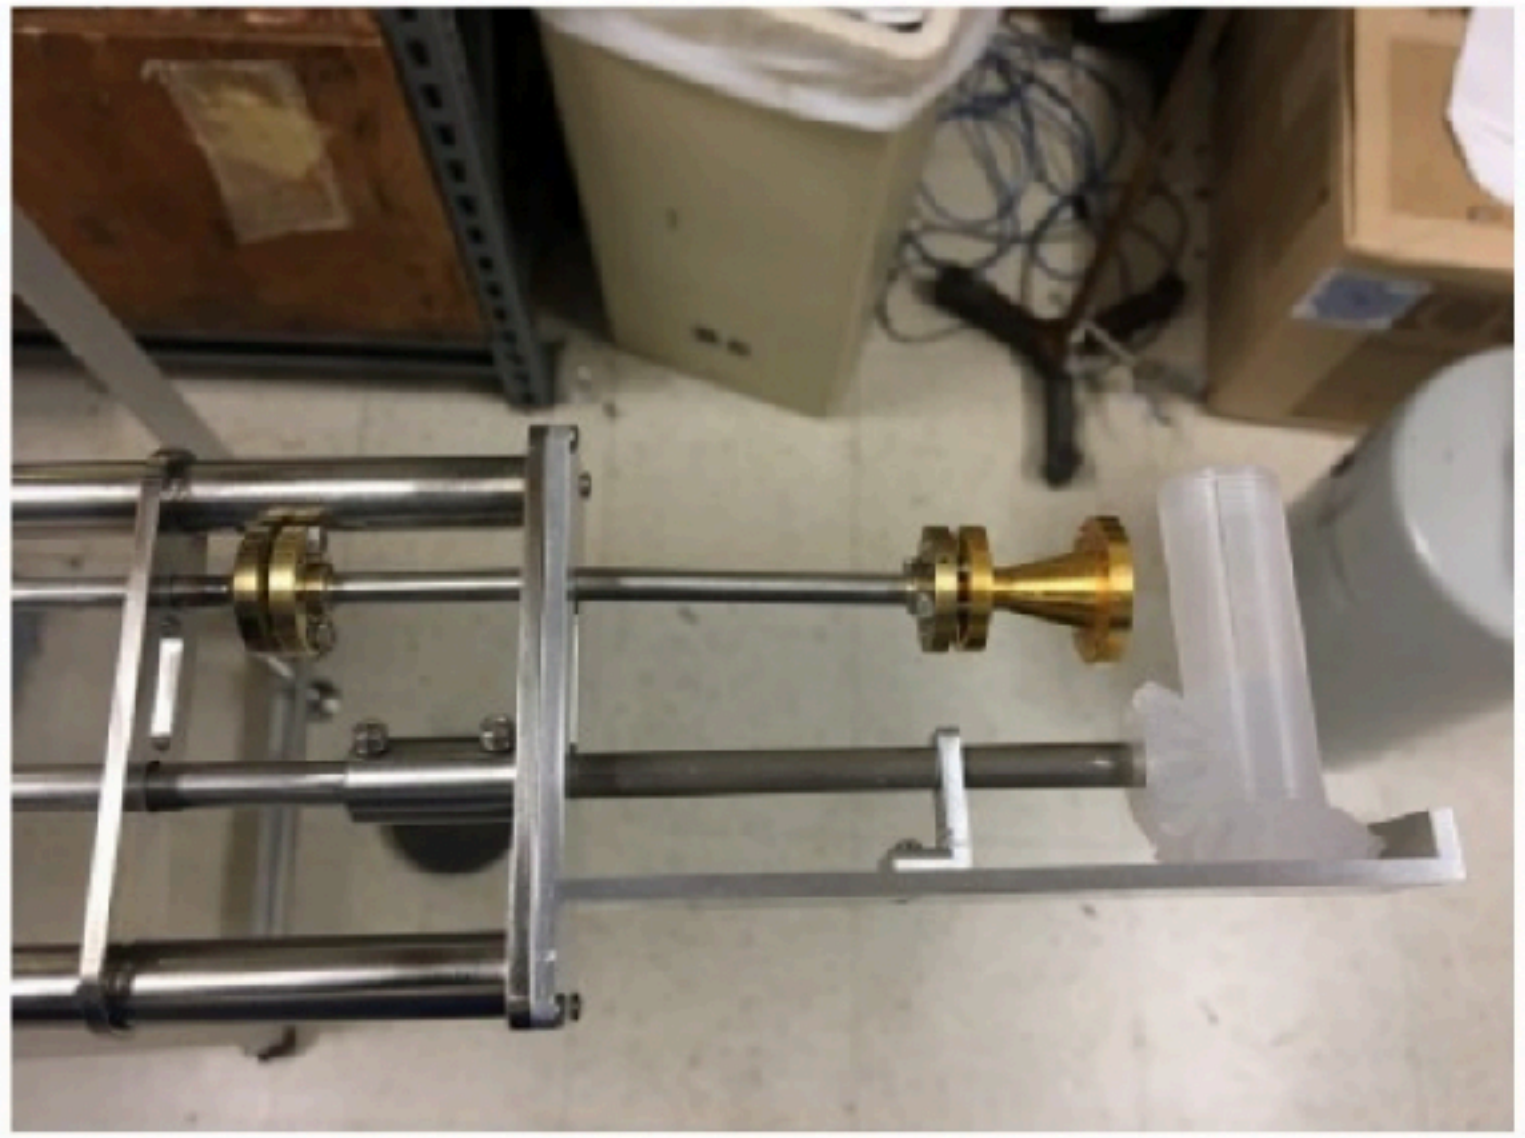

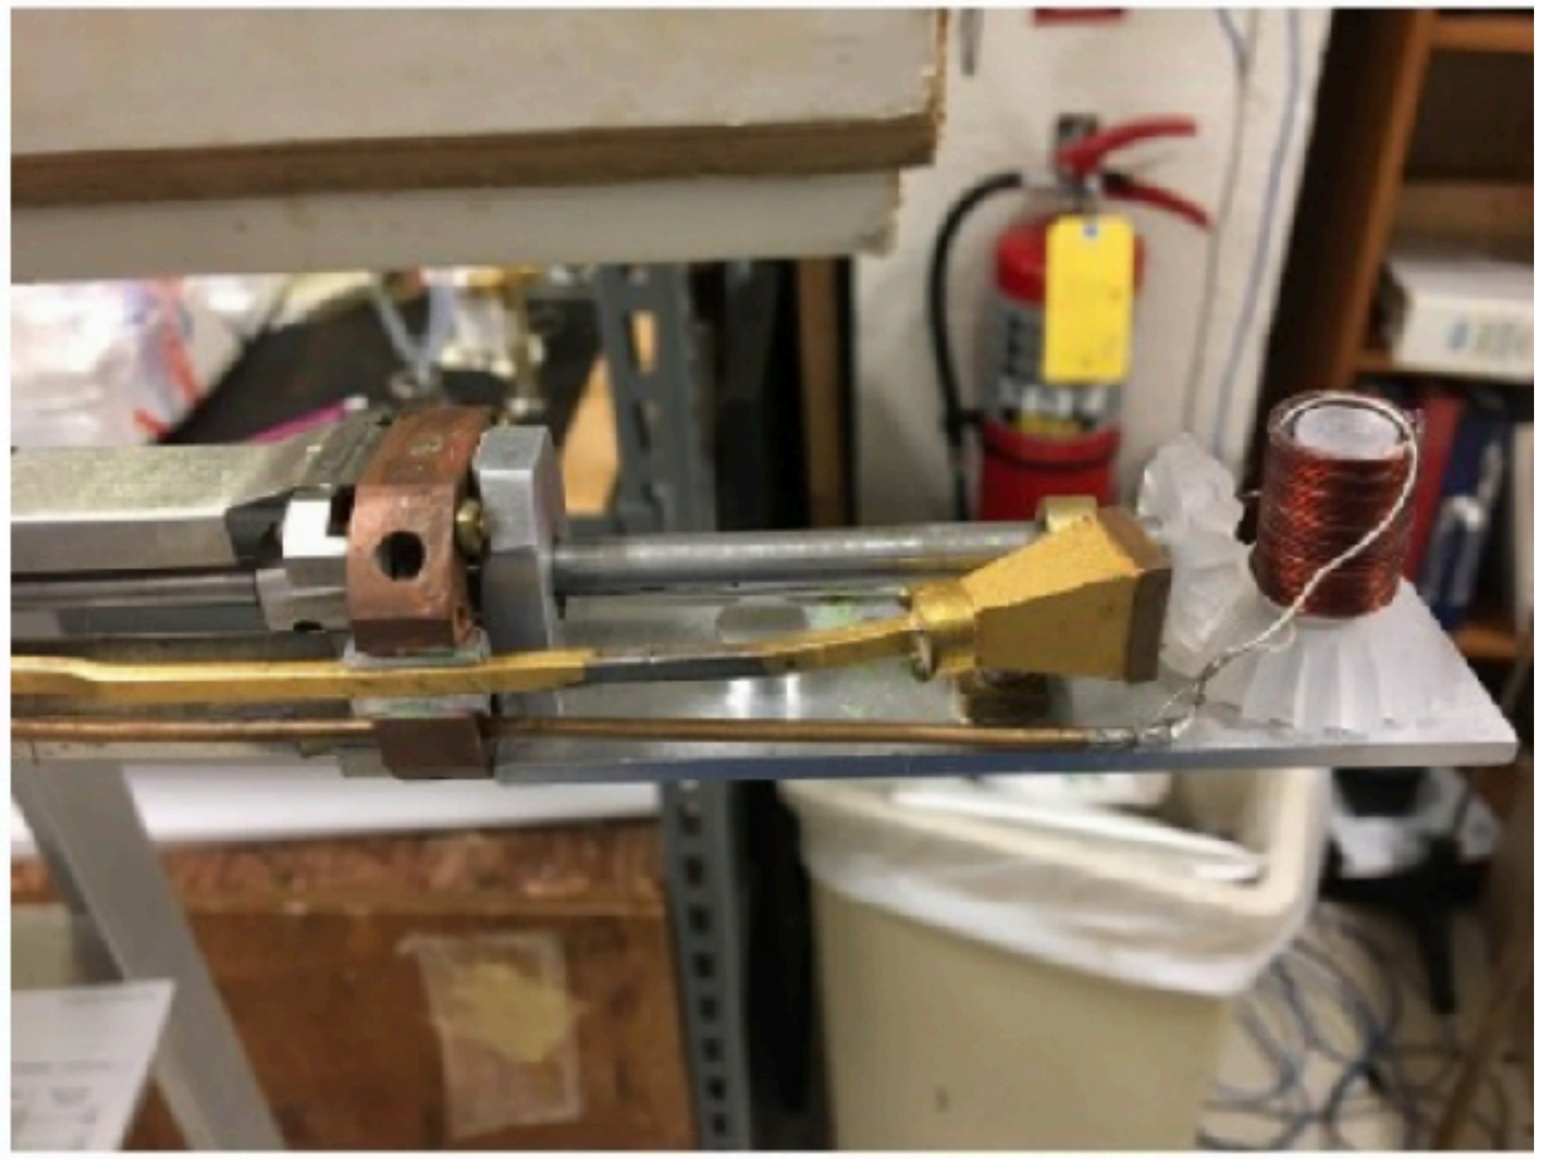

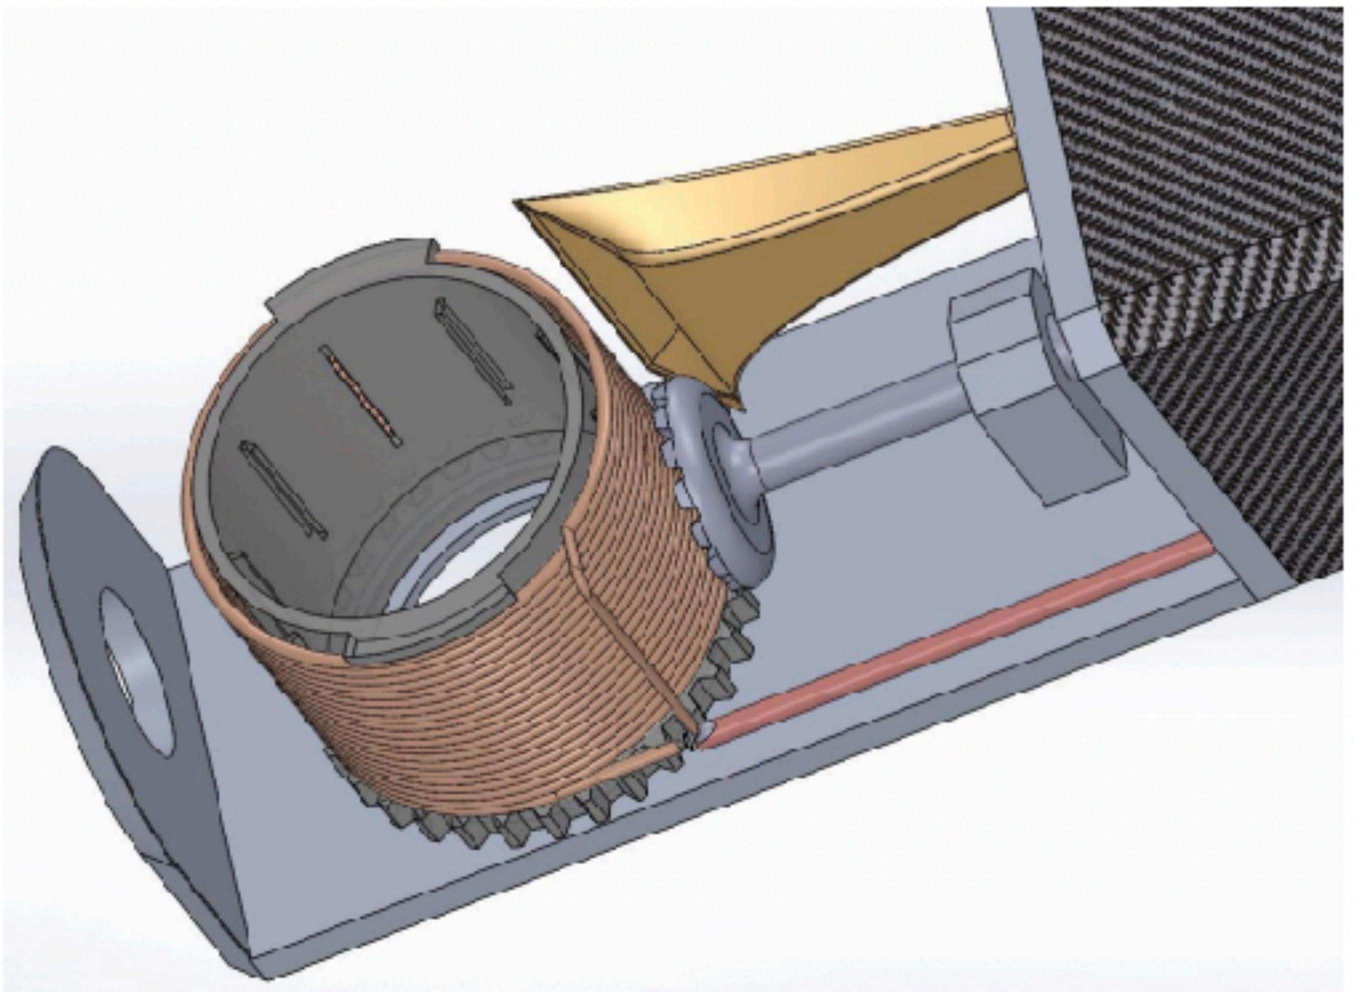

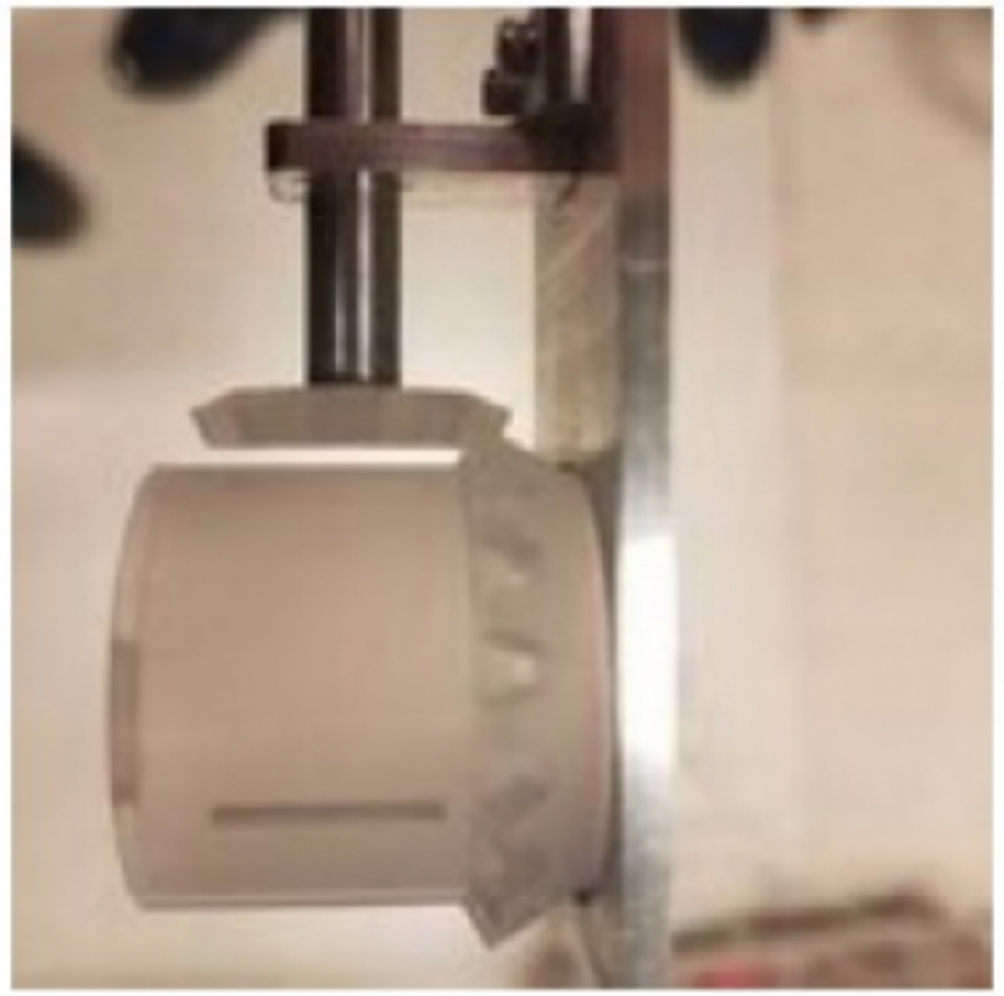

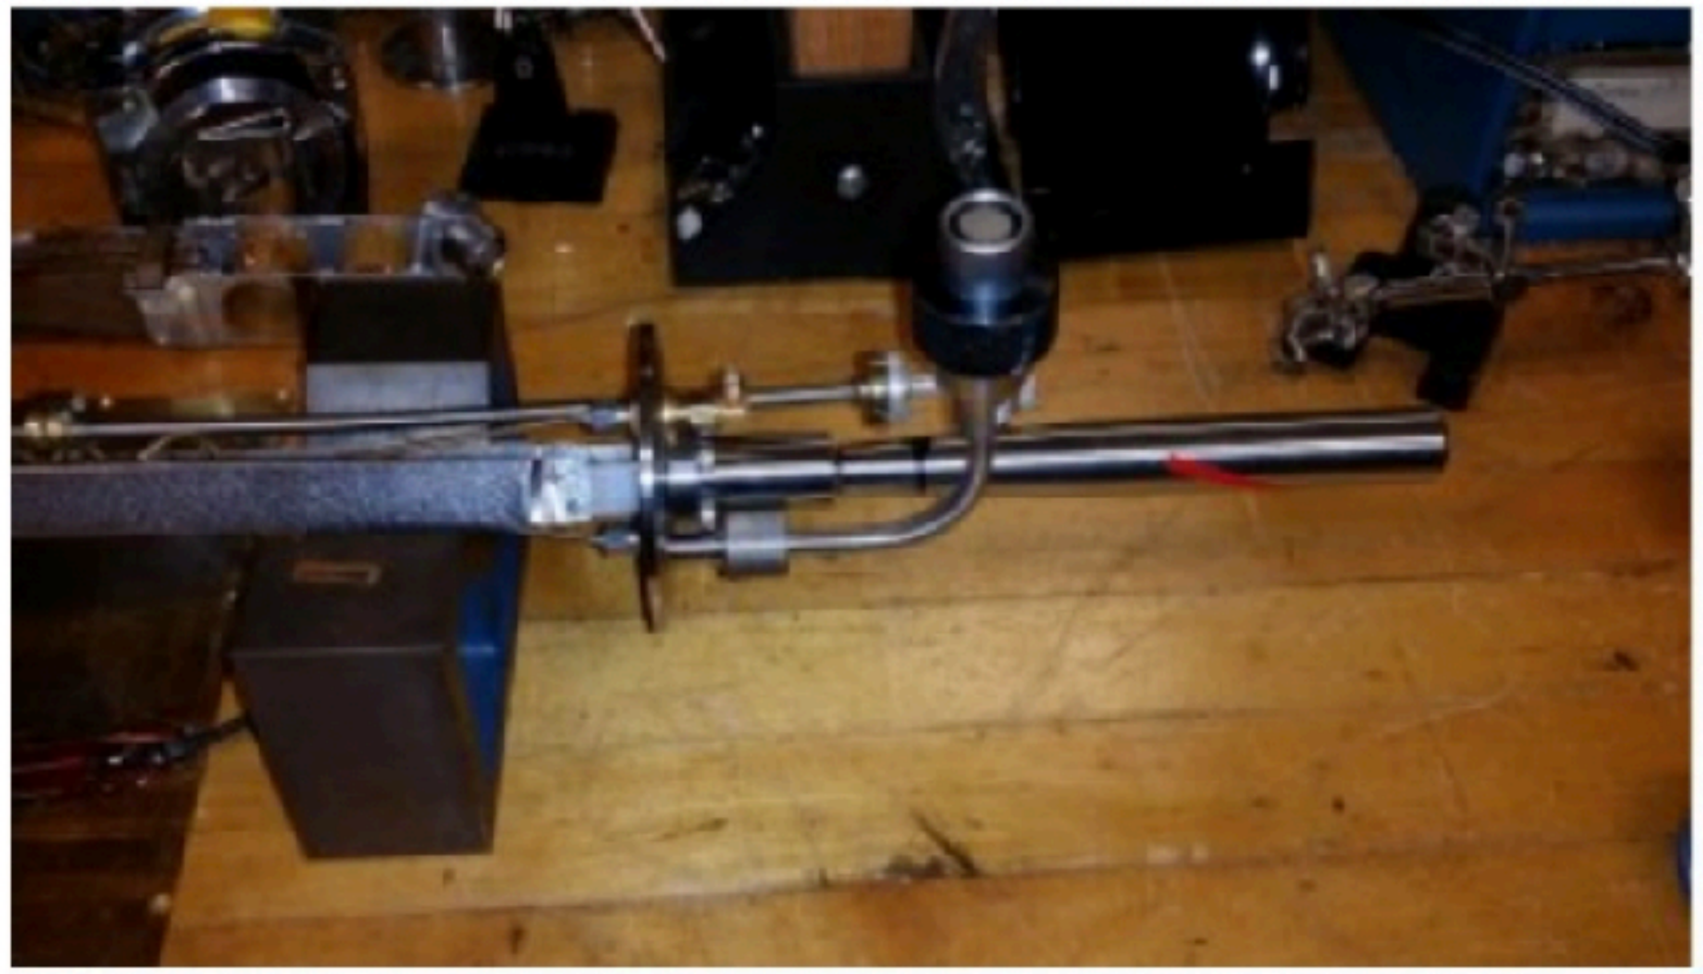

- Motor outside cryostat
- NMR coil around cup
- Already used with several designs at UVA
- 1 Hz achieved with no problem
- Fixed beam spot

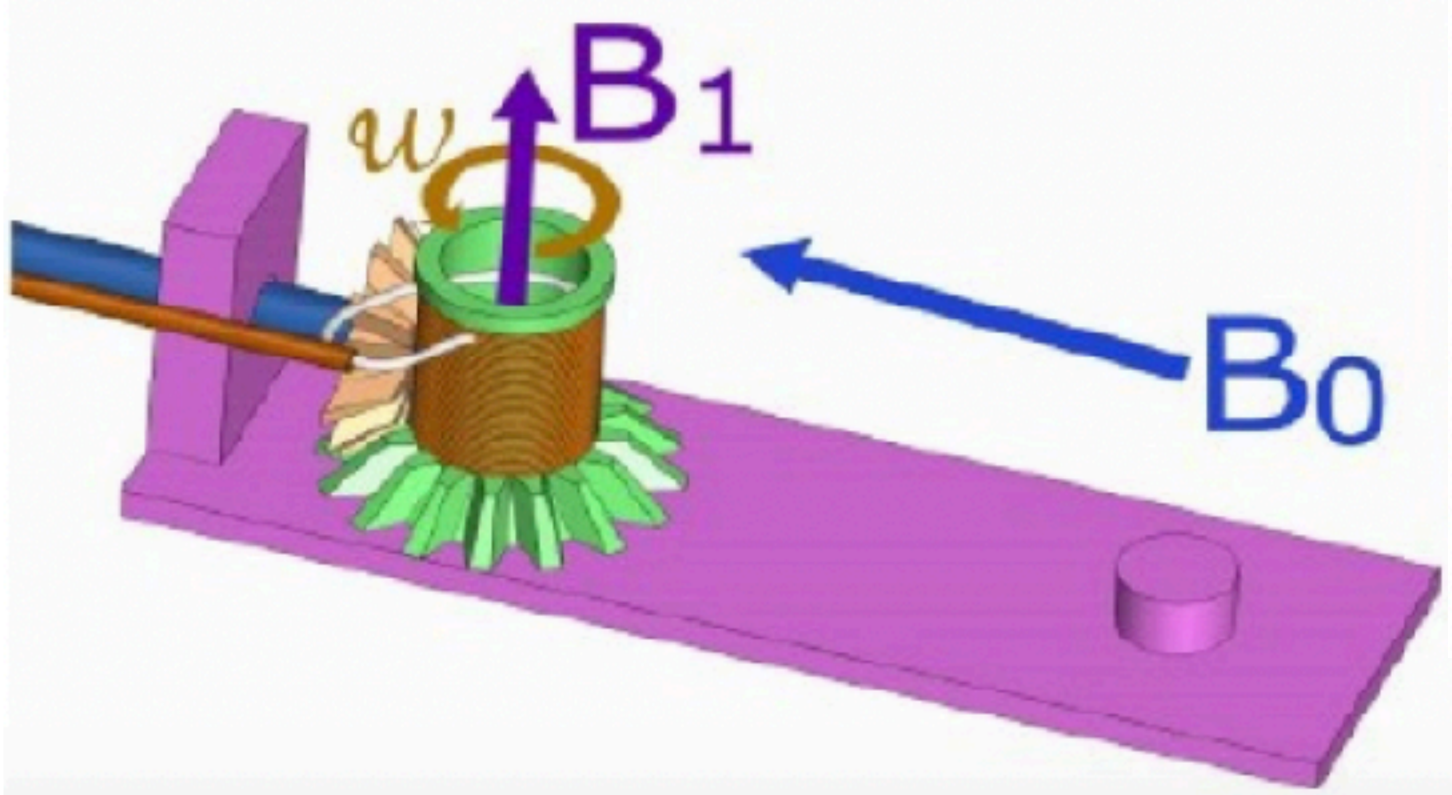

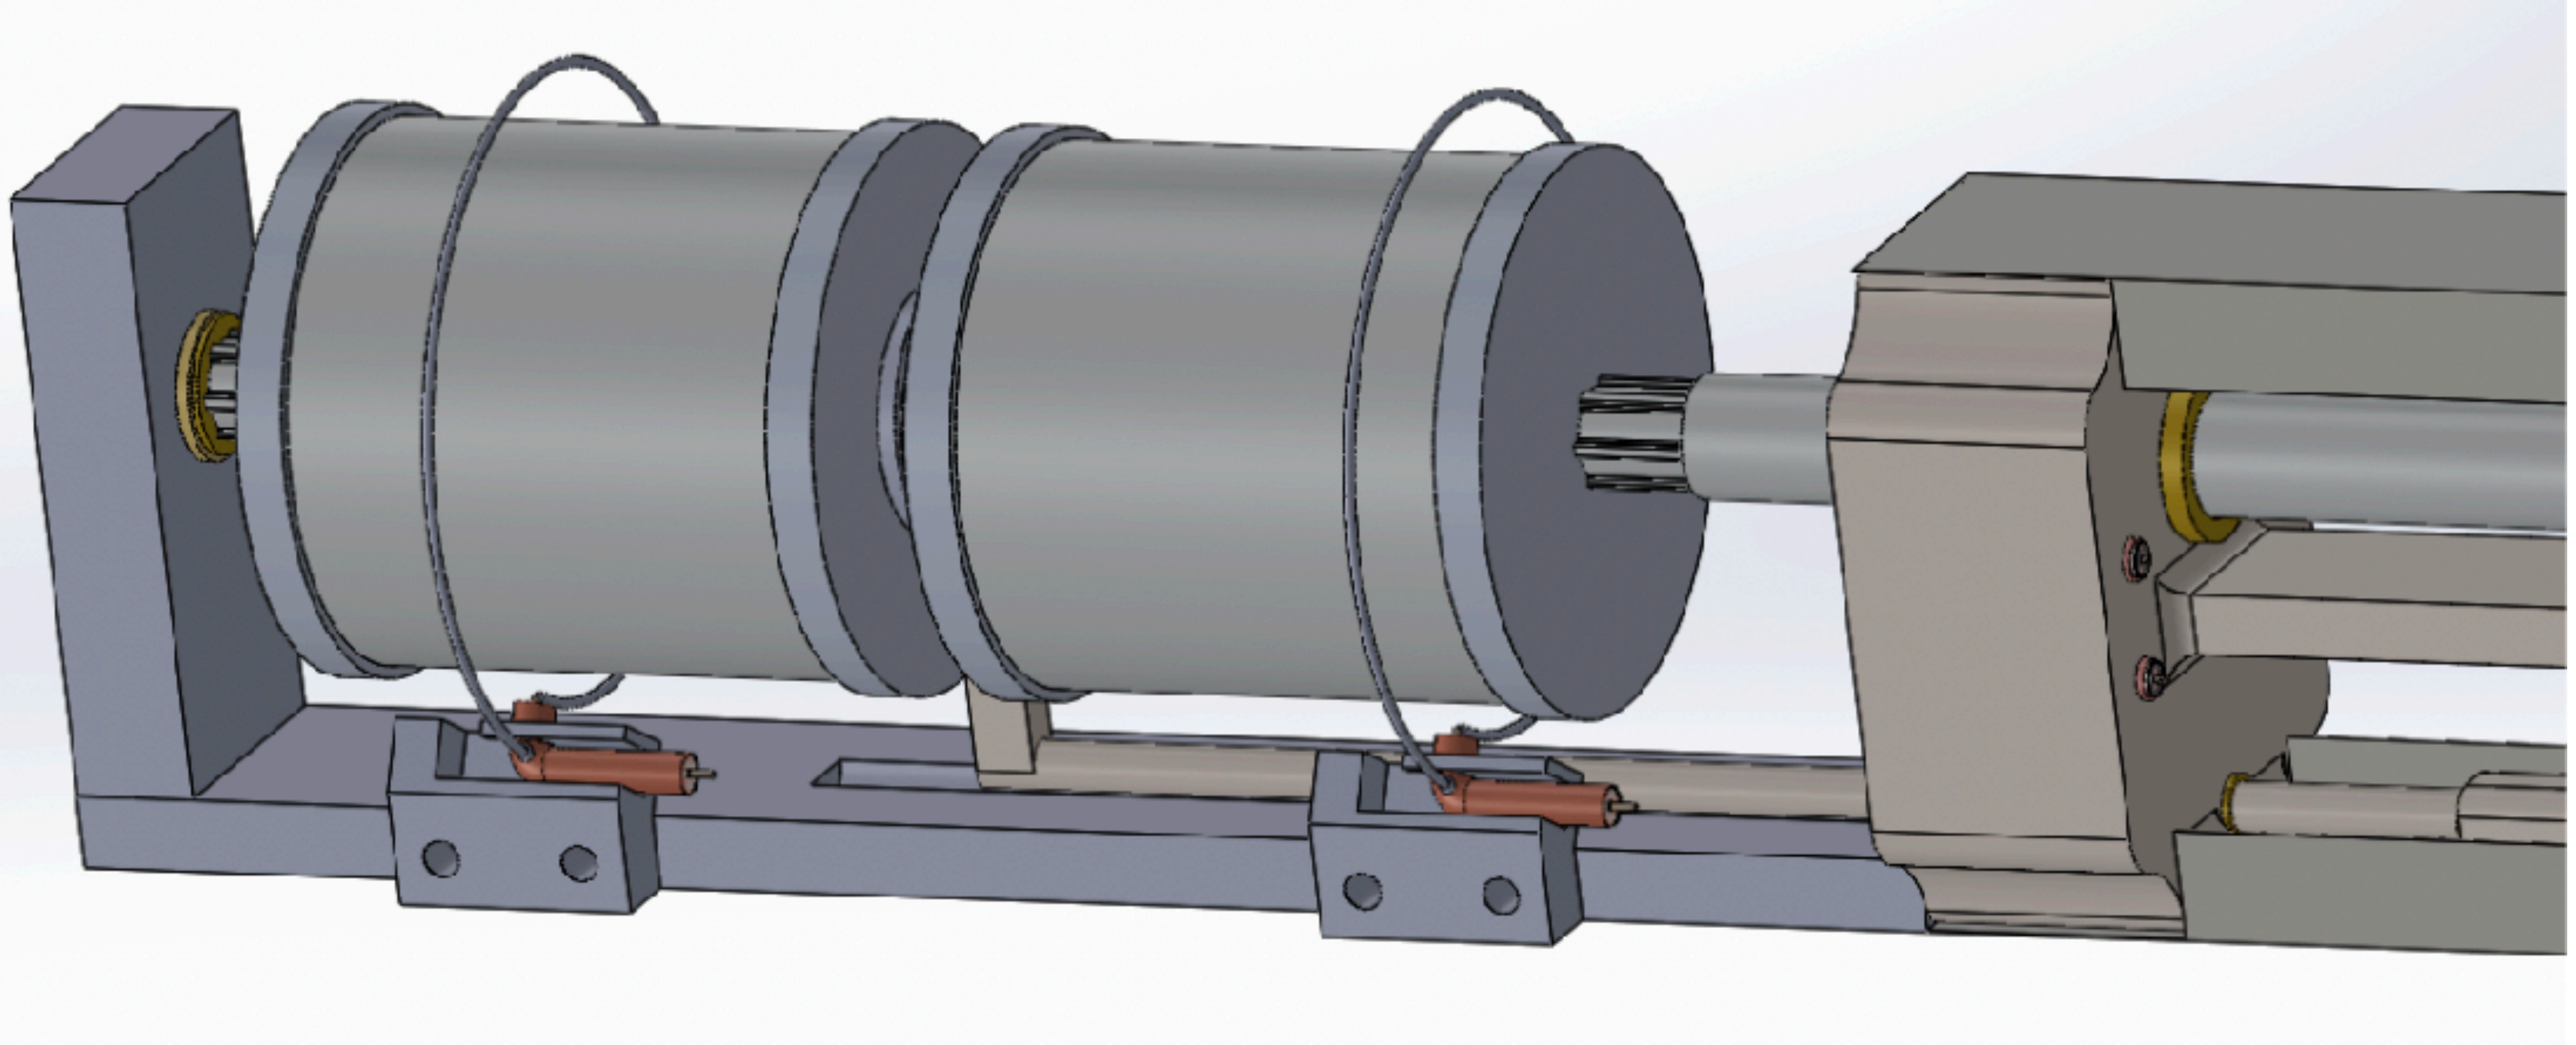

Rotation Results

10% relative uncertainty is the best we can do with rotation (no AI)

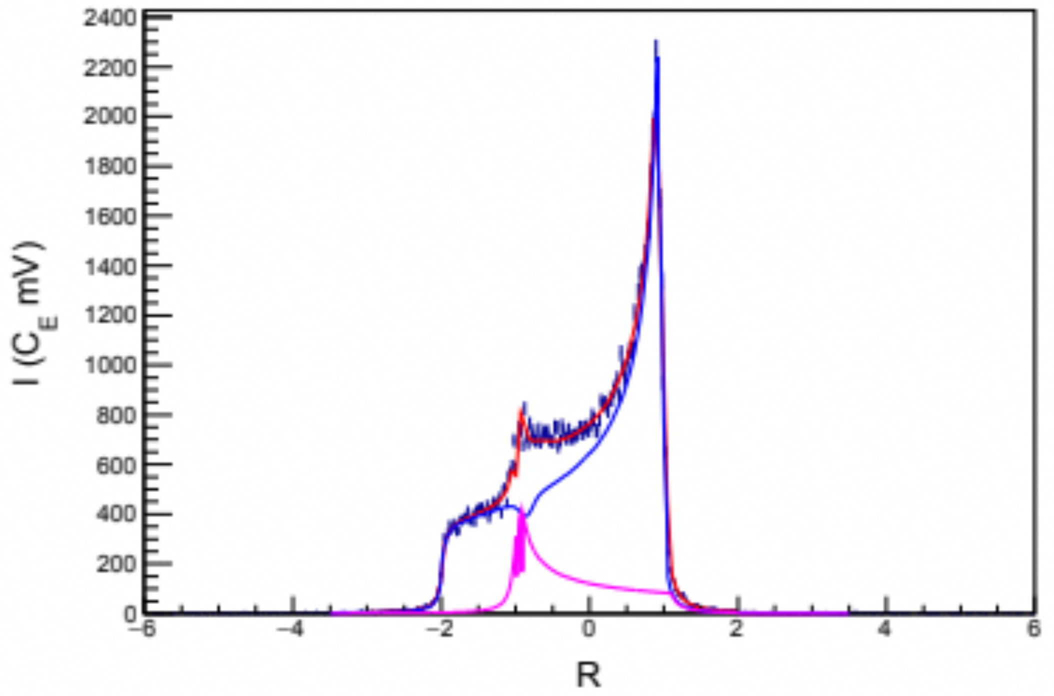

At <40° with respects to B

Rotation rate	rss-RF Enhanced Measurements					Relative error
Ω^{-1}	Peak (MHz)	Amp (mV)	Pedestal (MHz)	Amp (mV)	P_{zz} (%)	Error (%)
50	32.65(0.010)	15	32.85(0.015)	45	35.7	8.4
44	32.66(0.000)	10	32.88(0.015)	40	36.5	9.7
40	32.65(0.000)	15	32.88(0.015)	40	36.3	9.3

NIM A 981, (2020), 164504

Basic Prototype

Standard single rotating cup

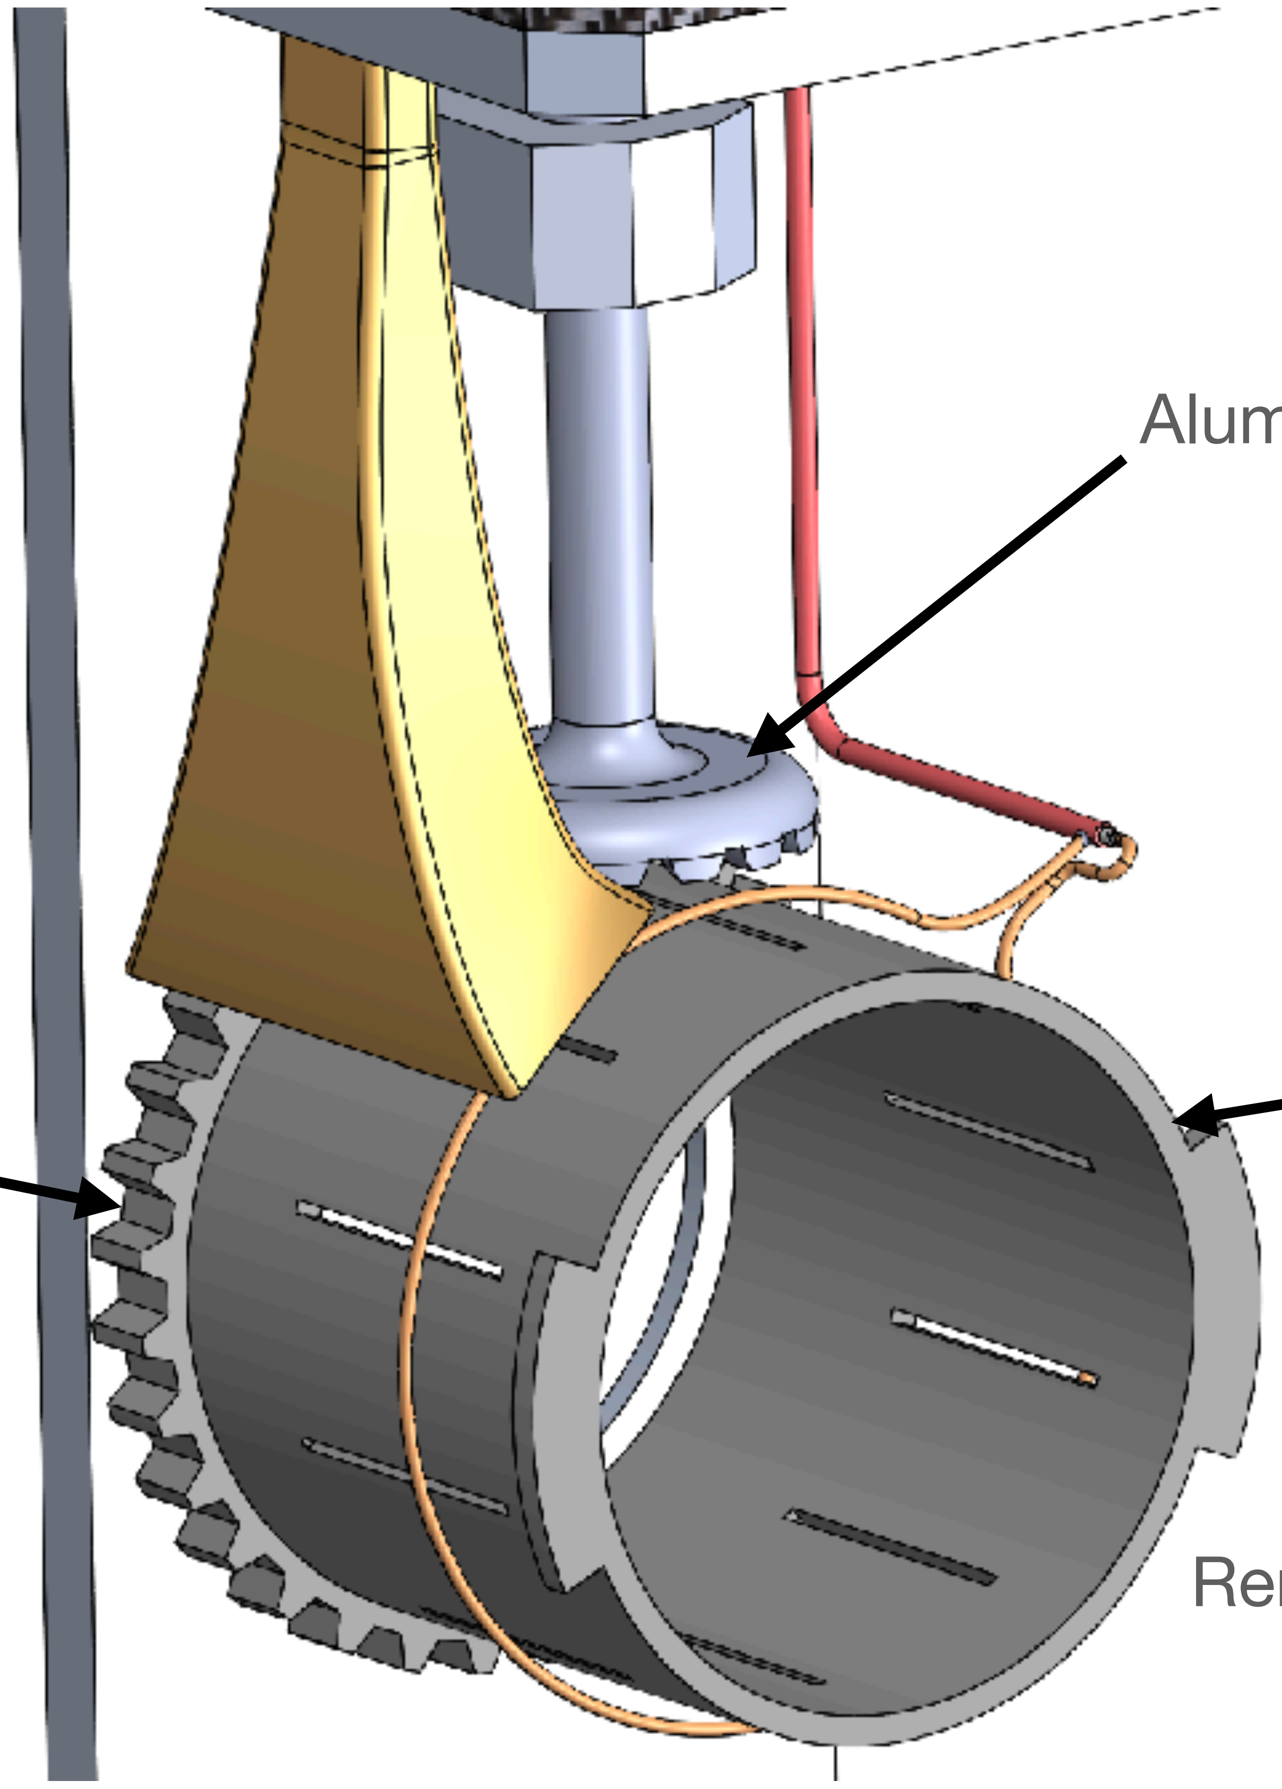

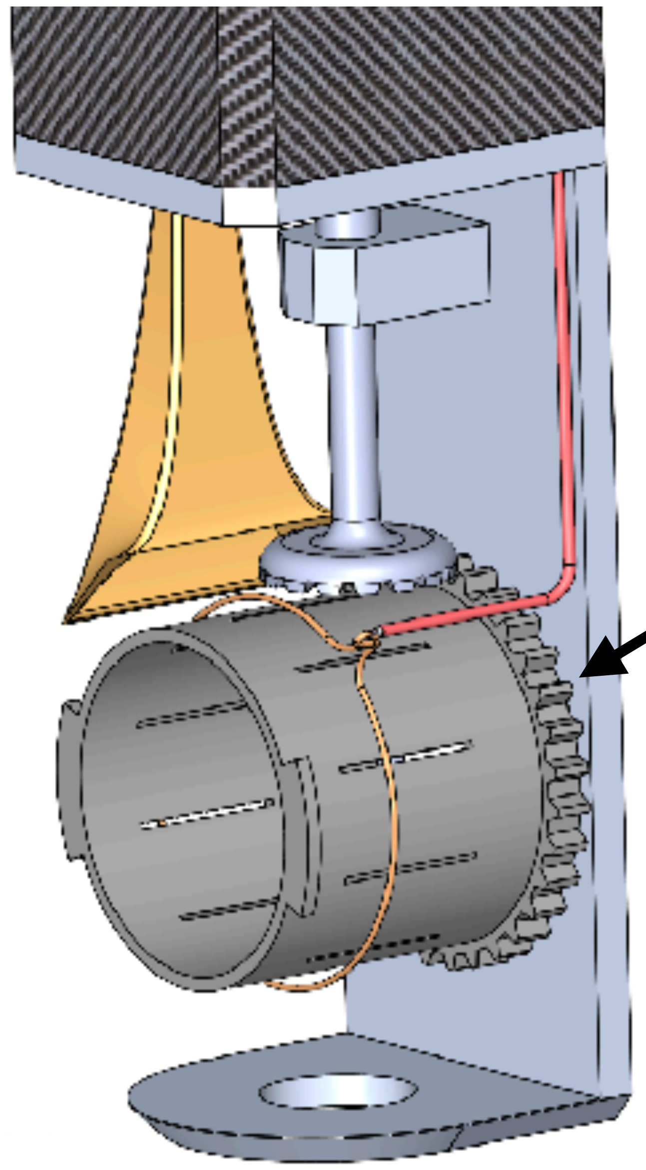

Polyether ether ketone

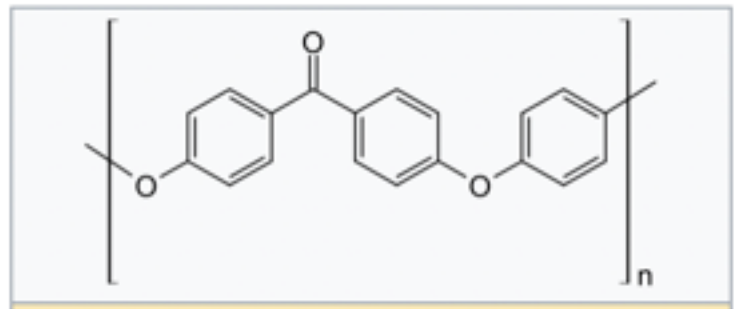

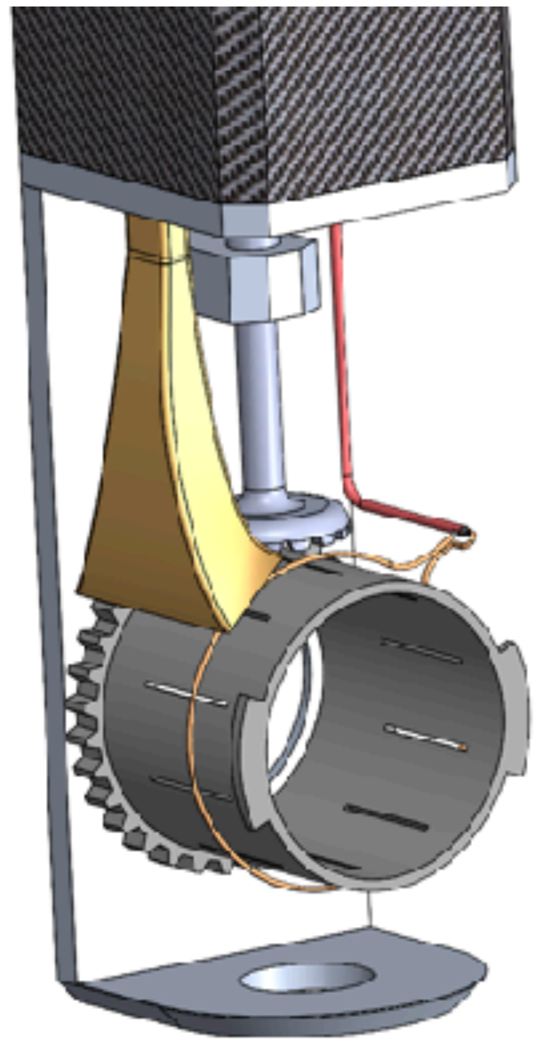

Wireless NMR coil for improved signal

Critical for Target Calibrations

- Poor signal to noise during TE diminishes FOM
- Inner coil inductively coupled to the static coil
    - Matched
    - Tuned
    - Balanced geometry

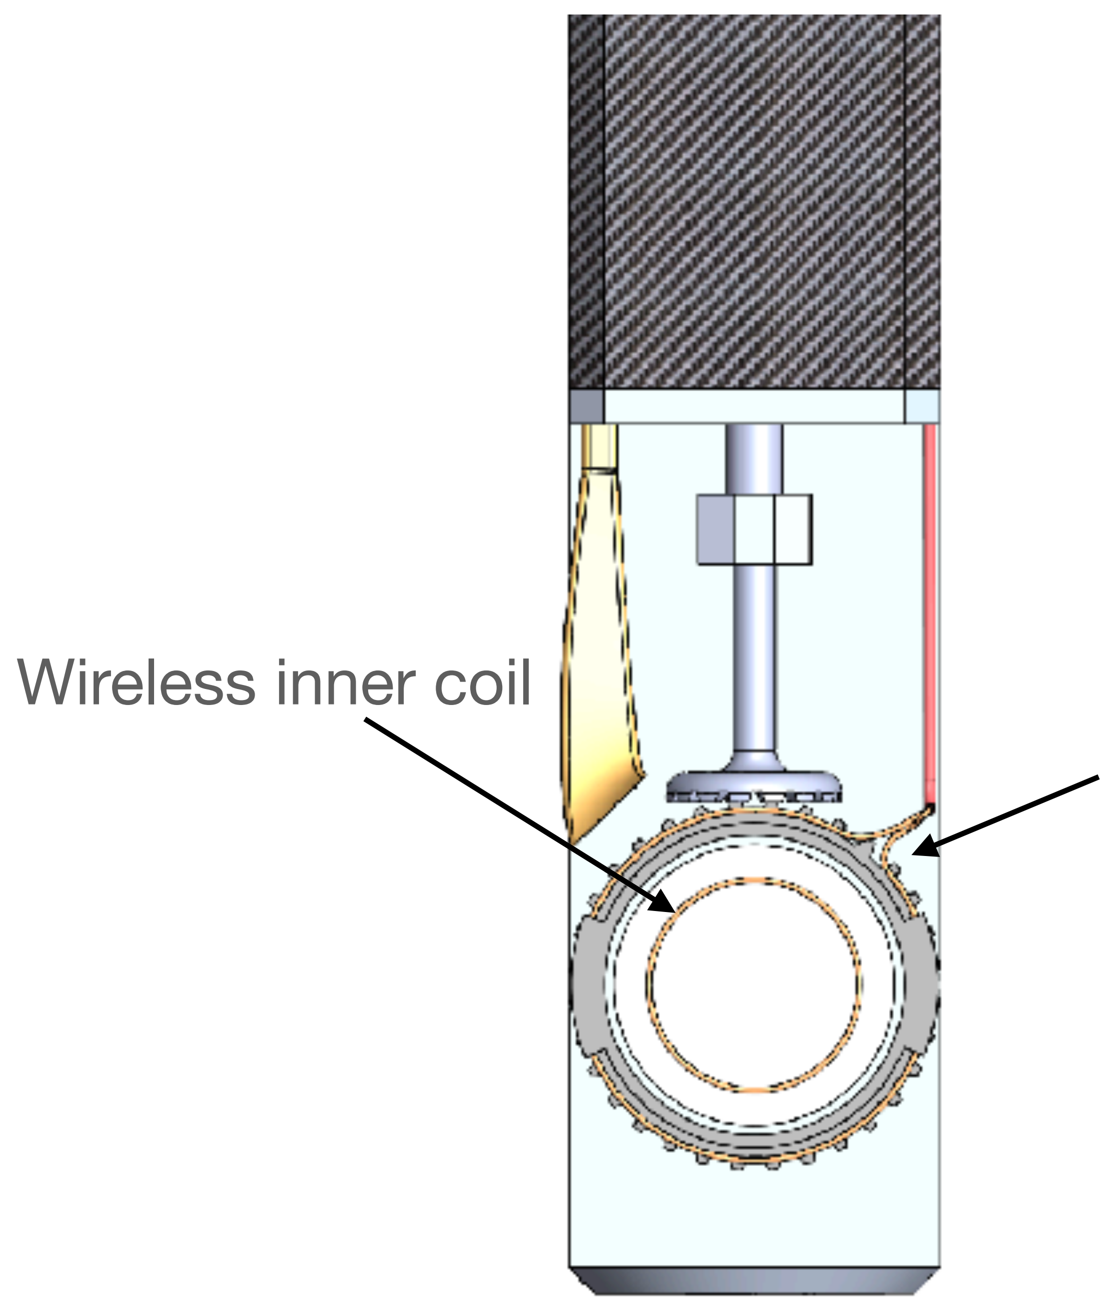

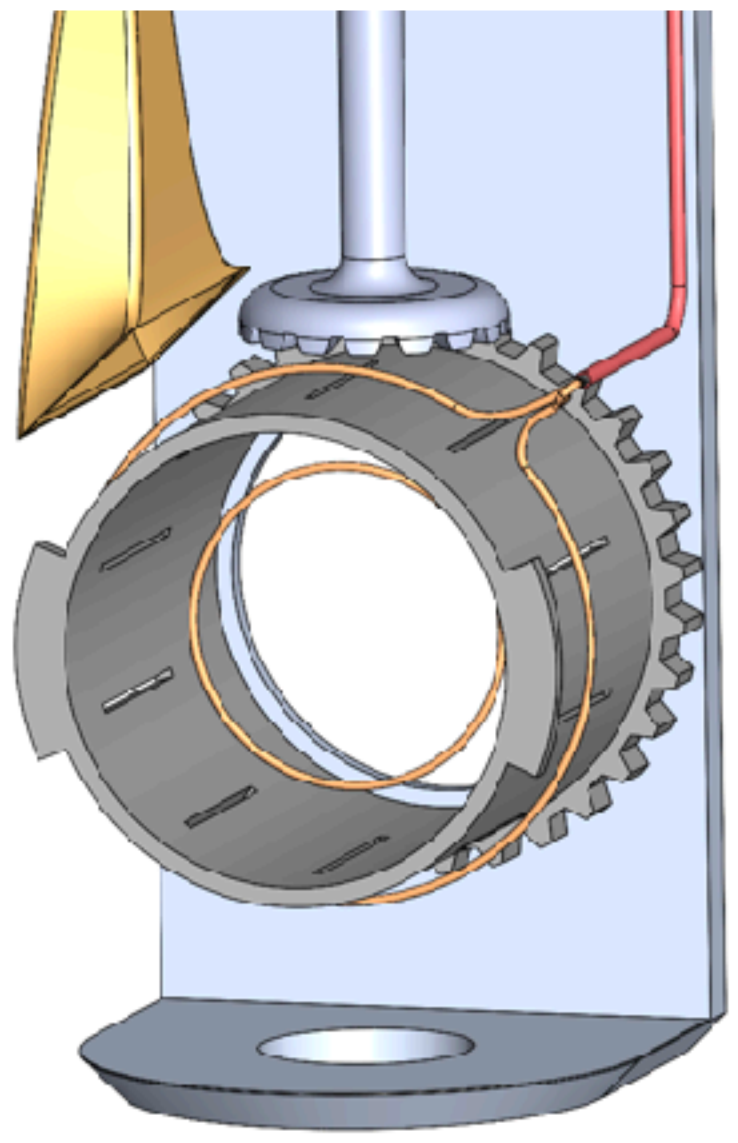

General Outlook

- First Generation
  - b1 structure function (Hall C)
  - Azz + T20 (Hall C)
  - DY Transversity (FNAL)
- Second Generation (beyond 3 years)
  - f_{1LL} in SIDIS (Hall A)
  - Tensor polarized DVCS (Hall C)
  - SIDIS Transversity (Hall A)
  - Photon Tensor polarized observables (Hall D)

Thank You

Backup stuff

AI in Polarization Measurements

Continuous Wave NMR

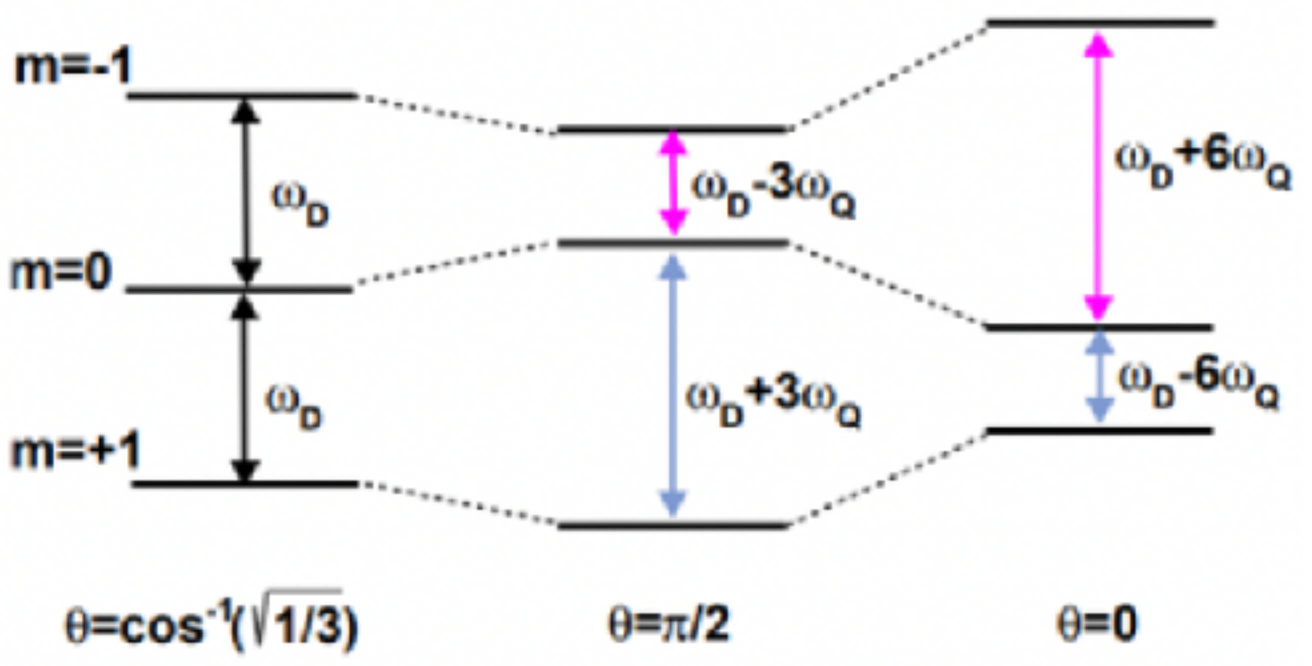

Deuteron Lineshape and Q-curve

$$\mathcal{F} = \frac{1}{2\pi\mathcal{X}}\left[2\cos(\alpha/2)\left(\arctan\left(\frac{\mathcal{Y}^2-\mathcal{X}^2}{2\mathcal{Y}\mathcal{X}\sin(\alpha/2)}\right)+\frac{\pi}{2}\right) + \sin(\alpha/2)\ln\left(\frac{\mathcal{Y}^2+\mathcal{X}^2+2\mathcal{Y}\mathcal{X}\cos(\alpha/2)}{\mathcal{Y}^2+\mathcal{X}^2-2\mathcal{Y}\mathcal{X}\cos(\alpha/2)}\right)\right],$$

$$\mathcal{X}^2 = \sqrt{\Gamma^2+(1-\varepsilon R-\eta\cos 2\phi)^2} \qquad \mathcal{Y} = \sqrt{3-\eta\cos 2\phi}$$

$$\cos\alpha = (1-\varepsilon R-\eta\cos 2\phi)/\mathcal{X}^2 \quad \Gamma\sim 0.05 \quad \eta\cos 2\phi\sim 0.04$$

Deuteron lineshape

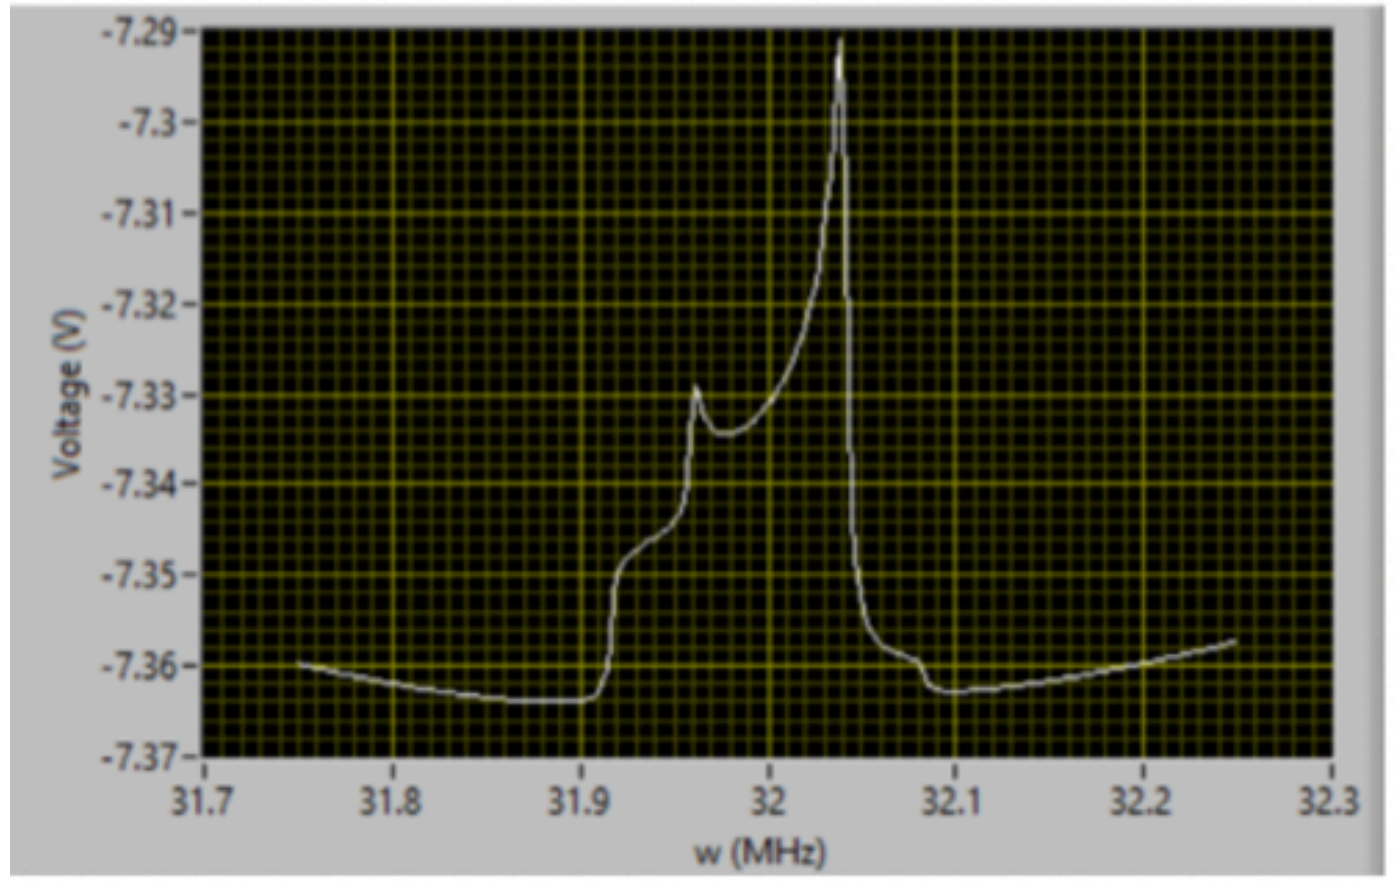

Q-Curve

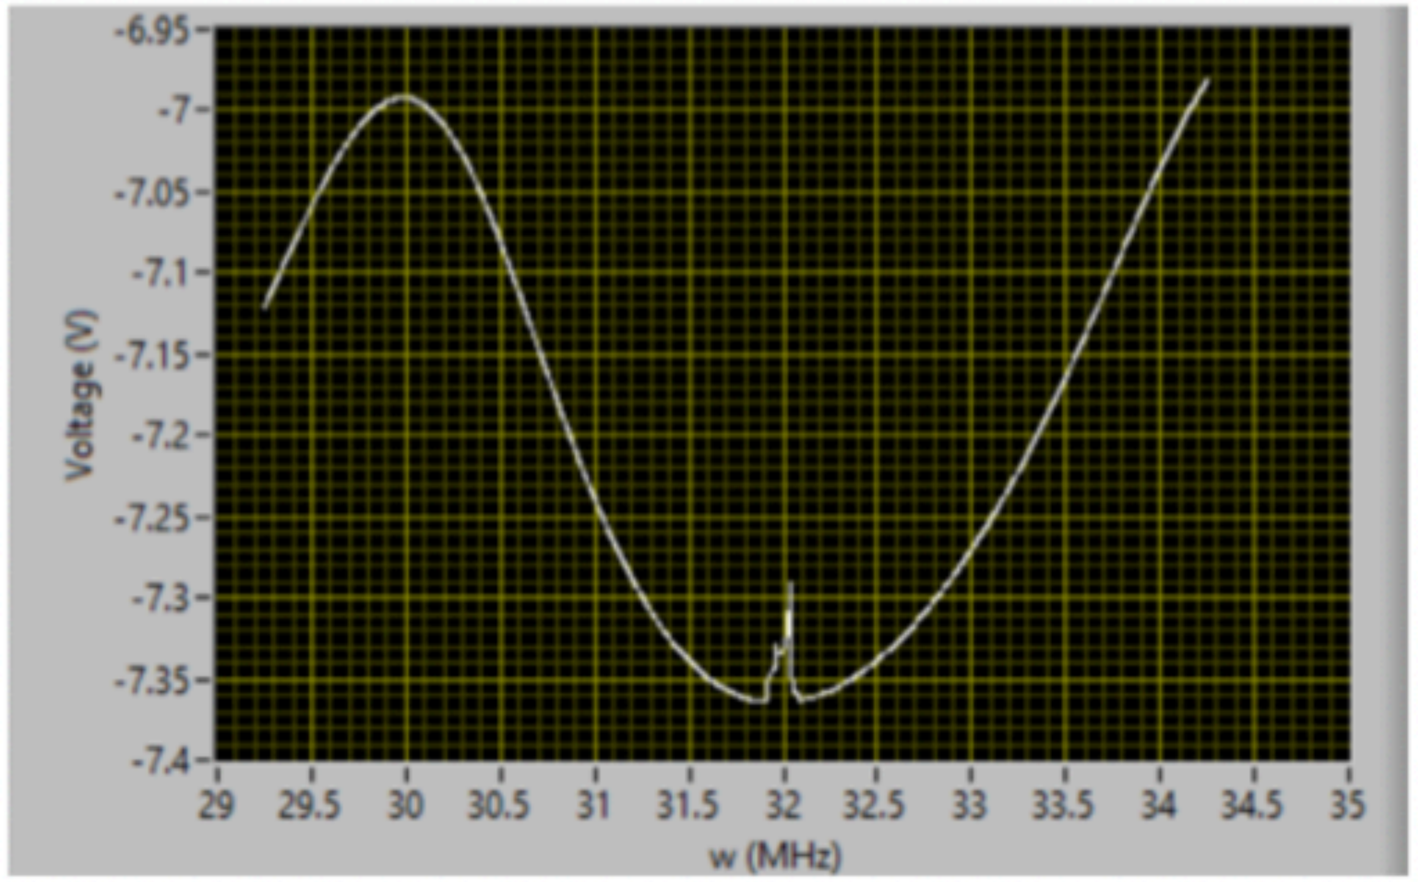

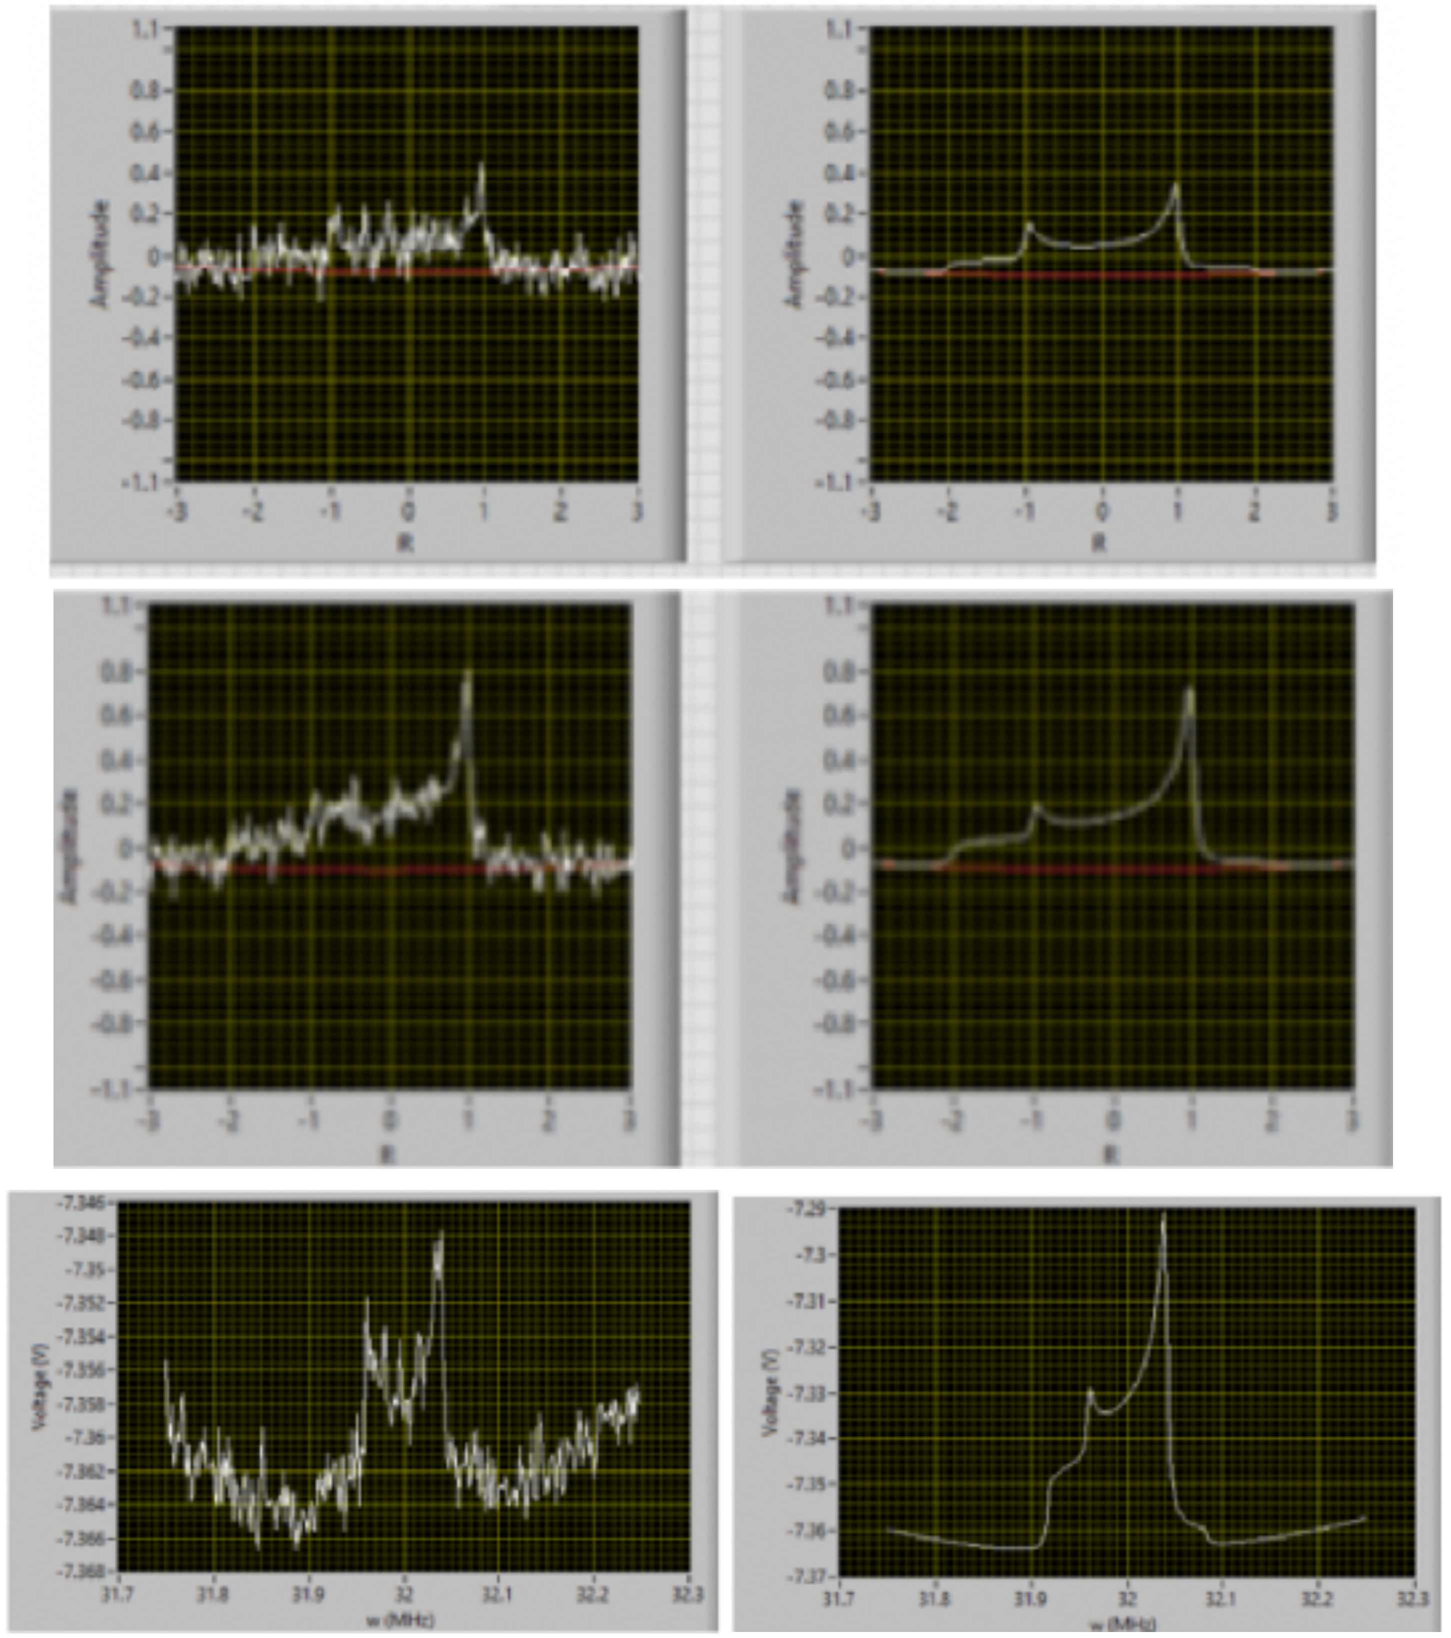

Devin Seay

Building a ANN

Simplified

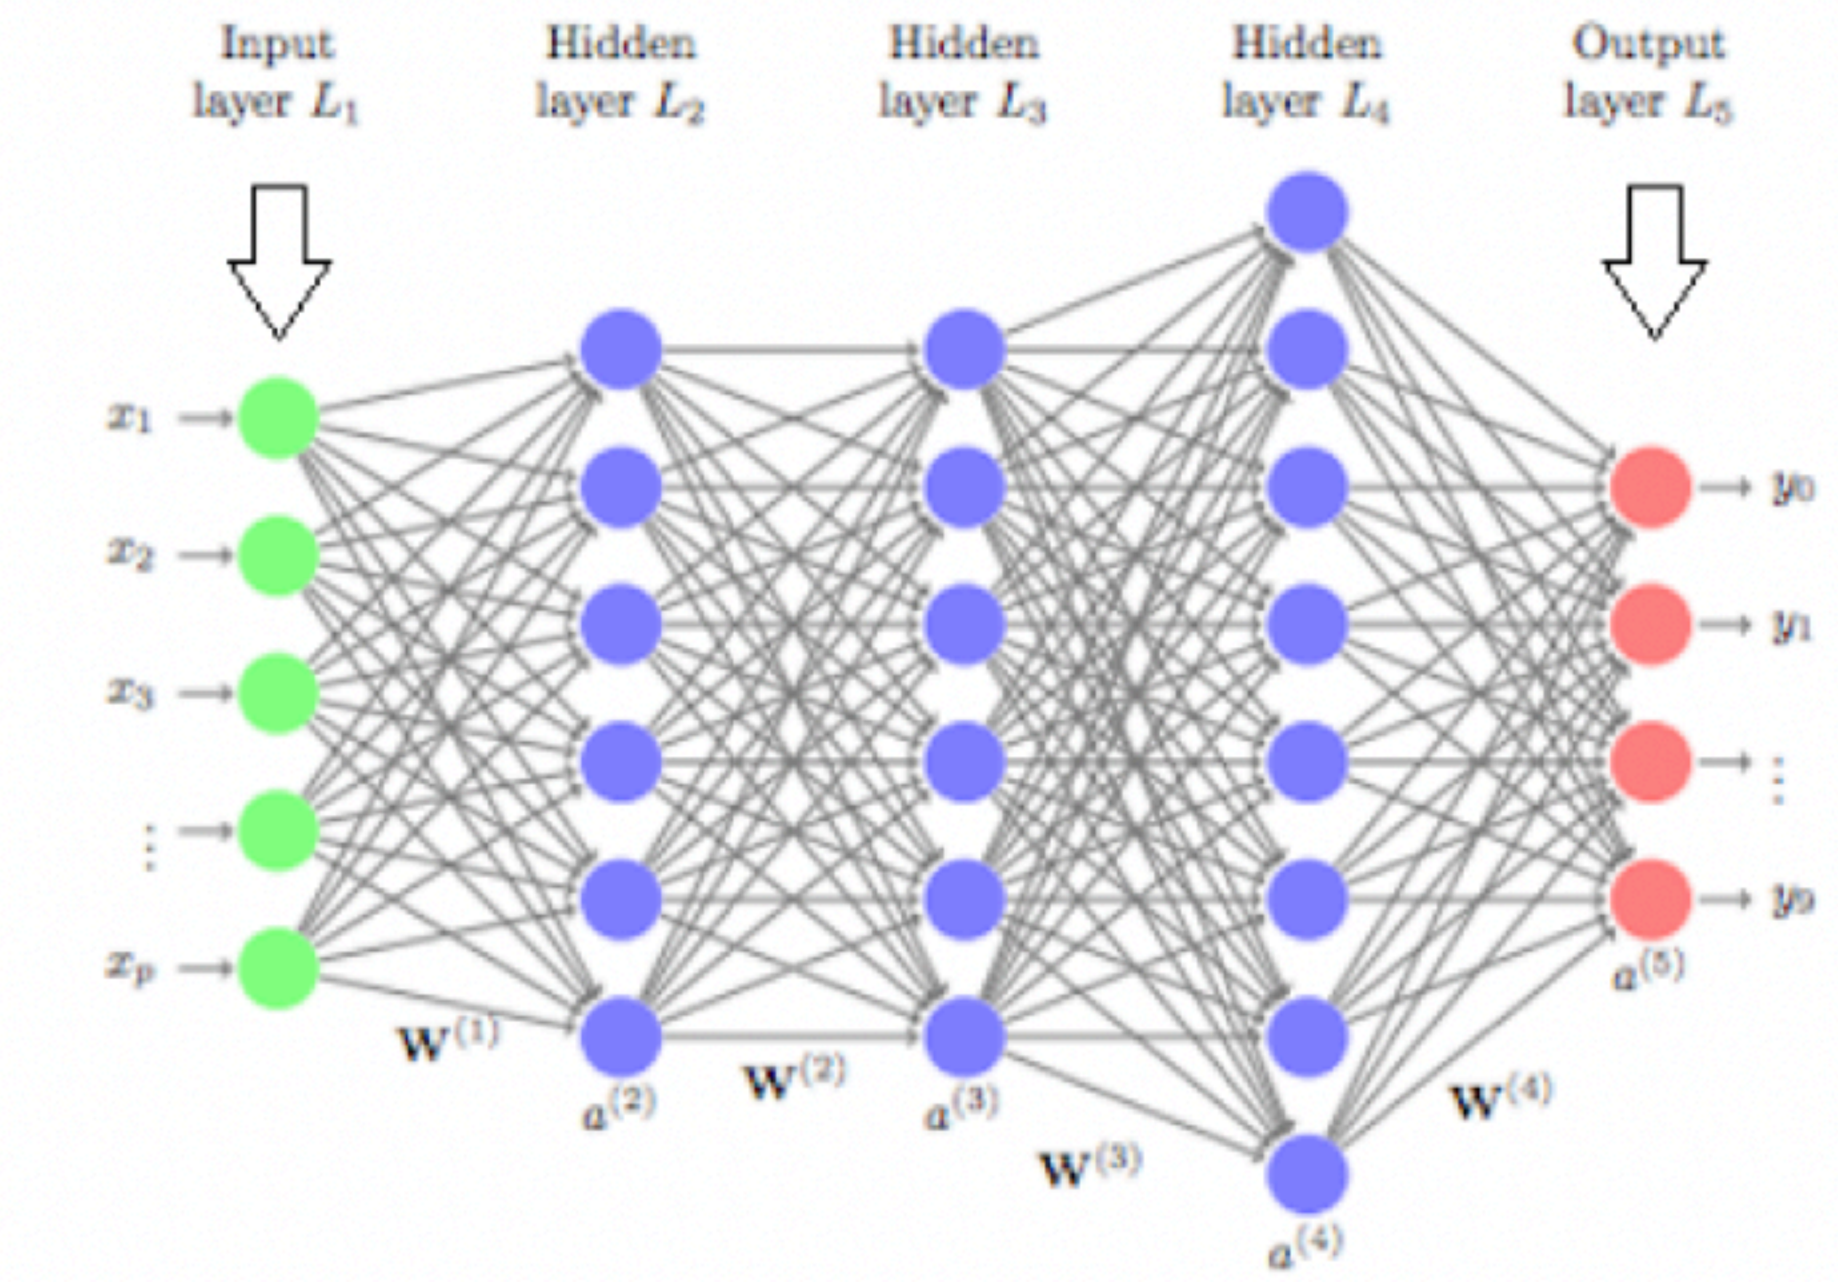

- Determine type of problem
- Prepare (engineer training) data
- Configure ANN (architecture, activation, loss, optimizer, learning rate, batch size)
- Train ANN
- Improve performance (tune hyperparameters)
- Test, study, probe limitations,…

Rate Response

Or why you can get rid of models

$$A_{gained} = \frac{1}{2}A_{lost}$$

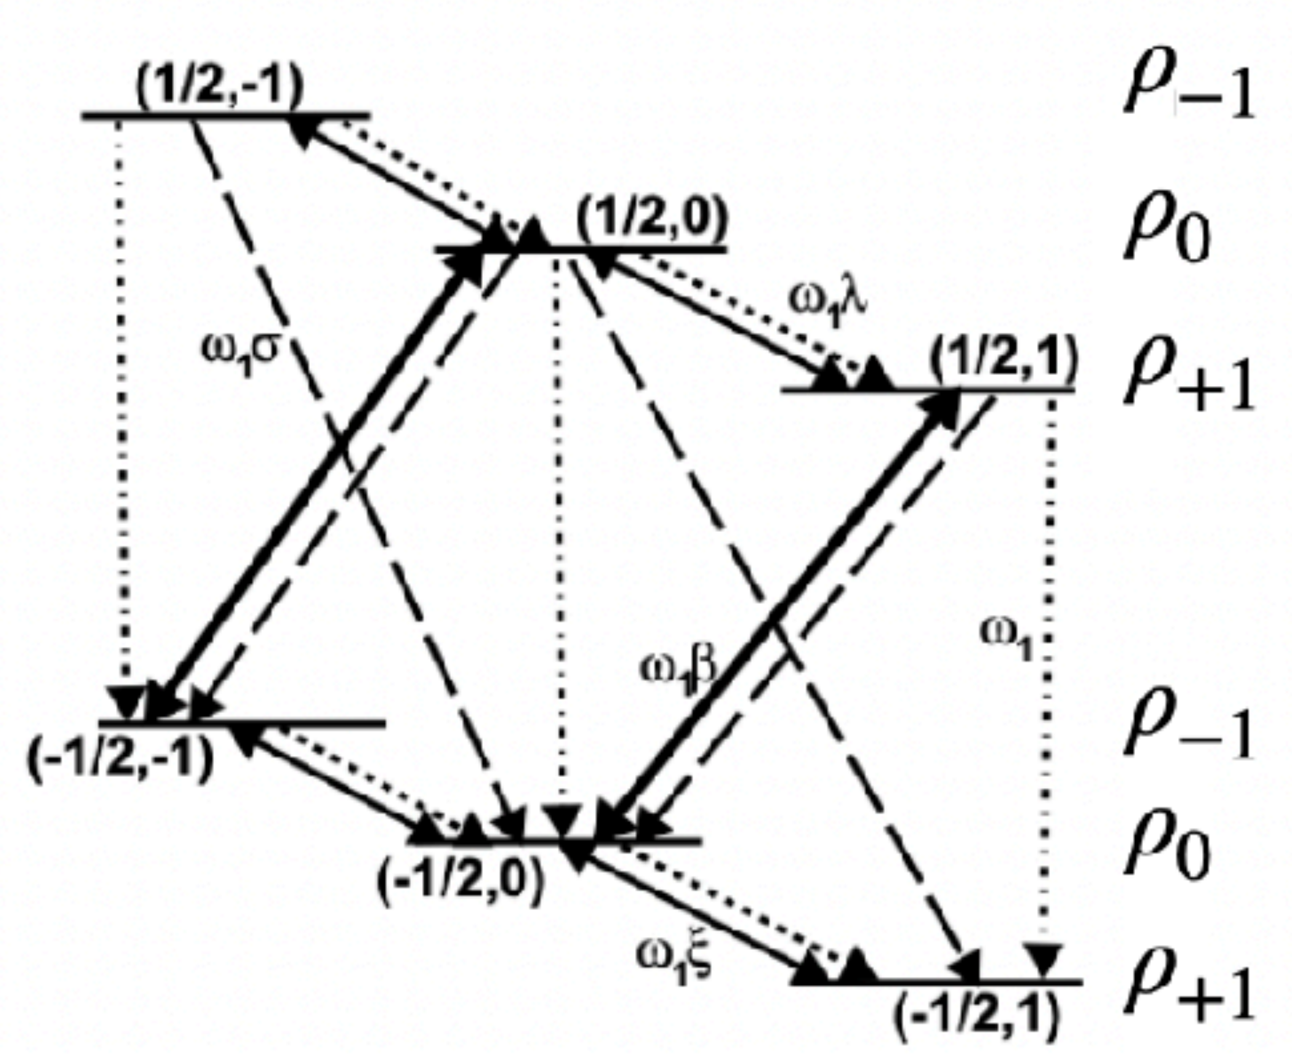

$$I_-(-\mathscr{R}) = C(\rho_0 - \rho_-)$$

$$I_-(-\mathscr{R}) - \dot{I}_-(-\mathscr{R})$$

$$= C[(\rho_0 - \xi\rho_0) - (\rho_- + \xi\rho_0)]$$

$$= C[(\rho_0 - \rho_-) - 2\xi\rho_0]$$

$$\dot{I}_-(-\mathscr{R}) = -2C\xi\rho_0$$

$$I_+(-\mathscr{R}) = C(\rho_+ - \rho_0)$$

$$I_+(-\mathscr{R}) - \dot{I}_0(-\mathscr{R})$$

$$= C[(\rho_+ - \xi\rho_+) - (\rho_0 + \xi\rho_+)]$$

$$= C[(\rho_+ - \rho_0) - 2\xi\rho_+]$$

$$\dot{I}_+(-\mathscr{R}) = -2C\xi\rho_+$$

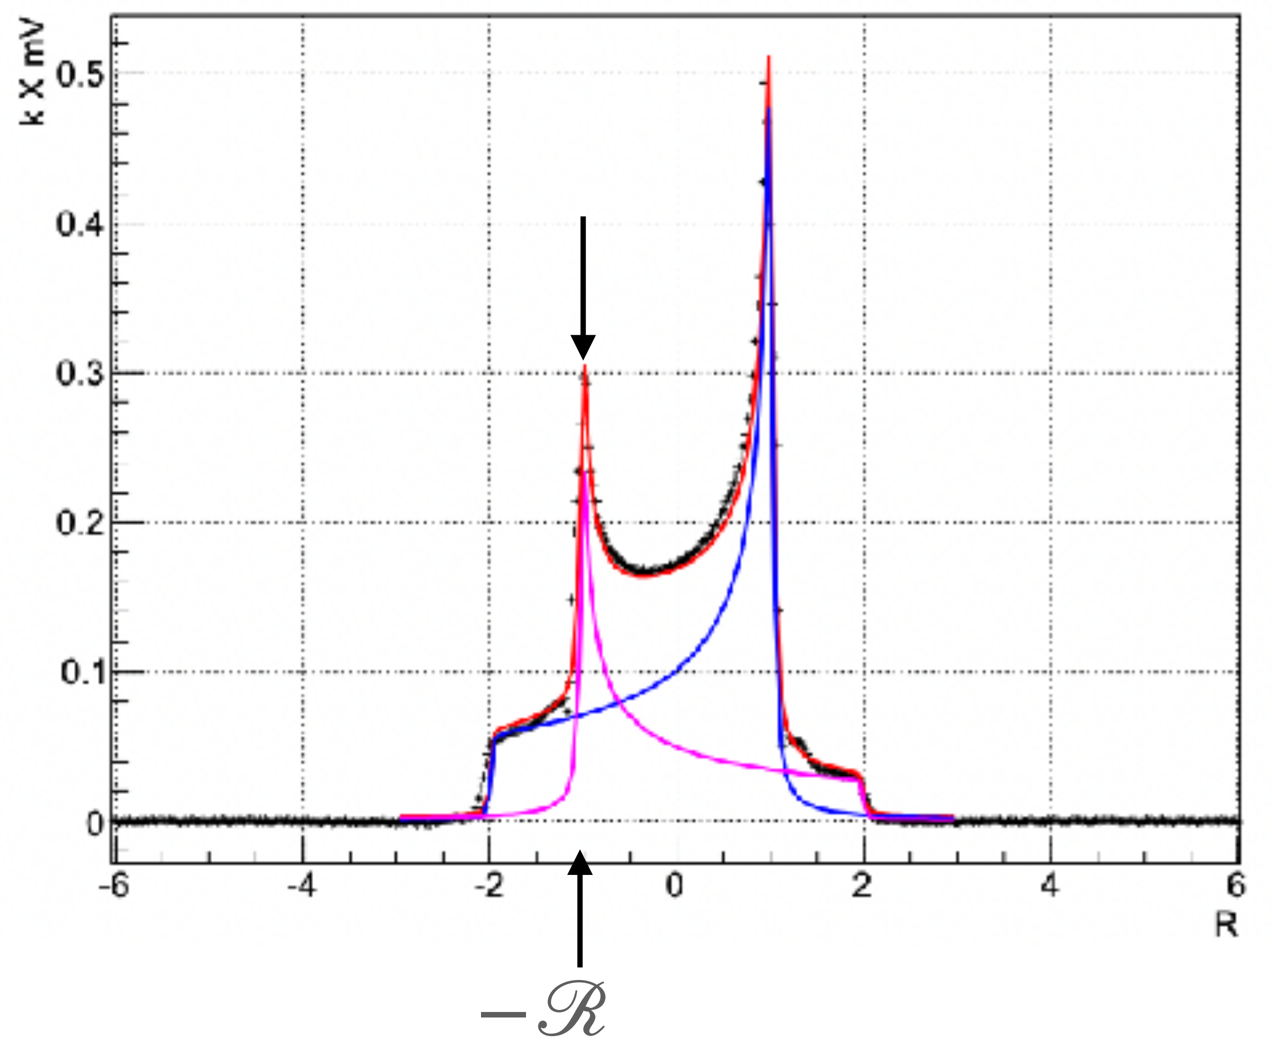

Rate Response

Or why you can get rid of models

$$I_-(-\mathscr{R}) = C(\rho_0 - \rho_-)$$

$$I_+(\mathscr{R}) + \dot{I}_+(\mathscr{R})$$

$$= C[(\rho_+) - (\rho_0 - \xi\rho_0)]$$

$$= C[(\rho_+ - \rho_0) + \xi\rho_0]$$

$$\dot{I}_+(\mathscr{R}) = C\xi\rho_0$$

$$I_+(-\mathscr{R}) = C(\rho_+ - \rho_0)$$

$$I_-(\mathscr{R}) + \dot{I}_-(\mathscr{R})$$

$$= C[(\rho_0 + \xi\rho_+) - (\rho_-)]$$

$$= C[(\rho_0 - \rho_-) + \xi\rho_+]$$

$$\dot{I}_-(\mathscr{R}) = C\xi\rho_+$$

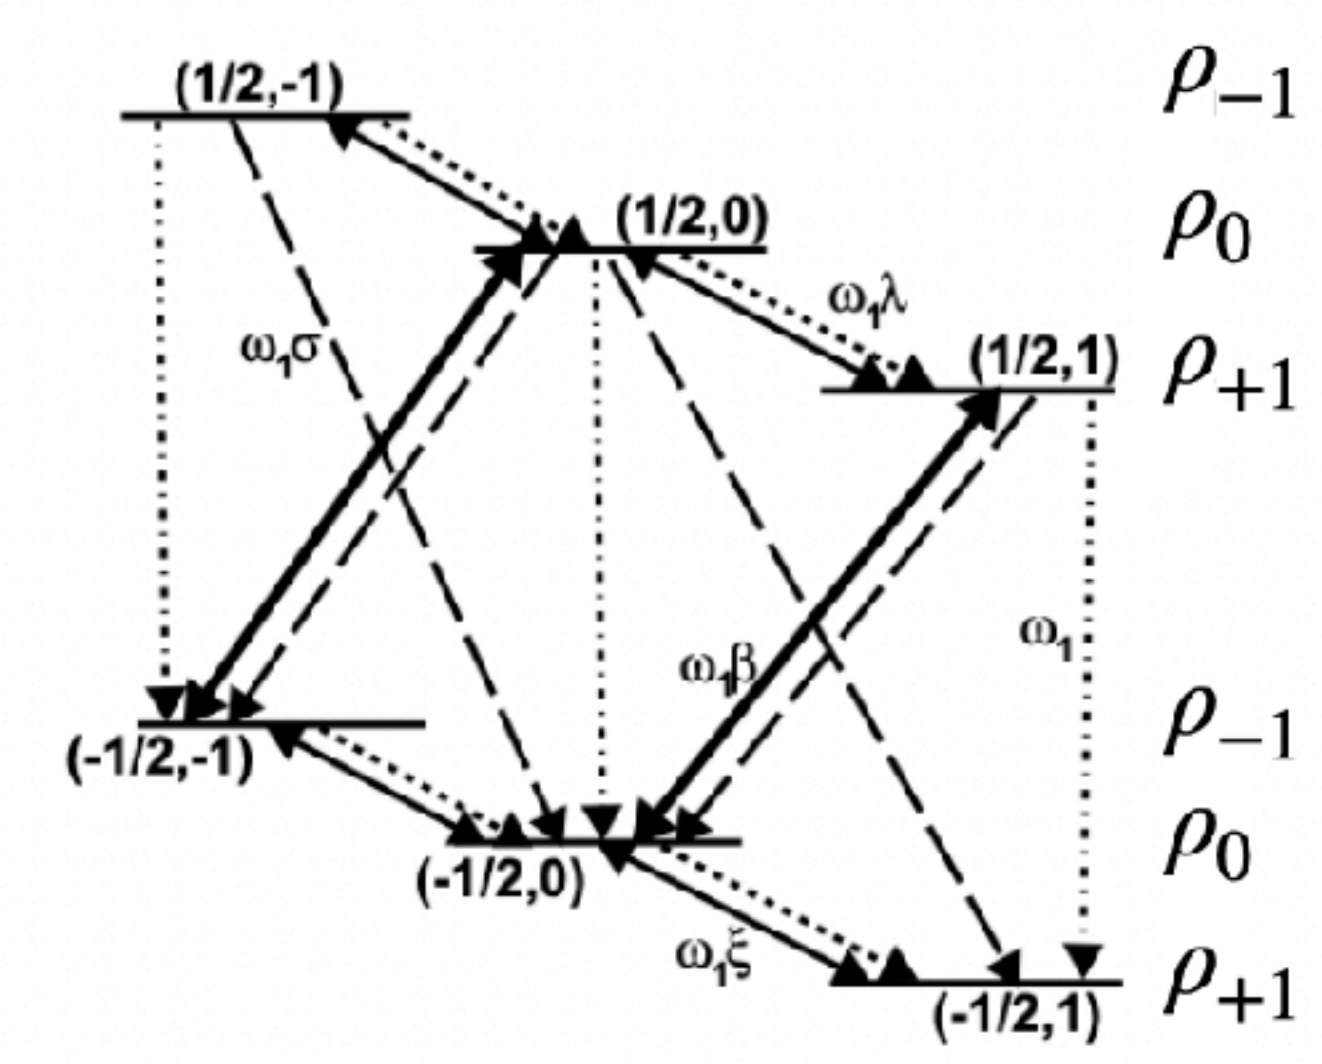

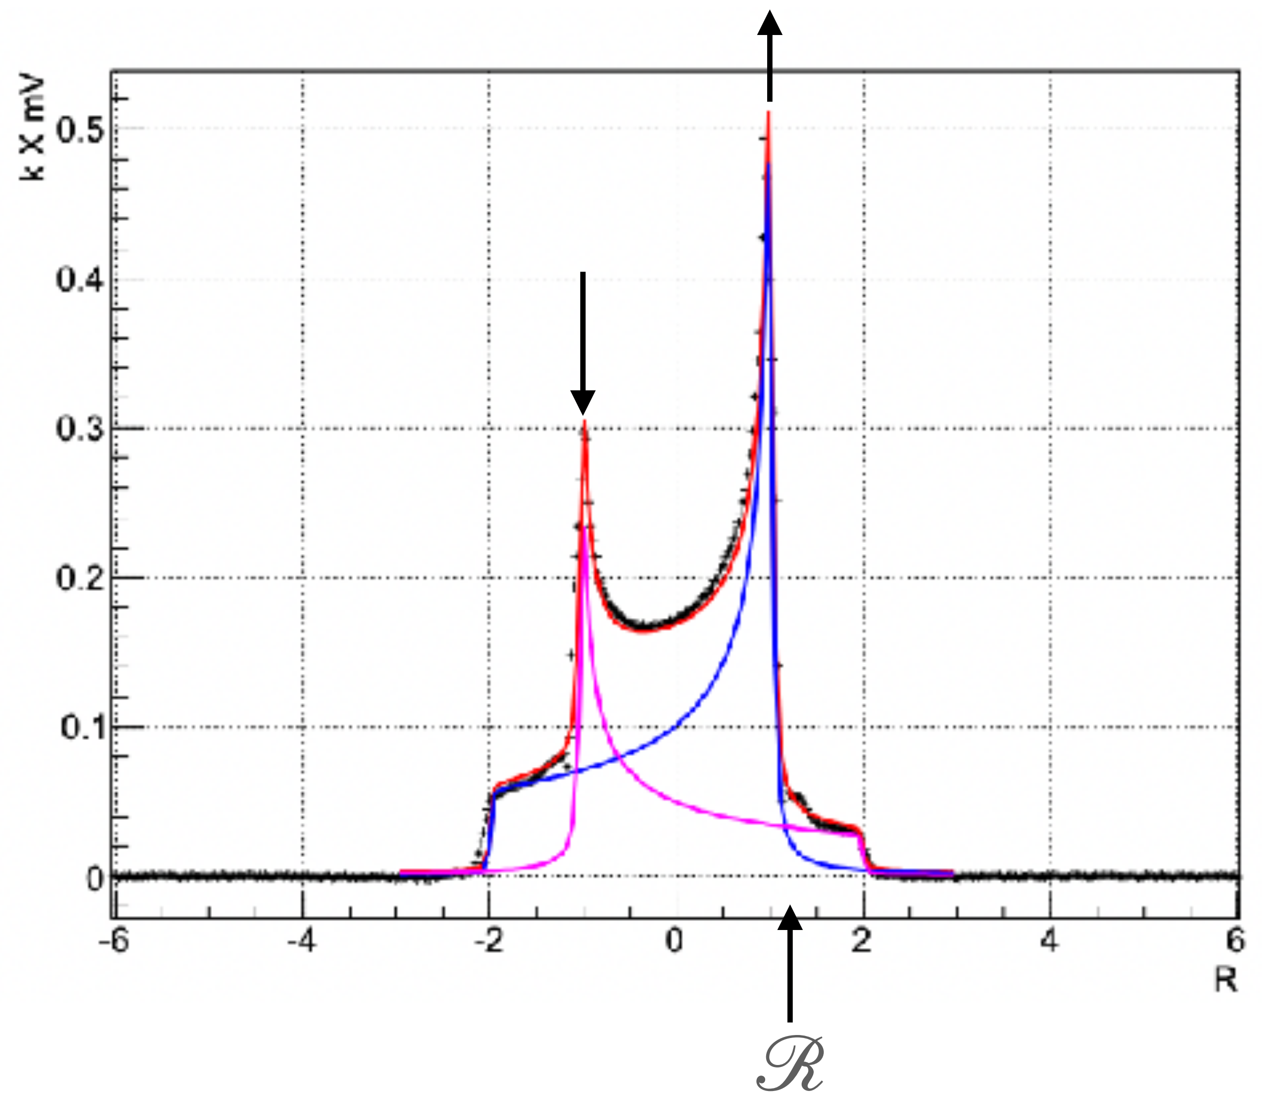

Some Important ANNs

With applications to our field

- Feedforward Neural Networks
- Radial basis function Neural Networks
- Self Organizing Neural Networks
- Recurrent Neural Networks
- Convolution Neural Networks
- Modular Neural Networks
- Graph Neural Networks

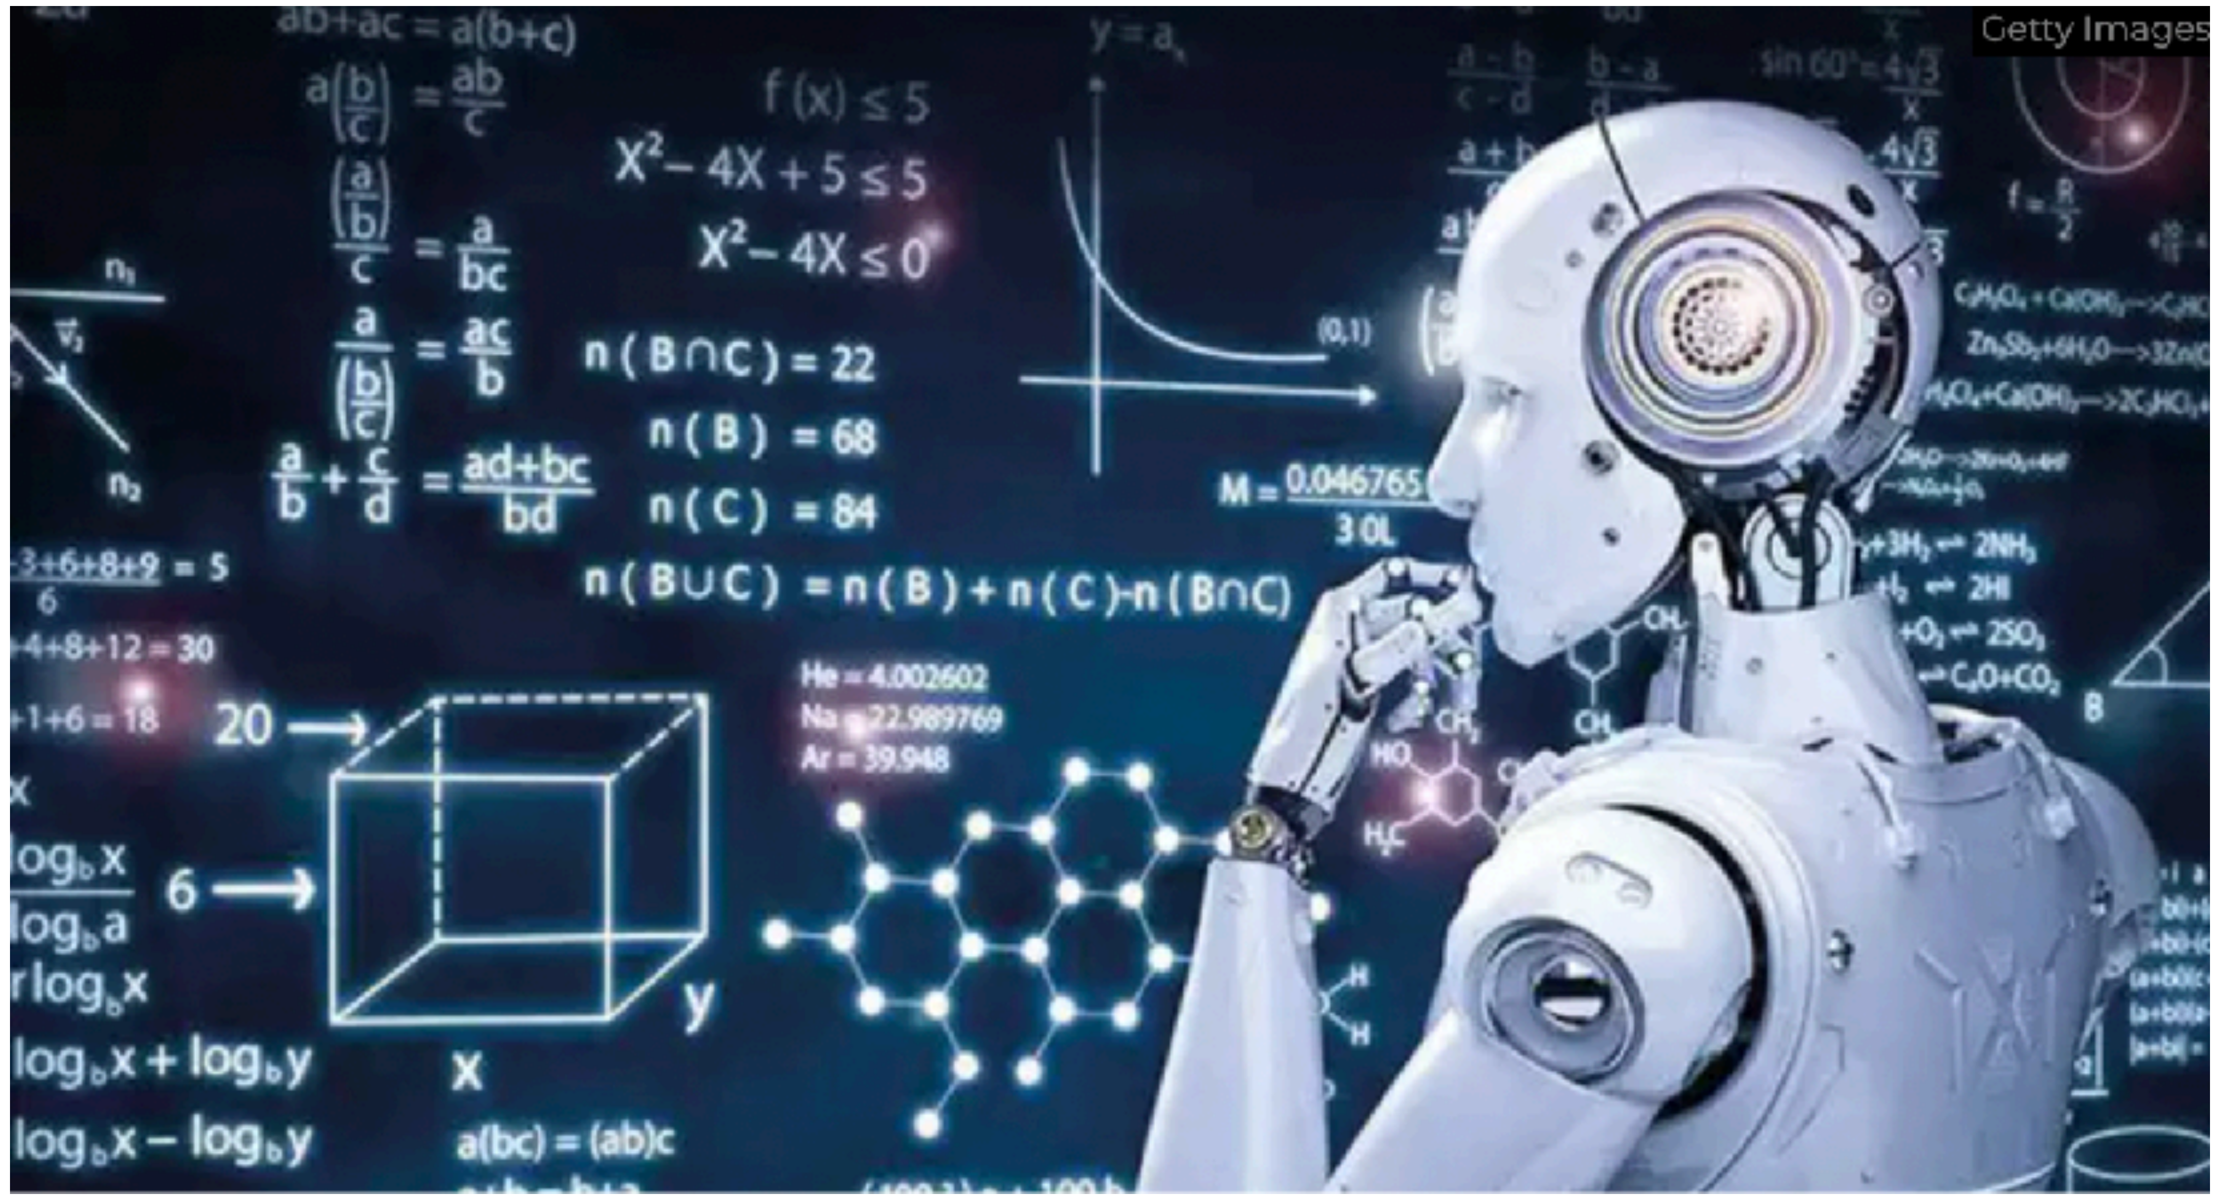

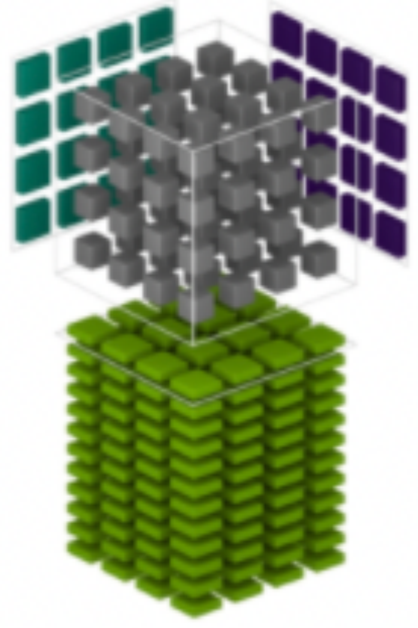

NVIDIA Tensor Cores

Caffe Caffe2 Chainer

MATLAB Microsoft Cognitive Toolkit mxnet

PaddlePaddle PyTorch

TensorFlow torch Wolfram Language

GPGPU Computing

Bogdan Oancea, Tudorel Andrei, Raluca Mariana Dragoescu
arXiv:1408.6923v1

Processing Units (AI Chips)

Acceleration in Training and Inference (driving the rate of progress)

- CPU: **Central Processing Unit** also called a **central processor**, **main processor** or just processor, is the electronic circuitry that executes instructions in a clock cycle. Need a farm to do serious AI.
- GPU: **Graphics Processing Unit** is a specialized processor with dedicated memory that conventionally perform floating point operations required for rendering graphics. Its a single-chip processor used for extensive graphical and mathematical computations for major parallelization.
- TPU: **Tensor Processing Unit** is an AI accelerator application-specific integrated circuit (ASIC) developed by Google specifically for neural network machine learning, particularly using Google's own TensorFlow software.
- NPU: N**eural Processing Unit** is a microprocessor that specializes in the acceleration of machine learning algorithms, typically by operating on predictive models such as artificial neural networks (ANNs) or random forests (RFs).
- QPU: Q**uantum processing Unit**, also referred to as a quantum chip, is a physical (fabricated) chip that contains a number of interconnected qubits. It is the foundational component of a full quantum computer, which includes the housing environment for the QPU, the control electronics, and many other components.