SAND2020-7929C

Bayesian Analysis Techniques Part 2: Deep Learning in the Inference Loop

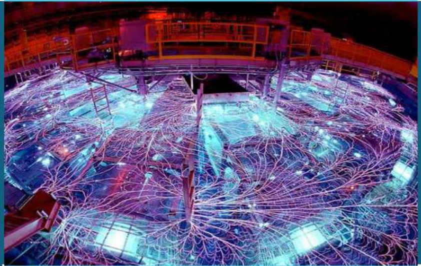

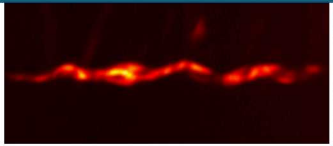

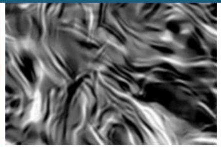

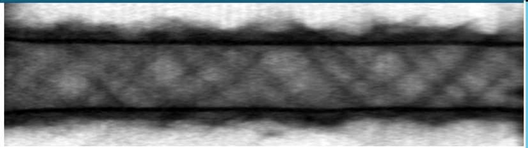

PRESENTED BY

William Lewis

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2020-####

Contributors

Pat Knapp

Matt Gomez

Adam Harvey-Thompson

Paul Schmit

Steve Slutz

Dave Ampleford

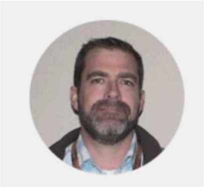

Magnetized liner inertial fusion relies on three stages to reach fusion relevant conditions.

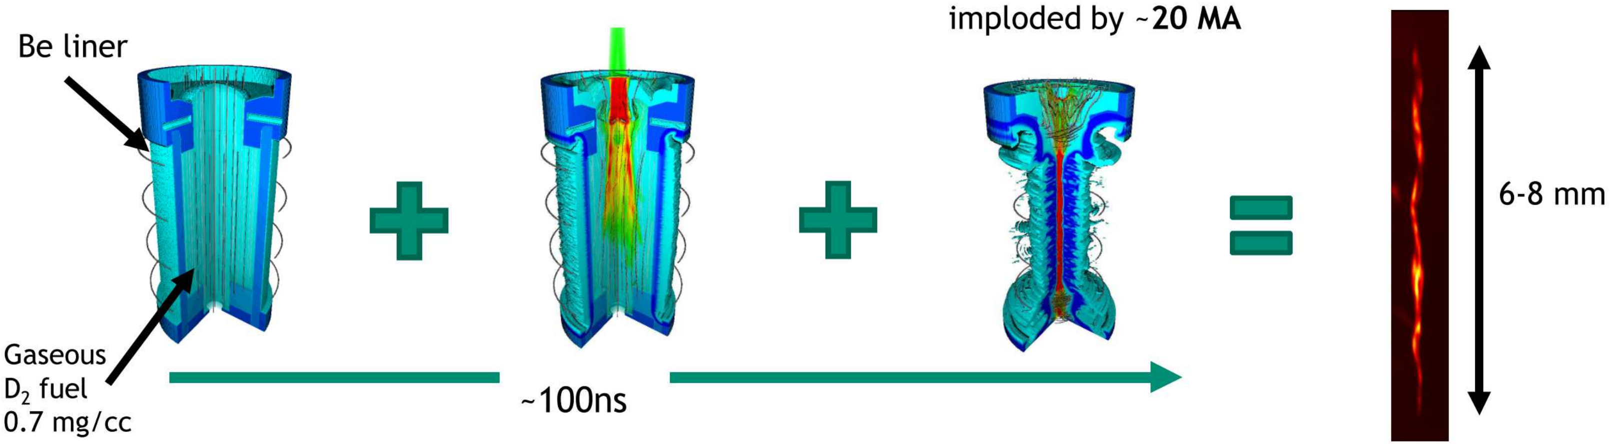

Magnetization

- Suppress radial thermal conduction losses

Preheat

- Increase fuel adiabat to limit required convergence
- Ionize fuel to lock in B-field

Implosion

- PdV work to heat fuel
- Amplify B-field through flux compression

Stagnation

- Several keV temperatures
- Several kT B-field to trap charged fusion products

S.A. Slutz *et. al.*, PoP (2010)

Secondary neutron yield and spectra are sensitive to BR and indicate a path toward measurement.

$$\text{Primary}: D + D \rightarrow \begin{cases} n + {}^3\text{He} & 2.45\ \text{MeV} \\ p + T, & 1.01\ \text{MeV} \end{cases}$$

$$\text{Secondary}: D + T \rightarrow n + \alpha. \quad 11.8\text{–}17.1\ \text{MeV}$$

$$\mathcal{P}_{DT} \propto \langle \rho_D \ell \rangle \sigma_{DT} \xrightarrow{\text{Magnetized}} \ell \propto f(BR)$$

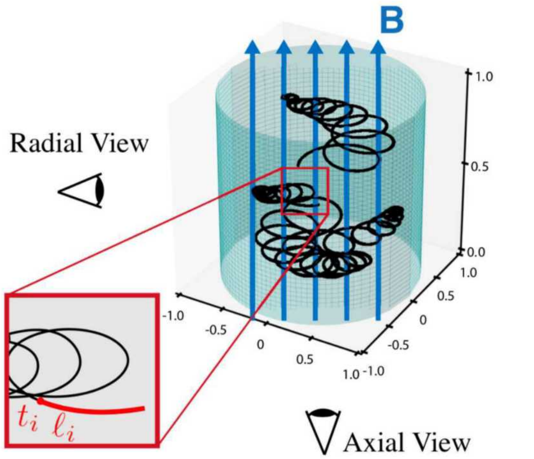

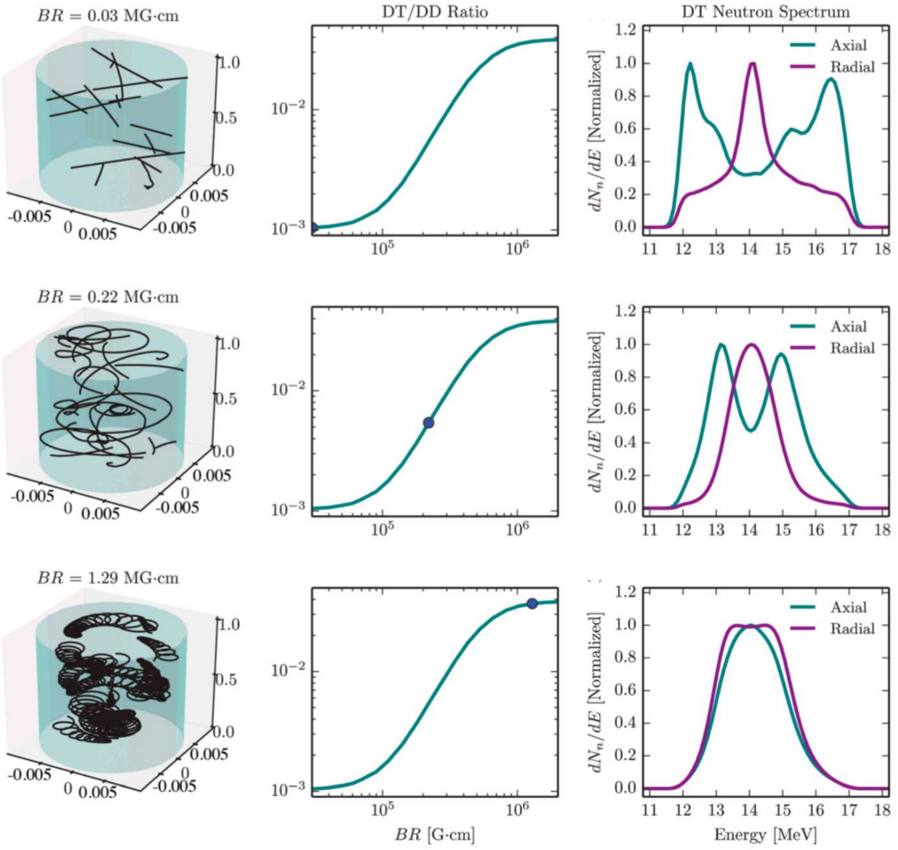

P.F. Knapp *et. al.*, PoP (2015)

We cannot directly measure neutron spectra, but we can observe the Time-of-Flight (ToF).

Neutron source

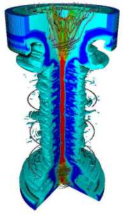

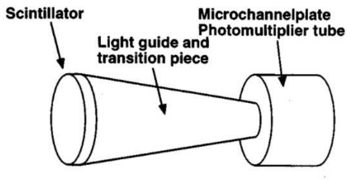

Fixed detector distance allows a change of variabls $E \rightarrow NToF$

Computed secondary neutron spectrum

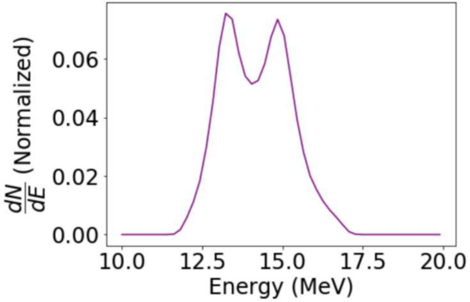

Computed NTOF

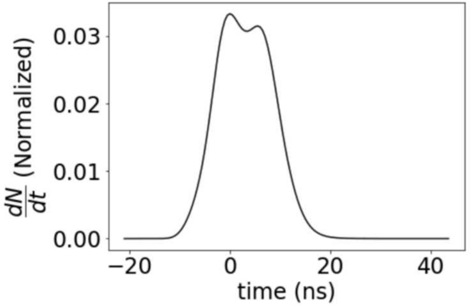

Experiment analysis for BR is now much closer to "push button"

Stage requiring minimal human input with low impact on results

Machine learning and Bayesian inference with no human input

Obtain secondary neutron signals

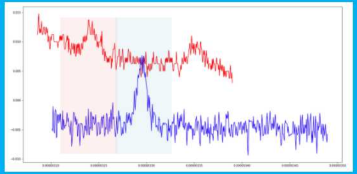

Select Signal and background ROIs

Provide DT and DD yields with uncertainty as well as nToF

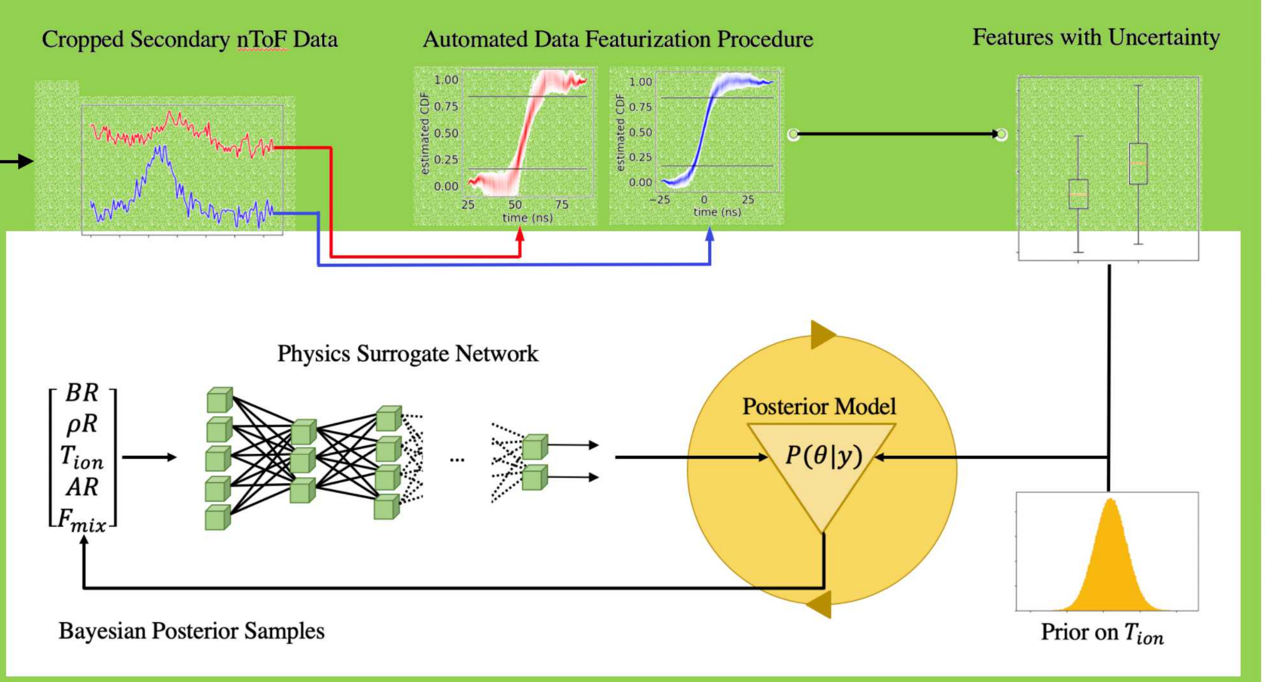

To extract range of parameters consistent with observations we perform Bayesian posterior inference.

Stage requiring minimal human input with low impact on results

Machine learning and Bayesian inference with no human input

Obtain secondary neutron signals

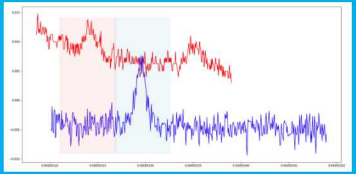

Select Signal and background ROIs

Provide DT and DD yields with uncertainty as well as nToF

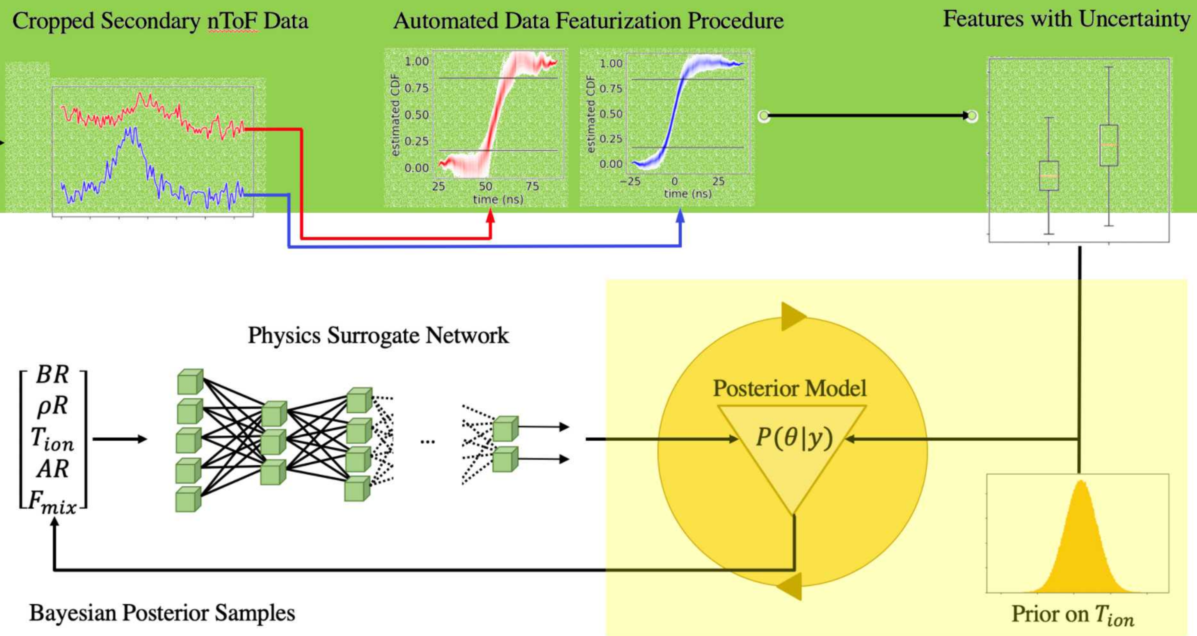

Bayes theorem allows us to easily incorporate multiple datasets to pin down model parameters with uncertainty.

Let A represent model parameters and B represent observed data we want to model

Likelihood: probability distribution for data under our model. E.g. measurement gaussian distributed about mean or obeys Poisson distribution

Prior: probability distribution for model parameters before incorporating data. Can encode physics constraints, ignorance, etc.

$$P(A \mid B) = \frac{P(B|A)P(A)}{P(B)}$$

Evidence: probability distribution for data. Typically unknown, but also typically unimportant.

Posterior: probability distribution for model parameters taking into account observations. This is typically what we want.

Our Bayesian model incorporates models for most sources of uncertainty.

- Uncertainty in forward model due to use of surrogate

$$\vec{y}(\theta) = \vec{y}_{nn}(\theta) + \mathrm{N}(0, \Sigma_{OOS})$$

- Uncertainty in observed values (DD yield, DT yield, quantile features)

$$\vec{y}_{Quant} = \langle \vec{y}_{Quant} \rangle + \mathrm{N}(0, cov[\vec{y}_{Quant}, \vec{y}_{Quant}])$$

- Statistical uncertainty between model and data

$$\vec{y}_{Quant} = \vec{y}(\theta) + \mathrm{N}(0, \Sigma_{unk}^{diag})$$

- Not included:
  - Possible systematic uncertainty from model (would need to assess performance of different models)
  - Doesn’t contain uncertainty in NN parameters (in principle possible, but not likely to be dominant source)
  - Background model assumption
- Additional testing of normal error models should be performed
- Testing of assumption of uncorrelated model-observation error needs testing

The basic outline of the Bayesian inference problem highlights 3 key ingredients.

- Forward model
- Dimensionality of likelihood and efficient sampling
  - Work in reduced dimensional feature space for efficient sampling
- Data preprocessing (both experiment and simulation)
  - nTоF instead of spectra
  - Experiment is exhibits significant noise, need robustly extracted features

$$P(B|A) \rightarrow ||B - B^*(A)||$$

If $Dim(A)$ large $P(A|B)$ requires careful sampling

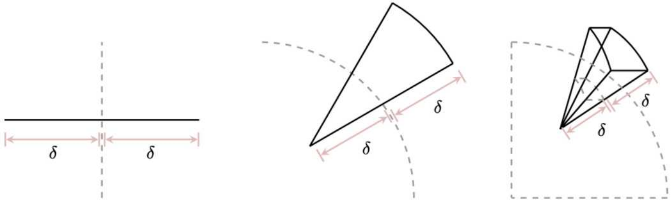

If $Dim(B)$ large $P(B|A)$ can have numeric issues (e.g. product of small numbers)

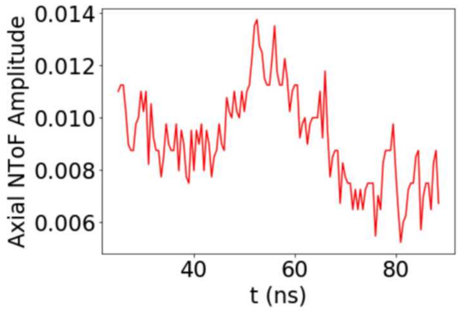

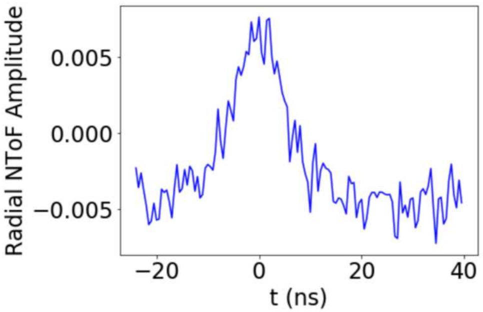

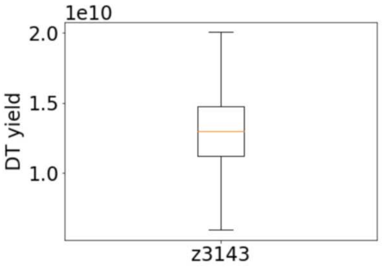

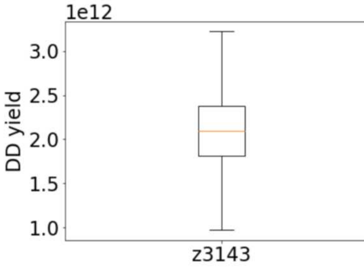

The physics model is treated with a neural network surrogate

Stage requiring minimal human input with low impact on results

Machine learning and Bayesian inference with no human input

Obtain secondary neutron signals

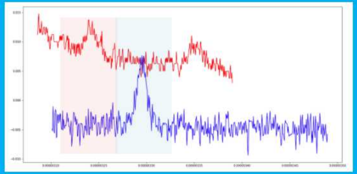

Select Signal and background ROIs

Provide DT and DD yields with uncertainty as well as nToF

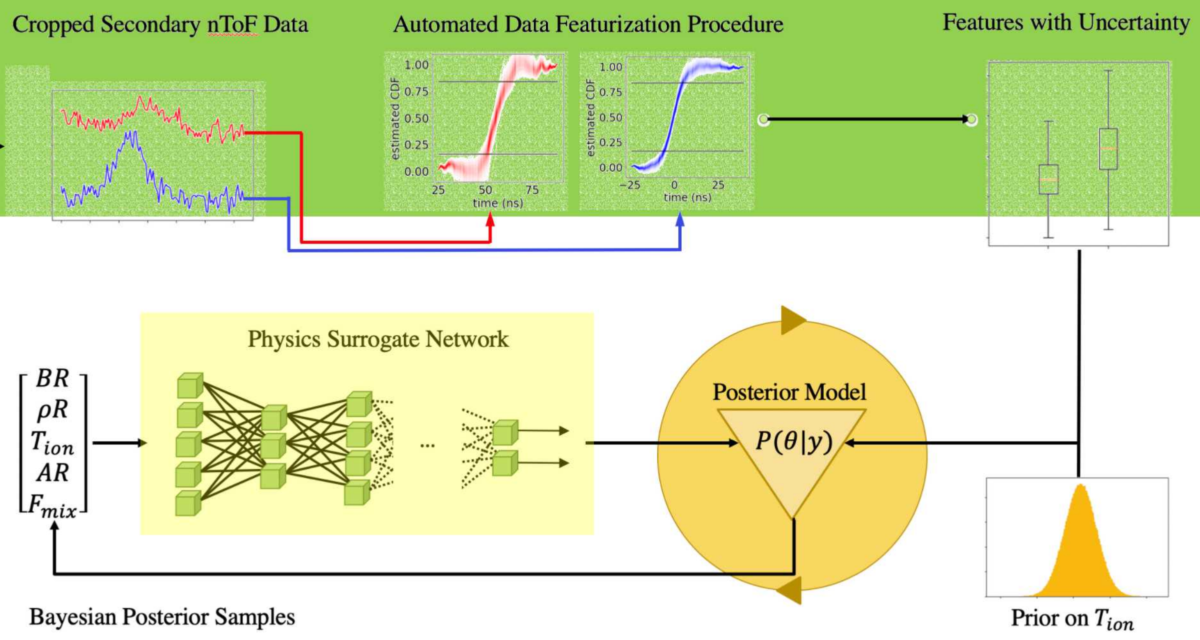

Surrogates allow incorporation of otherwise computationally unfeasible models into a Bayesian analysis.

- Single yield and spectra calculation takes O(1min-1hr)
- Bayesian inference requires potentially 10's-100's of thousands of forward model evaluations
  - Solution is to surrogate the calculation using a machine learning model to reduce forward model time to O(μs-ms)

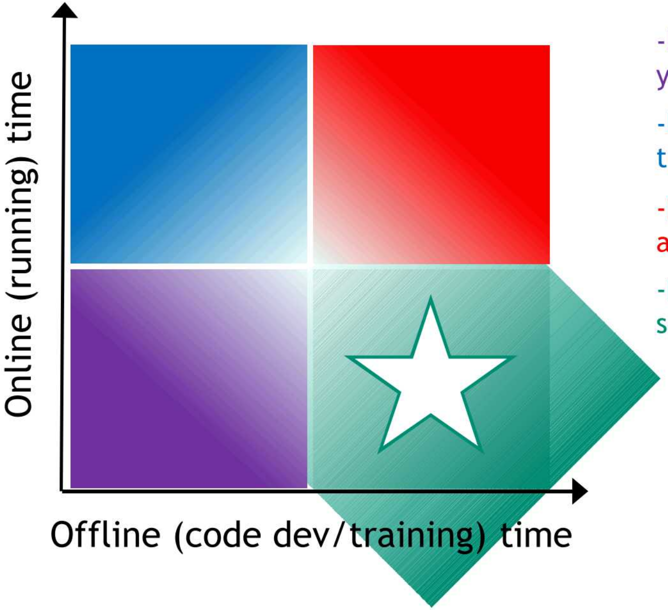

-Low offline low online: The problem is very easy, or you neglected/forgot to include the hard part

-Low offline high online: quick dev and little or no training but long running simulation

-High offline high online: The problem is very hard, and you are probably doing something inefficiently

-High offline low online: front load effort by training surrogate that runs fast. Pays off when

$$t_{FP} \geq \frac{t_{ML}}{1 - \frac{N_{Tr}}{N_B N_{Exp}}}$$

All machine learning methods may be boiled down to a few key components with which you will be familiar.

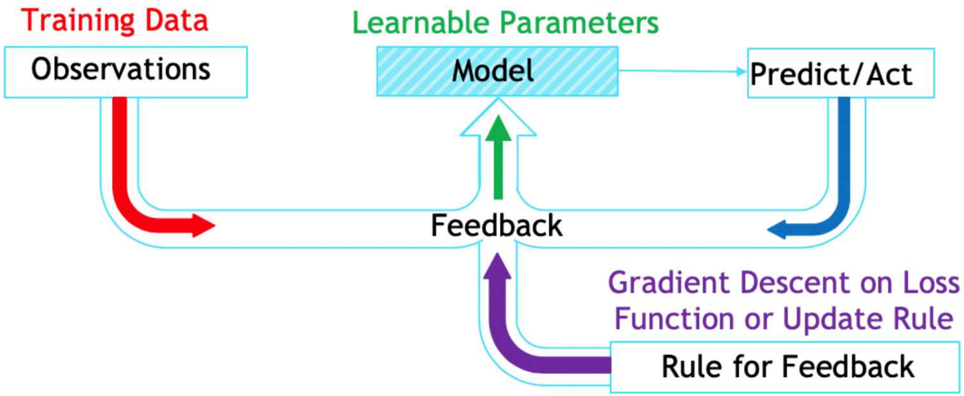

- Given a rule for using observation/predictions for feedback, we can update params
  - For regression/classification can define loss functions (e.g. MSE)
  - Update parameters by gradient descent to minimize loss

$$L_m = \frac{1}{N_{train}} \sum_{i=1}^{N_{train}} \left(y_i - f(x_i, \vec{\theta}_m) \right)^2 \qquad \theta_{m,i}^{new} = \theta_{m,i}^{old} - \eta \frac{\partial L_m}{\partial \theta_{m,i}^{old}}$$

- Will not go into detail, but can construct many reasonable loss functions
  - Should be bounded below (hence typically chosen to be non-negative)
  - Sometimes couched in statistical framework (see Bayesian stuff later)

An ~600 parameter neural network was trained on ~15k LHC sampled simulations regularized by inverse network.*

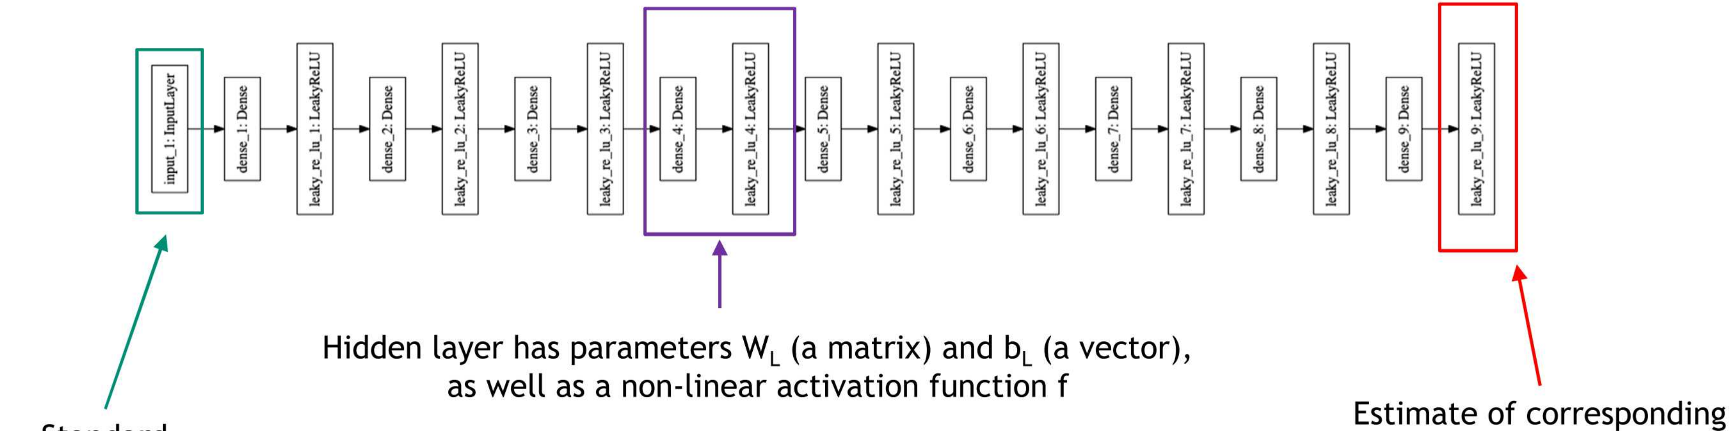

$$\vec{x}_L = f(W\vec{x}_{L-1} + b)$$

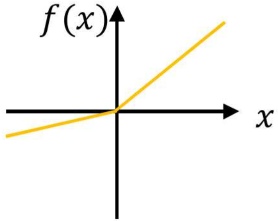

W and b found by minimizing square error between ground truth and prediction:

$$L = \frac{1}{N_{train}} \sum_{i=1}^{N_{train}} (y_i^{true} - y_{nn}^{pred}(\theta_i, W, b))^2 + (\theta_i^{true} - \theta_{nn}^{pred}(y_{nn}^{pred}, W_{inv}, b_{inv}))^2$$

*This has been shown to produce better behaved forward models (see e.g. R. Anirudh *et al.*, PNAS **117**, 9741-9746 (2019).)

Out-of-sample (OOS) validation shows we aren’t overfitting and provides OOS error estimates for the model fit.

Validation error ~ Training error indicating not overfitting

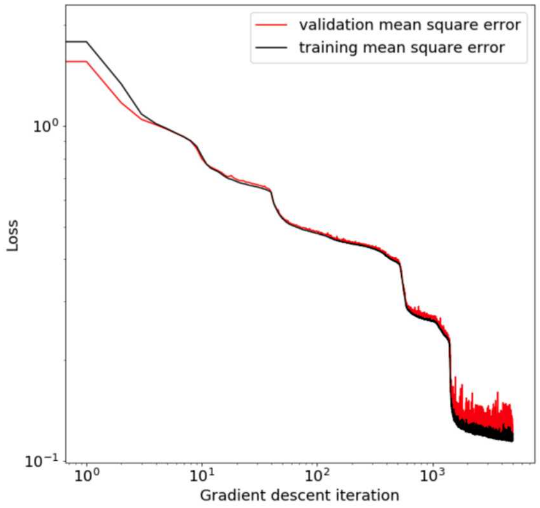

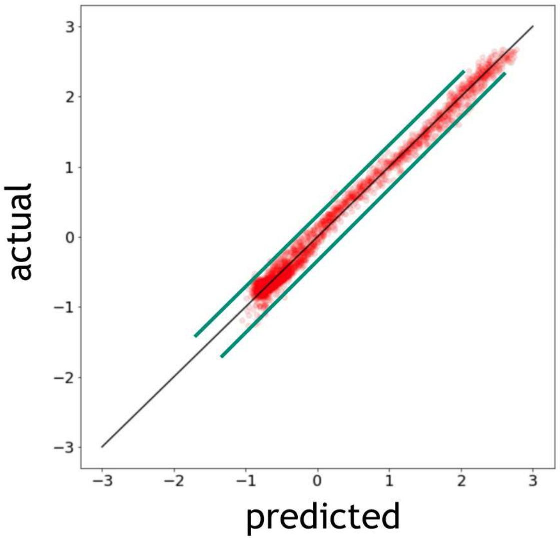

Estimate OOS covariance from performance on data not seen during training (~ 2k validation and ~14k training)

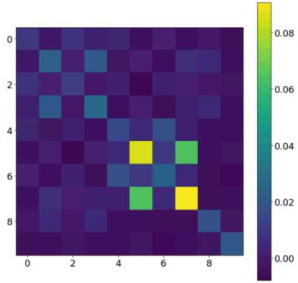

Used in Bayesian likelihood function

To minimize impact of human input, an automated data featurization is applied to signal

Stage requiring minimal human input with low impact on results

Machine learning and Bayesian inference with no human input

Obtain secondary neutron signals

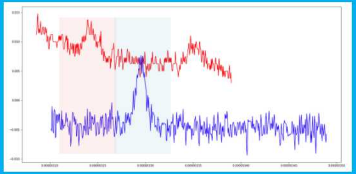

Select Signal and background ROIs

Provide DT and DD yields with uncertainty as well as nToF

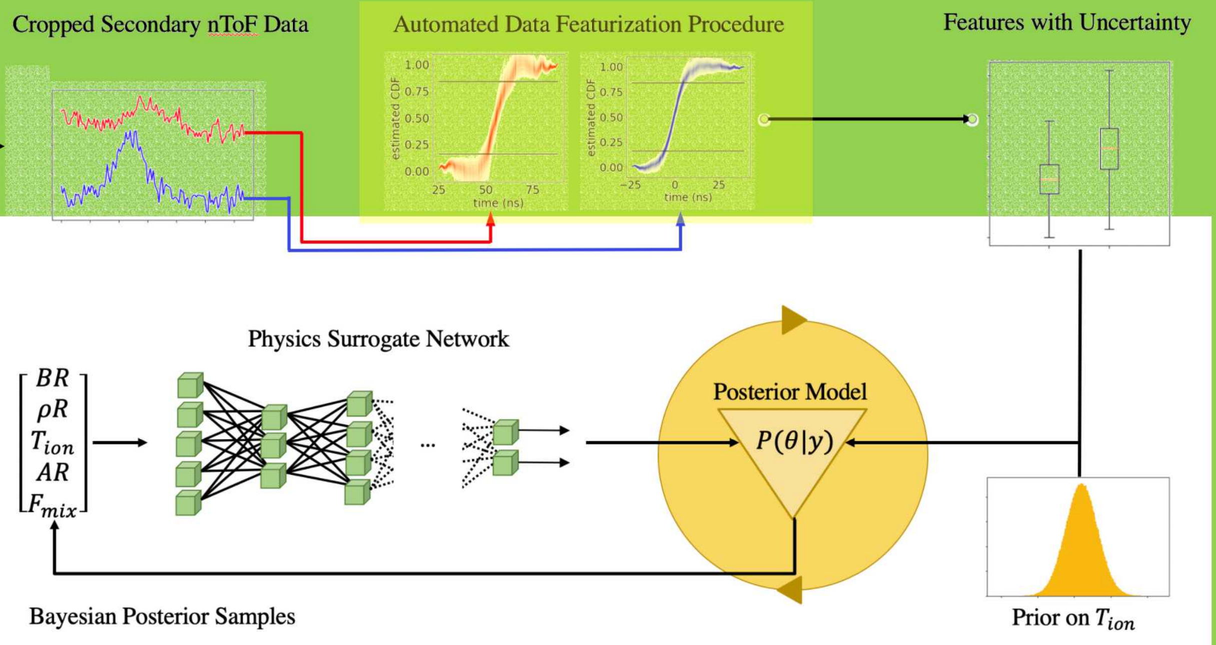

Experimental data exhibit significant noise which should be captured in uncertainty of features extracted.

Step 1: collect data from experiment

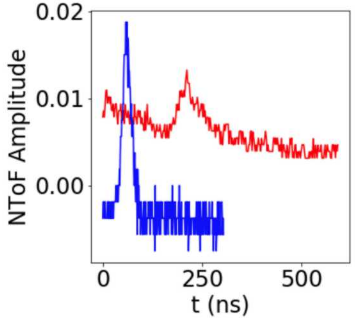

Step 2: crop and select background ROI

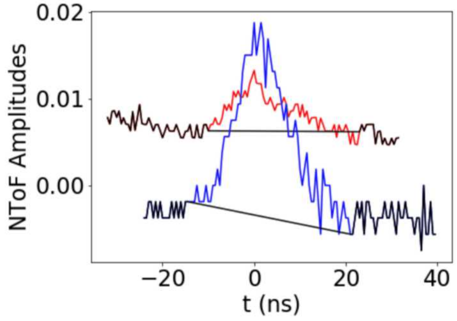

Step 3: Bayesian Background fit

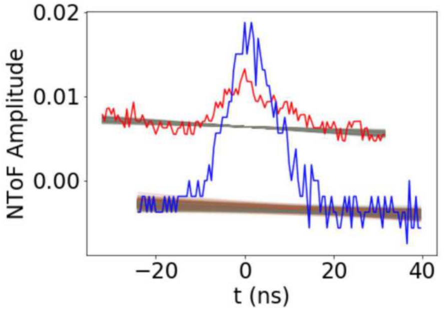

Step 4: Compute CDFs after subtraction

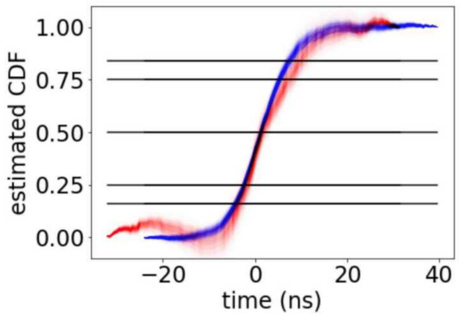

Step 5: Compute Quantile Features with Uncertainty

75%-25%

(75%-50%)-(50%-25%)

84%-16%

(84%-50%)-(50%-16%)

We can now perform Bayesian inference on model parameters given observations

Stage requiring minimal human input with low impact on results

Machine learning and Bayesian inference with no human input

Obtain secondary neutron signals

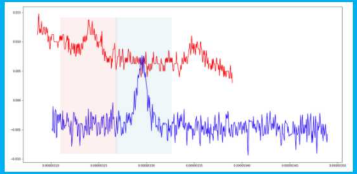

Select Signal and background ROIs

Provide DT and DD yields with uncertainty as well as nToF

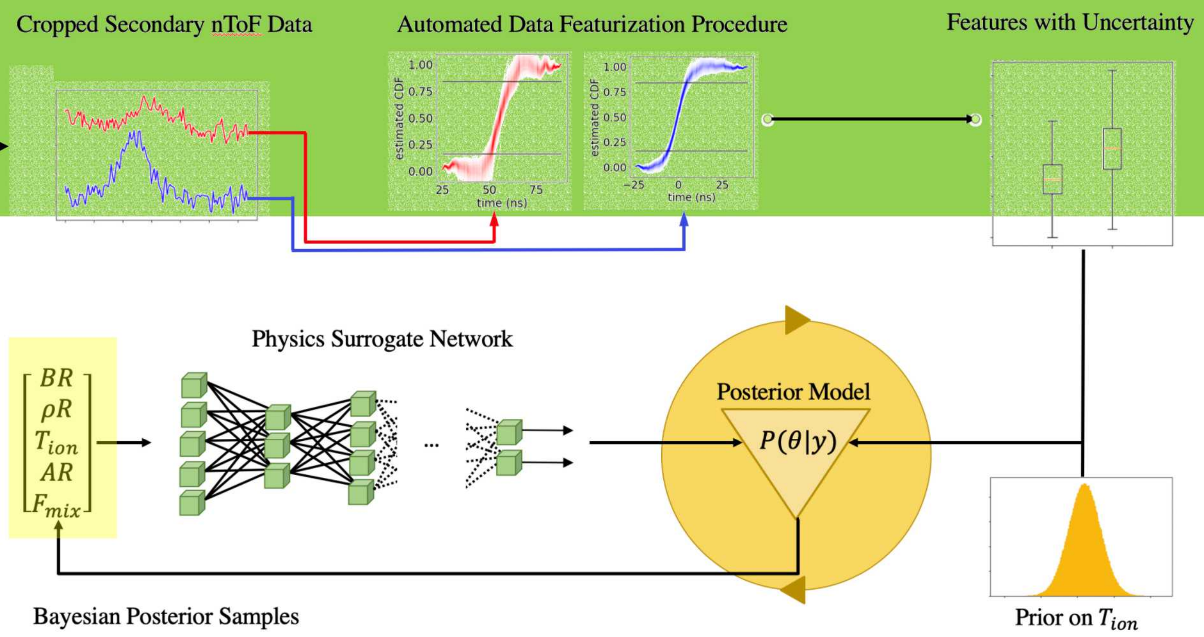

By utilizing a Markov Chain Monte Carlo (MCMC) sampling, we can obtain parameter distributions consistent with observed data.

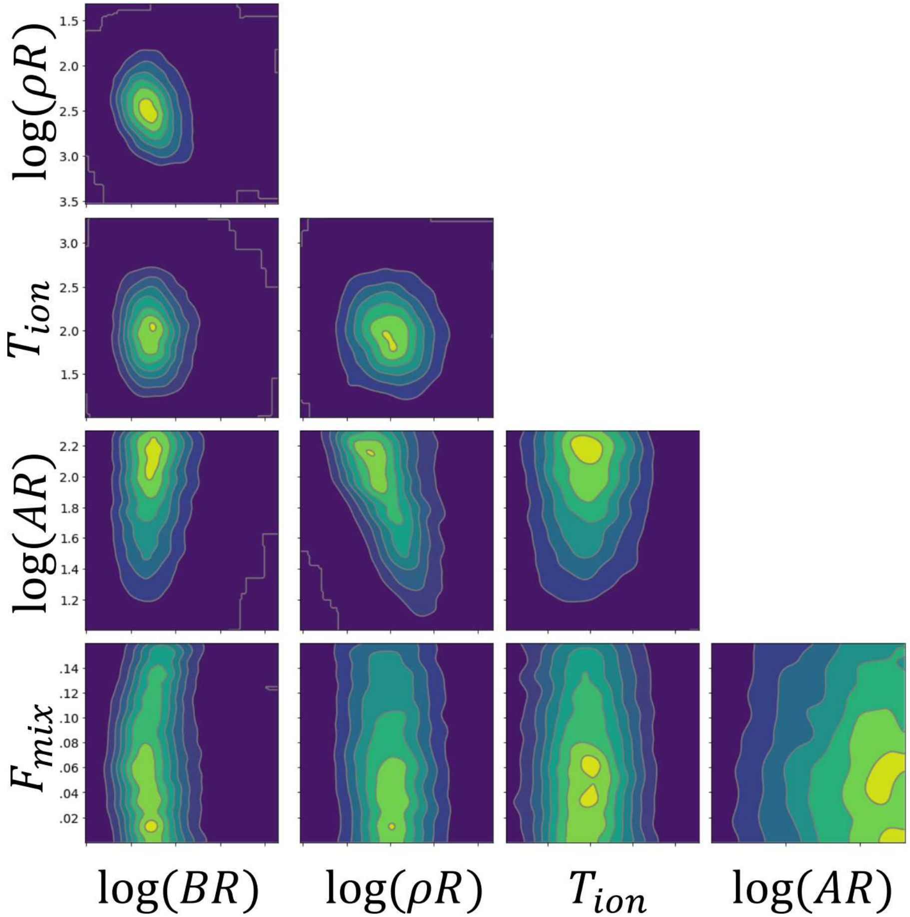

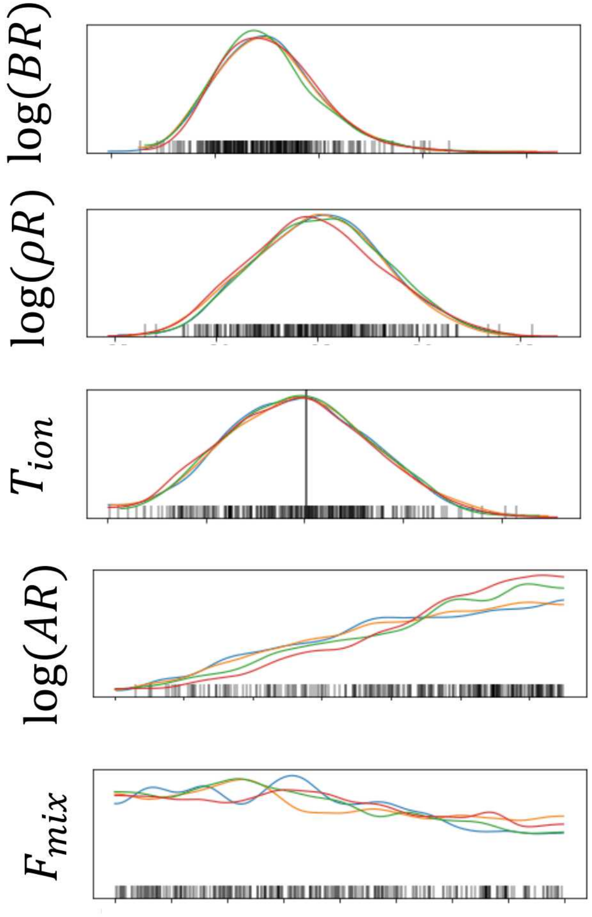

Experiments show trend consistent with Nernst effect.

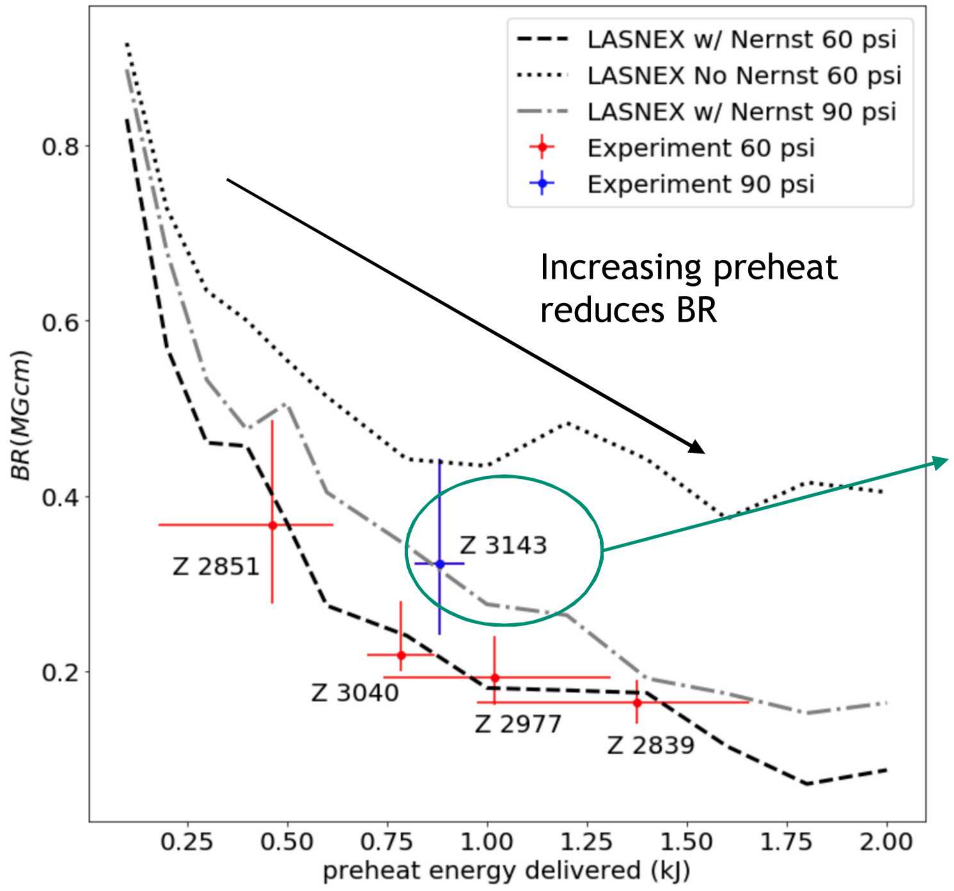

Caveat: General trend of declining BR vs preheat is extremely robust. Precise BR shows some sensitivity e.g. to featurization/surrogate architecture choices. May be observing density effect here which is exciting, but additional experimental data points and more detailed study of model selection should be carried out.

Closing remarks

- DL enabled Bayesian inference of BR for MagLIF shots
- Want to develop a database of BR for MagLIF shots to mine for trends
  - Already see interesting physics consistent with Nernst effect
- Plans to investigate
  - Fill density (already early indications?)
  - 3D nature of plasma
  - Instabilities
  - Mix
  - Impact of uncertainty
  - Scaling aspects of Nernst effect

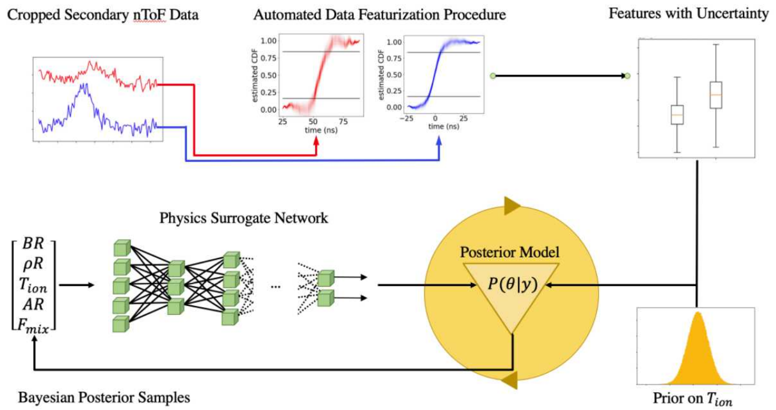

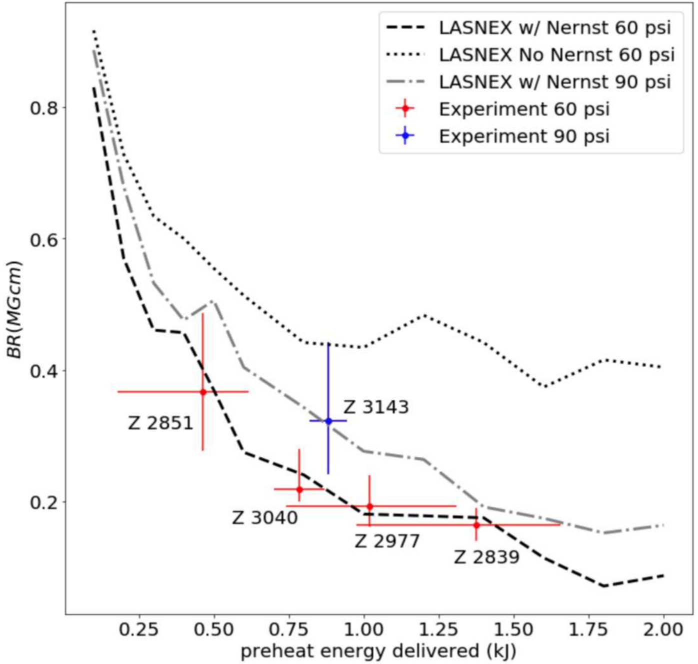

Backup

ToF is measured, so we convert model spectra to NToF.

Computed NTOF

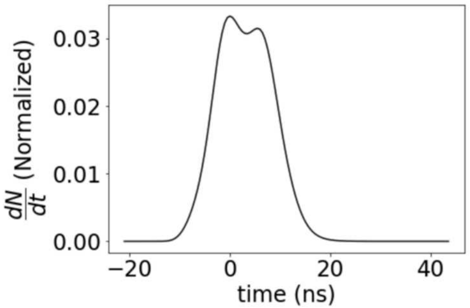

Computed secondary neutron spectrum

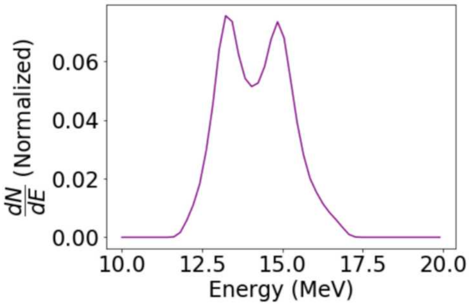

Jacobian from special relativity

$$\left.\frac{dN}{dt}\right|_{exp} = BH(t) * IRF(t) * \frac{dN}{dt} = BH(t) * IRF(t) * LO(E) \cdot \frac{dN}{dE}\,\frac{dE}{dt}$$

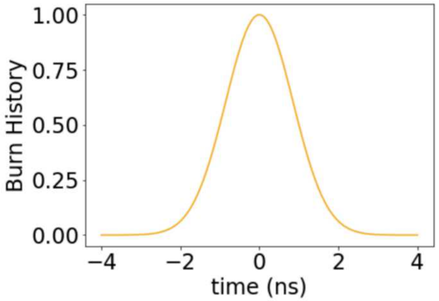

Burn history (neutron production localized in time)

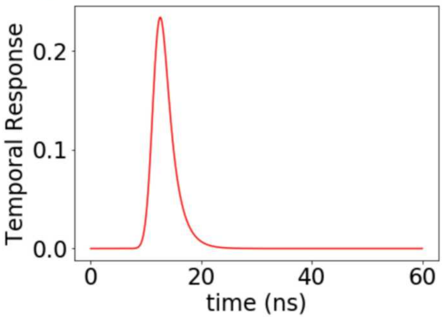

Instrument temporal response

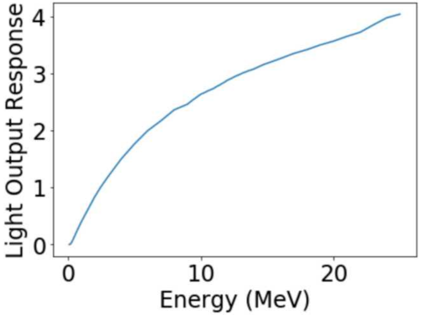

Light output energy response

Surrogate should be designed with an eye towards reduced dimensional representation

- For efficient Markov Chain Monte Carlo (method to sample from P(A|B))
  - Access to analytic gradients (*e.g.* neural networks)
  - Low dimensional
    - work in a compressed "latent space"
      - Critical idea: small number of features descriptive for the specified task

Signals

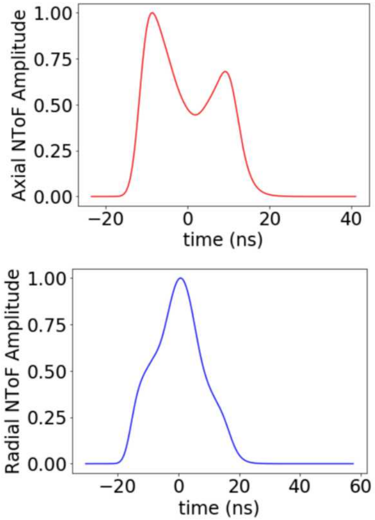

Small number of features*

Standard deviation (width of NToF)

$$\sigma = \sqrt{E[(t-\mu)^2]}$$

Skewness (asymmetry of NToF)

$$\varsigma = E\left[\left(\frac{t-\mu}{\sigma}\right)^3\right]$$

Kurtosis ("tailedness" of NToF)

$$\kappa = E\left[\left(\frac{t-\mu}{\sigma}\right)^4\right]$$

*I found shape descriptors to be reliably extracted from noisy data, but in principle many choices are possible, and I tried several different features.

Surrogate should be designed with an eye towards application to real data.

z3143 demonstrates the typical data quality

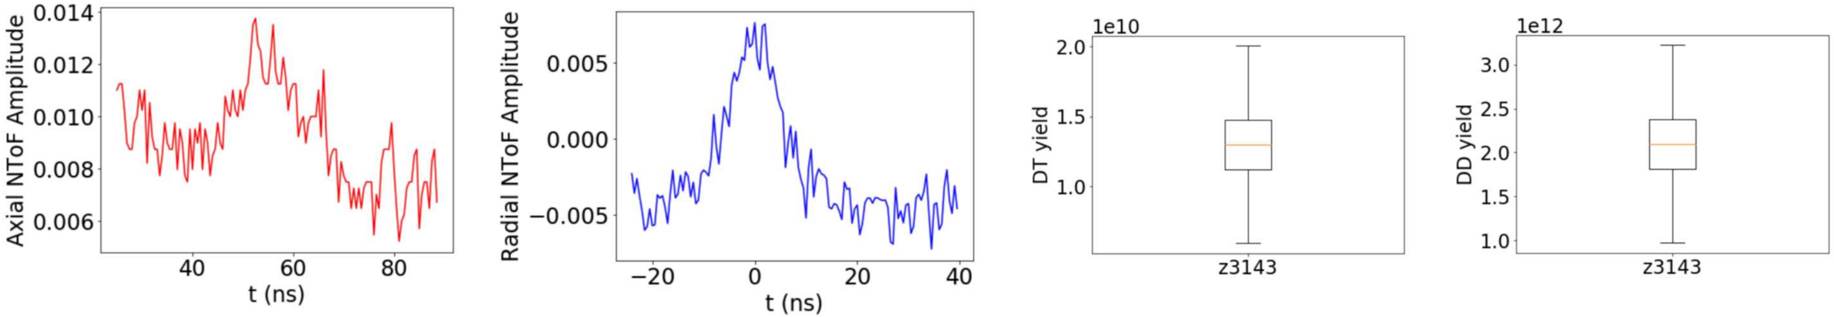

- NToF exhibit
  - Significant digitization noise
  - Approximately linear background in region-of-interest
- Wish to do feature extraction in a fashion
  - that is robust to noise and background (statistical moments have issues)
  - has a meaningful characterization of uncertainty
  - Quantile features of the NToF signal were found to be a good featurization

We can gain some additional intuition for constraining parameters from measurements.

$$\log \frac{DT}{DD}$$

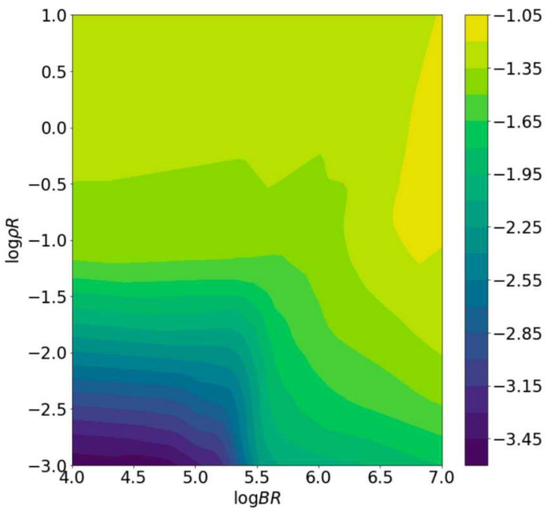

$Radial\ ToF\ width$

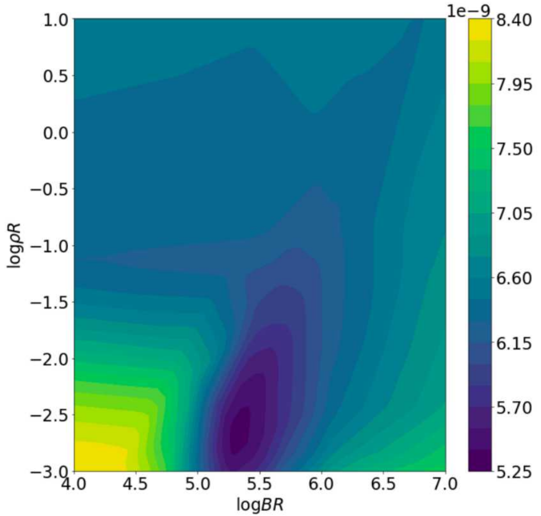

$Axial\ ToF\ width$

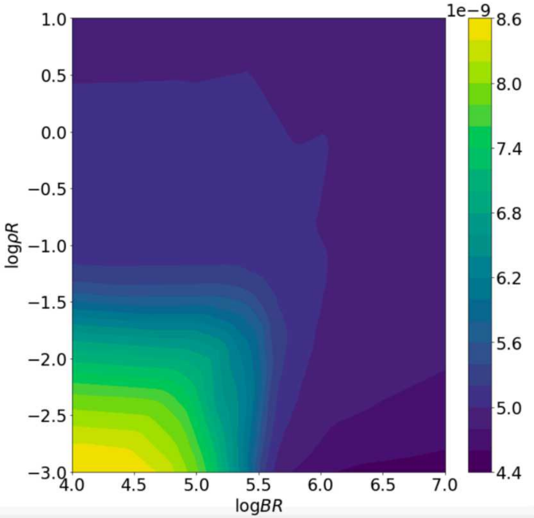

Sanity check calculations using posterior distribution show good agreement with data

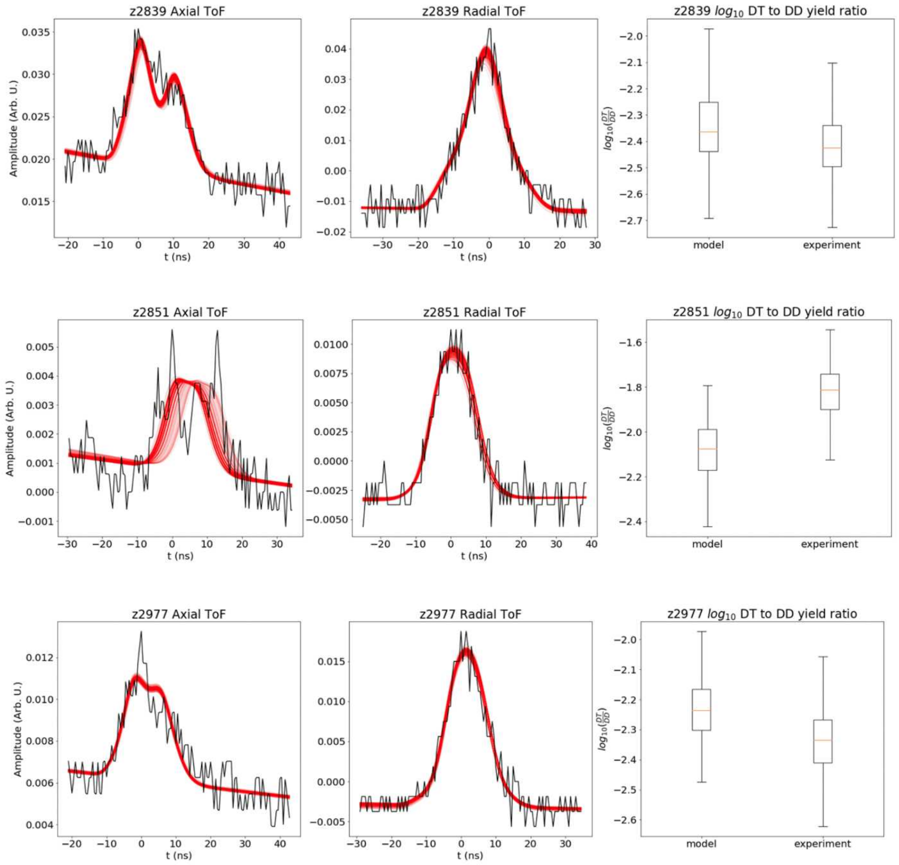

Sanity check calculations using posterior distribution show good agreement with data