



On Relating Uncertainties in Machine Learning and High Energy Physics

Michael Kagan
SLAC


PHYSTAT / CERN Data Science Seminar
November 16, 2022



We demand rigidly defined areas of doubt and uncertainty!



- Douglas Adams,
The Hitchhikers Guide to the Galaxy



I DON'T KNOW HOW TO PROPAGATE
ERROR CORRECTLY, SO I JUST PUT
ERROR BARS ON ALL MY ERROR BARS.

Snowmass Machine Learning Report

CompF3: Machine Learning

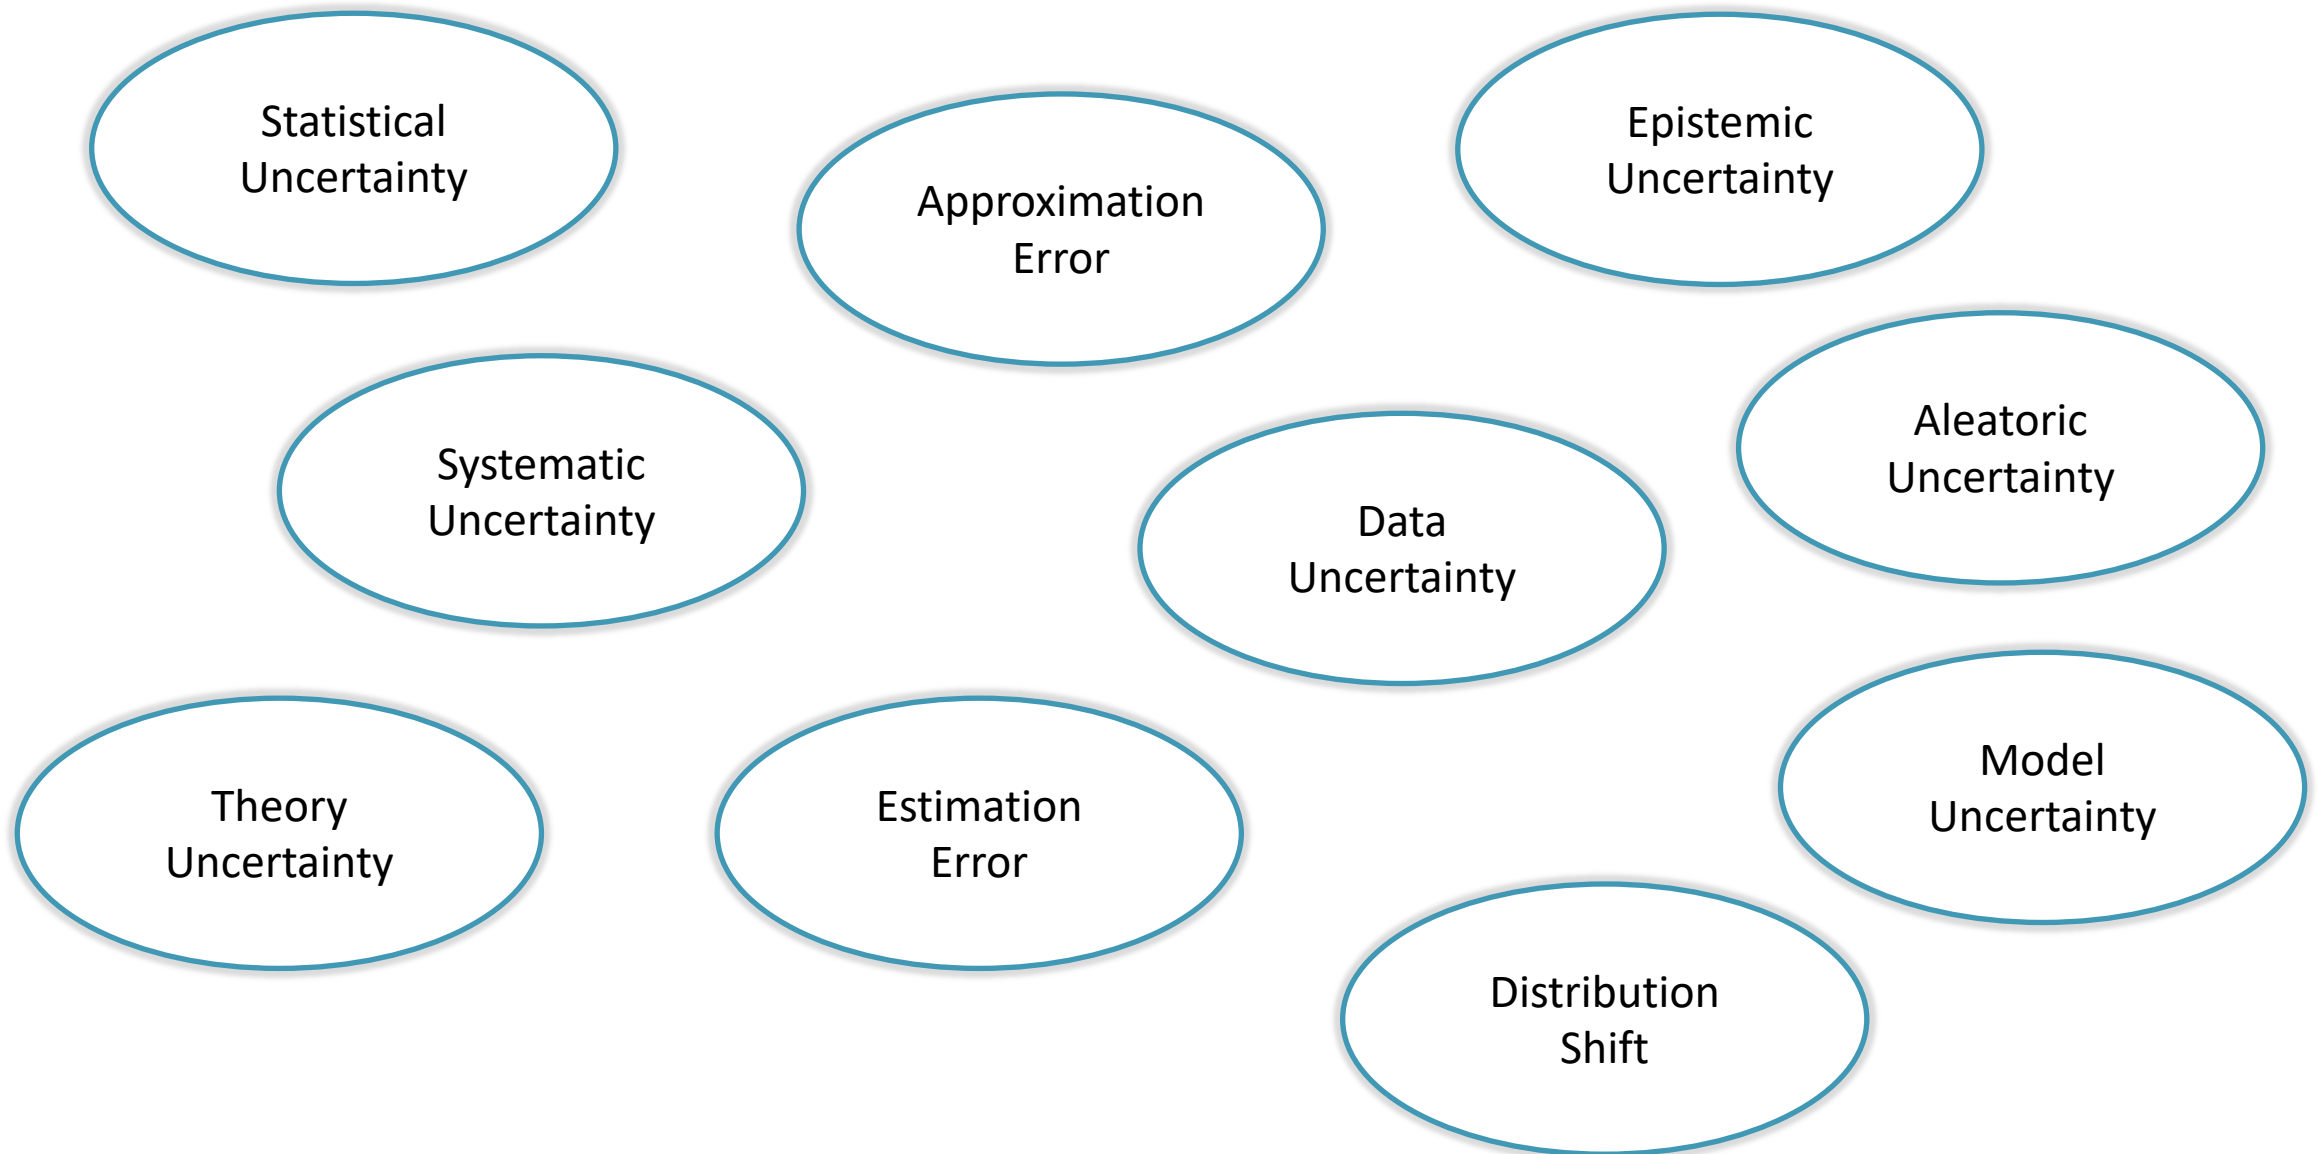
Phiala Shanahan, Kazuhiro Terao, Daniel Whiteson (Editors)

Including contributions from White Paper authors:

Gert Aarts^{1,2}, Andreas Adelmann³, N. Akhurin⁴, Andrei Alexandru^{5,6}, Oz Amram⁷, Anders Andreassen⁸, Artur Apresyan⁹, Camille Avestruz¹⁰, Rainer Bartoldus¹¹, Keith Bechtol¹², Kees Benkendorfer^{13,14}, Gabriele Benelli⁵⁹, Catrin Bernius¹¹, Alexander Bogatskiy¹⁵, Blaz Bortolato¹⁶, Denis Boyda^{17,18}, Gustaaf Brooijmans¹⁹, Paolo Calafiura¹³, Salvatore Cali^{20,18}, Florencia Canelli²¹, Grigorios Chachamis²², S.V. Chekanov¹⁷, Deming Chen²³, Thomas Y. Chen⁴⁰, Aleksandra Ciprijanovic⁹, Jack H. Collins¹¹, Andrew J. Connolly²⁴, Michael Coughlin²⁵, Biwei Dai²⁶, J. Damgov⁴, Gage DeZoort²⁷, Daniel Diaz²⁸, Barry M. Dillon^{16,29}, Ioan-Mihail Dinu⁷, Zhongtian Dong³⁰, Julien Donini³¹, Javier Duarte²⁸, S. Dugad³², Cora Dvorkin¹³, D. A. Faroughy²¹, Matthew Feickert²⁸, Yongbin Feng⁹, Michael Fenton⁵⁸, Sam Foreman¹⁷, Felipe F. De Freitas³⁴, Lena Funcke^{20,18,35}, P. G. C⁴, Abhijith Gandrakota⁹, Sanmay Ganguly³⁶, Lehman H. Garrison¹⁵, Spencer Gessner¹¹, Aishik Ghosh⁵⁸, Julia Gonsk¹⁹, Matthew Graham⁴⁸, Lindsey Gray⁹, S. Grönroos³⁷, Daniel C. Hackett^{20,18}, Philip Harris²⁰, Scott Hauck²⁴, Christian Herwig⁹, Burt Holzman⁹, Walter Hopkins¹⁷, Shih-Chieh Hsu²⁴, Jin Huang³⁸, Yi Huang³⁸, Xiao-Yong Jin¹⁷, Michael Kagan¹¹, Alan Kah¹⁹, Jernej F. Kamenik^{16,39}, Raghav Kansal²⁸, Georgia Karagiorgi⁴⁰, Gregor Kasieczka⁴¹, Erik Katsavounidis²⁰, Elham E. Khoda²⁴, Charanjit K. Khosa^{42,43}, Thomas Kipr⁴⁴, Patrick Komiske²⁰, Matthias Komm³⁷, Risi Kondor⁴⁵, Evangelos Kourlitis¹⁷, Claudius Krause⁴⁶, K. Lamichhane⁴, Luc Le Pottier^{13,10}, Meifeng Lin³⁸, Yin Lin^{20,18}, Mia Liu⁴⁷, Nan Lu⁴⁸, Biagio Lucini^{49,1}, J. Martinez⁴, Pablo Martín-Ramiro^{13,50}, Andrej Matevc^{16,39}, William Patrick McCormack²⁰, Eric Metodiev²⁰, Vinicius Mikuni²¹, David W. Miller⁴⁵, Siddharth Mishra-Sharma^{33,18,6}, Samadrita Mukherjee³², Daniel Murnane¹³, Benjamin Nachman^{13,51}, Gautham Narayan²³, Mark Neubauer²³, Jennifer Ngadiuba⁹, Scarlet Norberg⁶⁰, Brian Nord^{9,4}, Inês Ochoa⁵², Jan T. Offermann⁴⁵, Sang Eon Park²⁰, Kevin Pedro⁹, Cristian Peña⁹, Alexx Perloff⁶¹, Mariel Pettee¹³, Maurizio Pierini³⁷, T. Quast³⁷, Dylan Rankin²⁰, Yihui Ren³⁸, Marcel Rieger³⁷, Jean-Roch Vlimant⁴⁸, Avik Roy²³, Veronica Sanz^{42,53}, Nilai Sarda²⁰, Claire Savard⁶¹, Alexander Scheinker⁵⁴, Uroš Seljak^{13,51,26}, Brian Sheldon²⁸, David Shih⁴⁶, Chase Shimmin⁵⁵, Aleks Smolkovic¹⁶, George Stein^{13,26}, Cristina Mantilla Suarez⁹, Manuel Szwec⁵⁶, Savannah Thais²⁷, Jesse Thaler²⁰, Dmitrii Torbunov³⁸, Nhan Tran⁹, Steven Tsan²⁸, Silviu-Marian Udrescu²⁰, S. Undeeb⁴, Louis Vaslin³¹, Francisco Villaseca-Navarro^{15,27}, V. Ashley Villar⁵⁷, Brett Viren³⁸, Jean-Roch Vlimant⁴⁸, A. Whitbeck⁴, Daniel Williams¹⁹, Daniel Winkler²⁰, Si Xie⁴⁸, Tingjun Yang⁹, Haiwang Yu³⁸, and Mikael Yunus²⁰

Uncertainty Quantification, Validation and Interpretability It is vital that physicists can validate the decisions of ML models and quantify their uncertainty, a goal made easier if the inner workings of the models are conceptually accessible to physicists. HEP is not alone in this concern, and can benefit from work in the wider community. The HEP community should support continued research into interpretable AI and uncertainty quantification (UQ), including making public benchmark data sets for rigorous testing and comparison of approaches to physically interpretable AI-UQ for physics, and supporting challenges and competitions to create and compare methods of uncertainty quantification, including bias mitigation.

Many Terminologies Around Uncertainty



Machine Learning in the Data Analysis Pipeline

θ – Physics Parameters

Matrix Elements,
PDF's

Fragmentation /
Hadronization

Detector
Interactions

Data Generation

$\hat{\theta}$ – Parameter Estimates








Hypothesis Testing

Event Selection /
Reconstruction

Particle
Reconstruction

Trigger

Data Analysis



Primary Questions of Concern


How to deal with *Systematic Uncertainties* when using Machine Learning Models?

Is there uncertainty from using the ML Model?

ML “Model Uncertainty”:

What if the ML model did not “perfectly” fit the data?

When does it matter?



Supervised Learning Setup

Training Data:

- $\mathcal{D} = \{x_i, y_i\}$ = features and target
- $x, y \sim p(x, y)$


Goal:

- Learn $f_w(x) = \hat{y}$
- w = model weights

Learning:

- Optimize Loss

$$w^* = \arg \min_w L = \arg \min_w \frac{1}{N} \sum_i \mathcal{L}(y, f_w(x))$$



Optimal vs. Correct


Reconstruction, data selection, event classification enable us to define powerful summary statistics

$$T_{w^*,\phi}(x): \mathbb{R}^{10^8} \rightarrow \mathbb{R}$$

Estimate likelihood for frequentist parameter inference:

$$p(T_{w^*,\phi}(x) | \lambda(\theta))$$

- θ = physics parameters of interest
- w^* = Learned params, ϕ = Reco / analysis params
- $\lambda(\cdot)$ = parameters of probability density
e.g. mean of Poisson / Gaussian density



Optimal vs. Correct

Reconstruction, data selection, event classification enable us to define powerful summary statistics

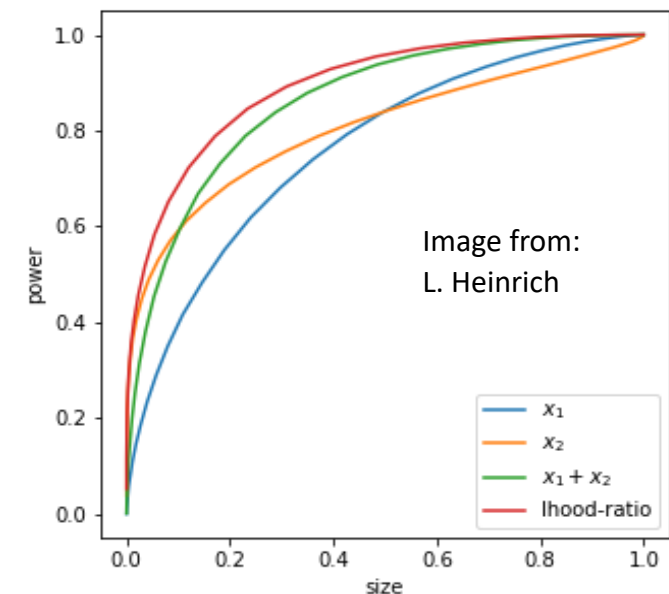

$$T_{w^*,\phi}(x): \mathbb{R}^{10^8} \rightarrow \mathbb{R}$$

Estimate likelihood for frequentist parameter inference:

$$p(T_{w^*,\phi}(x) | \lambda(\theta))$$

Changing summary statistic $T(x)$ affects optimality of result, but not correctness

- Reconstruction, event classification, ...
- **Not a question of ML model uncertainty**



Optimal vs. Correct

Reconstruction, data selection, event classification enable us to define powerful summary statistics


$$T_{w^*,\phi}(x): \mathbb{R}^{10^8} \rightarrow \mathbb{R}$$

Estimate likelihood for frequentist parameter inference:

$$p(T_{w^*,\phi}(x) | \lambda(\theta))$$

ML models that affect $\lambda(\cdot)$

- Background estimation, simulations, ...
- Affects compatibility of statistical model with data
- **Quality of ML model could lead to uncertainty, Or requires additional systematic uncertainties**




The Effect of Systematic Uncertainties

Systematic Uncertainties

- Simulation used for training $f_w(x)$
- Simulation not a perfect model of data
- $p_{SIM}(x, y) \neq p_{DATA}(x, y)$

Problem:

- Evaluating $f_w(x)$ will result in different distributions in simulation and data



Must consider how to handle systematic uncertainties for all ML models

How do Machine Learners think about uncertainty?

What kinds of uncertainty is relevant?

How do we estimate these uncertainties, when we need to?

How can we incorporate systematic uncertainties in HEP ML models?

This talk: An incomplete look at an ongoing research area

- [Uncertainties workshop](#) at Learning to Discover → this talk started there
- Great new ML review in PDG: [\[Cranmer, Seljak, Terao, 2021\]](#)
- Snowmass paper on uncertainty for ML in HEP: [\[2208:03284\]](#)
- Book Chapter: [\[Dorigo, de Castro Manzano\]](#)

Uncertainties in Machine Learning


Types of Uncertainties

Aleatoric Uncertainty:

Inherent variations in data, e.g. due to randomness of the process

Epistemic Uncertainty:

Due to lack of knowledge, lack of data, incomplete information




Types of Uncertainties

Aleatoric Uncertainty:

Inherent variations in data, e.g. due to randomness of the process

Epistemic Uncertainty:

Due to lack of knowledge, lack of data, incomplete information



Domain Shift:

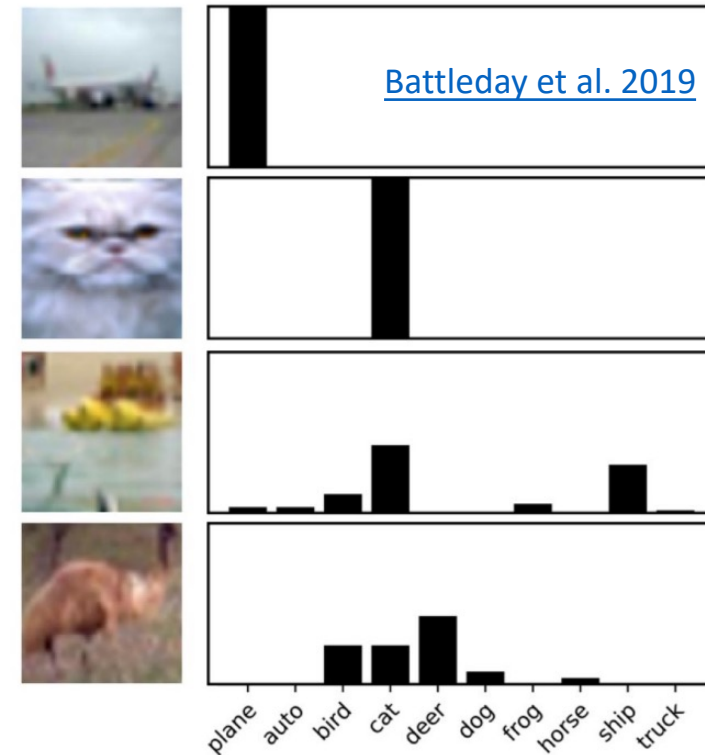

Test data is different from training data

Aleatoric Uncertainty

Often called “Statistical Uncertainty”

Variability in outcome of experiment due to inherently random effects

Often considered “irreducible”



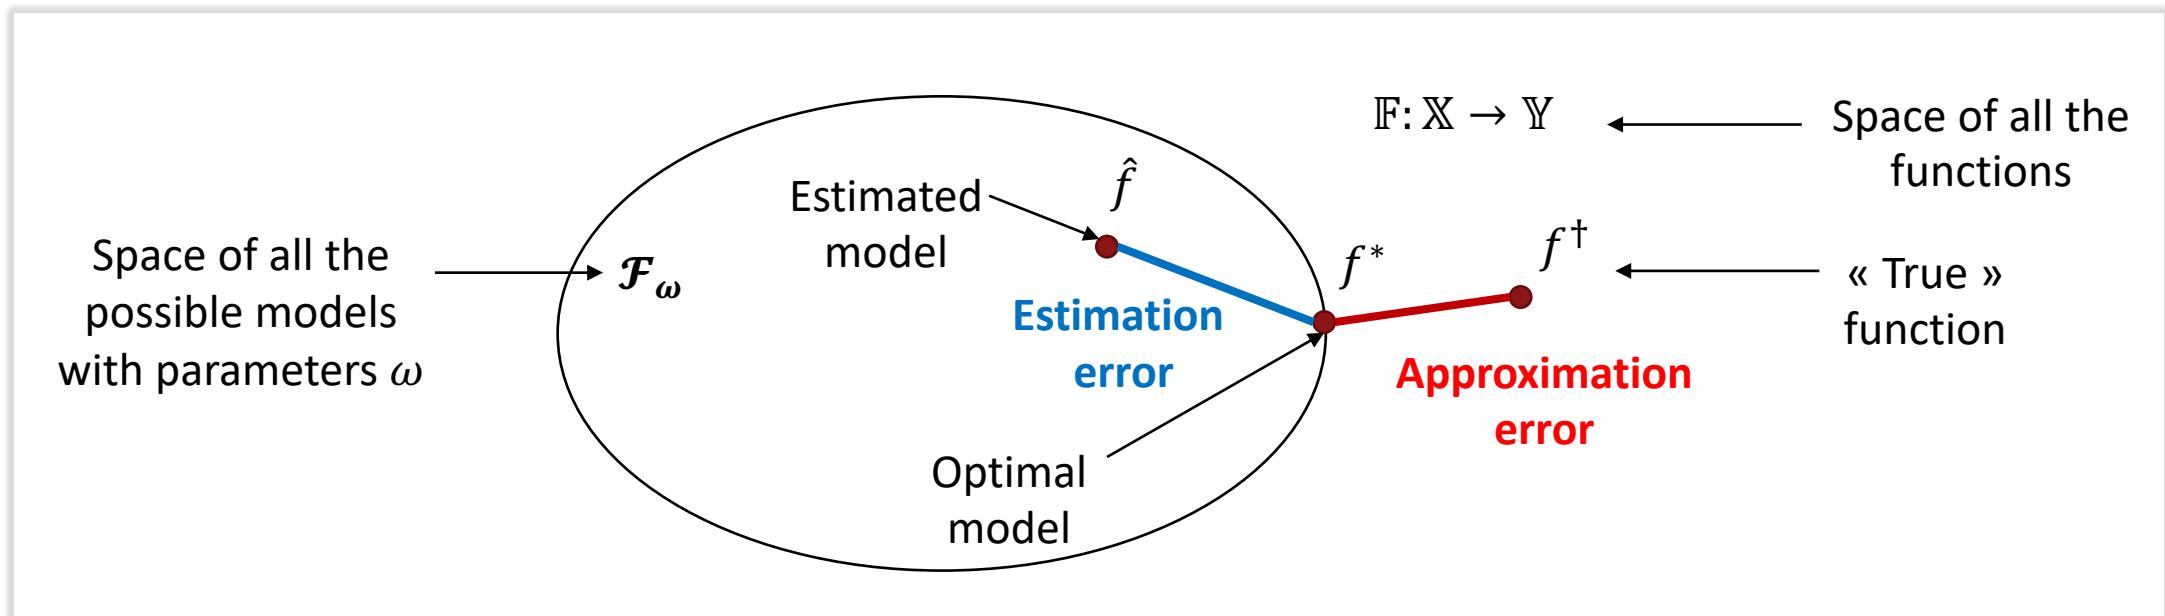
Epistemic Uncertainty

Lack of knowledge about the best model

Main origins in ML

- Estimation error: Training data just a sample of possible observations
- Approximation error: no model (in model class) can capture unknown true model

Often considered “reducible” with more data or more complex model

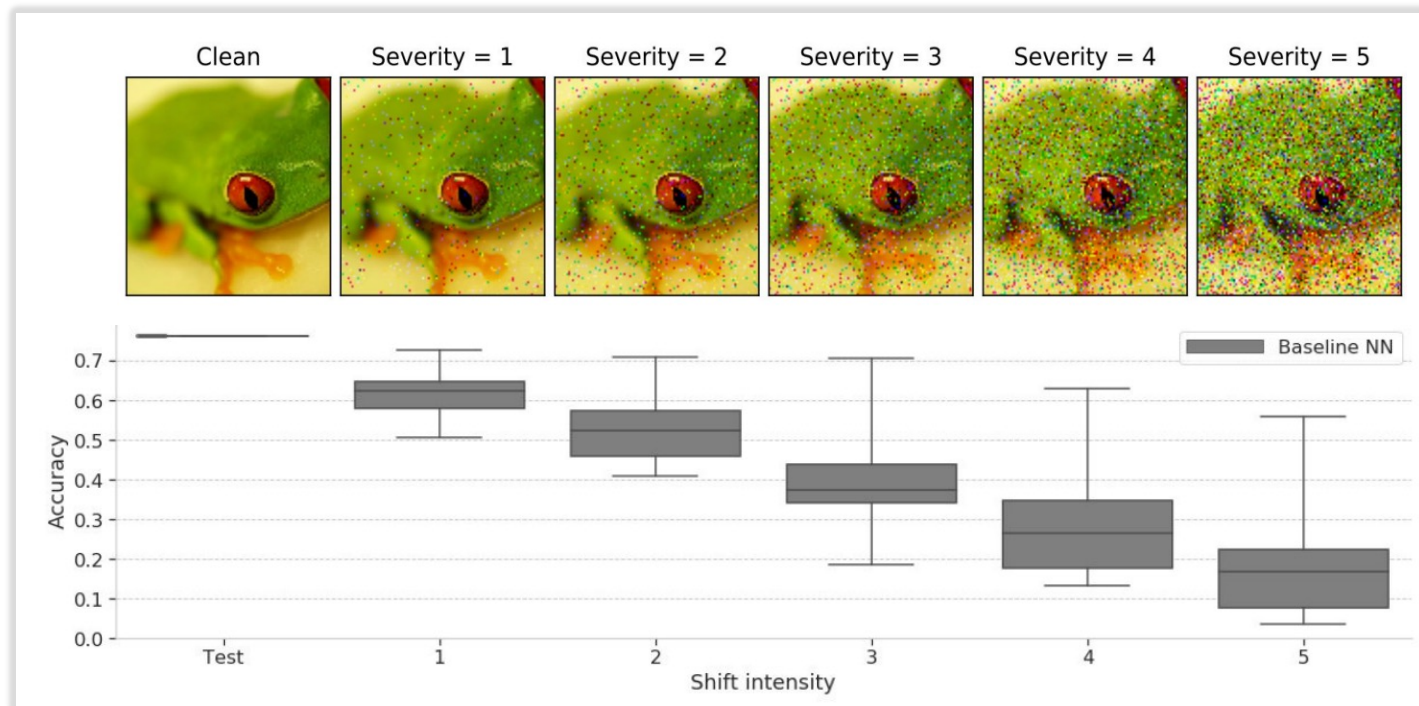


Domain / Distribution / Dataset Shift

$$p_{TEST}(x, y) \neq p_{TRAIN}(x, y)$$

Examples:

- Covariate Shift: $p(y|x)$ fixed but $p_{TEST}(x) \neq p_{TRAIN}(x)$
- Label Shift: $p(x|y)$ fixed but $p_{TEST}(y) \neq p_{TRAIN}(y)$
- Concept Shift: $p(y)$ fixed but $p_{TEST}(x|y) \neq p_{TRAIN}(x|y)$



Imperfect Correspondence: My View*

Machine Learning

Aleatoric uncertainty

- “Statistical” / “Data” Uncertainty
- Uncertainty Inherent to data
- Not reduced w/ more data

Epistemic uncertainty

- “Model” Uncertainty
- Uncertainty from Imperfect knowledge
- Reduces with more data


Domain Shift

- Imperfect model of data generation process

HEP

Detector Noise
Resolutions

Stat. errors in HEP

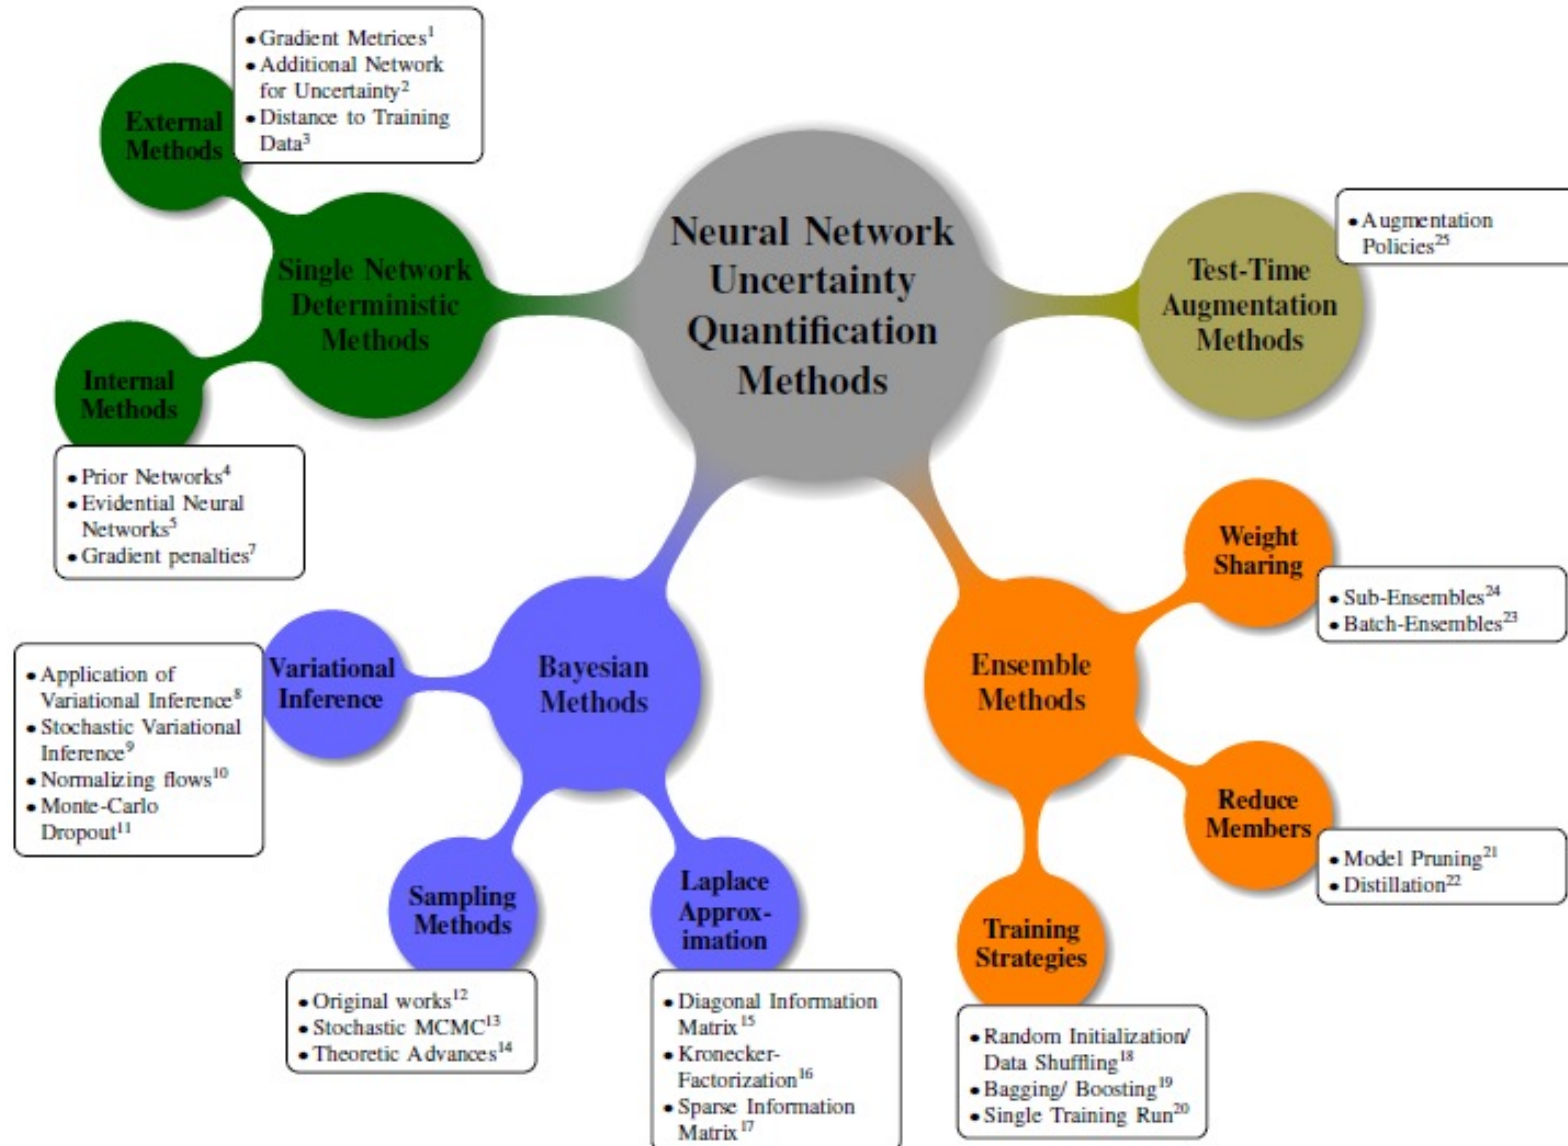


Systematic errors induced by ML model training on finite stats.

Systematic Uncertainties from data / simulation differences

*Even within the ML community, these terms can be ambiguous

Uncertainty Estimation Approaches in Deep Learning



Aleatoric Uncertainty


Randomness of data → Typically described by probability distributions

Density Networks

Define density $p_\phi(y|x)$
with parameters ϕ

Train neural network
to predict per-example
parameters

$$f(x) \rightarrow \phi(x)$$



Mixture density network

Randomness of data → Typically described by probability distributions

Generative Models:

Aim to approximate a density, $p(x)$

Train NN to transform noise $z \sim p(z)$ into data:

$$\hat{x} = f_w(z), \quad p(\hat{x}) \approx p_{data}(x)$$

Implicit models:

can only generate sample synthetic data, e.g. GANS

Explicit models:

can also evaluate density, e.g. Normalizing Flows

StyleGAN v2


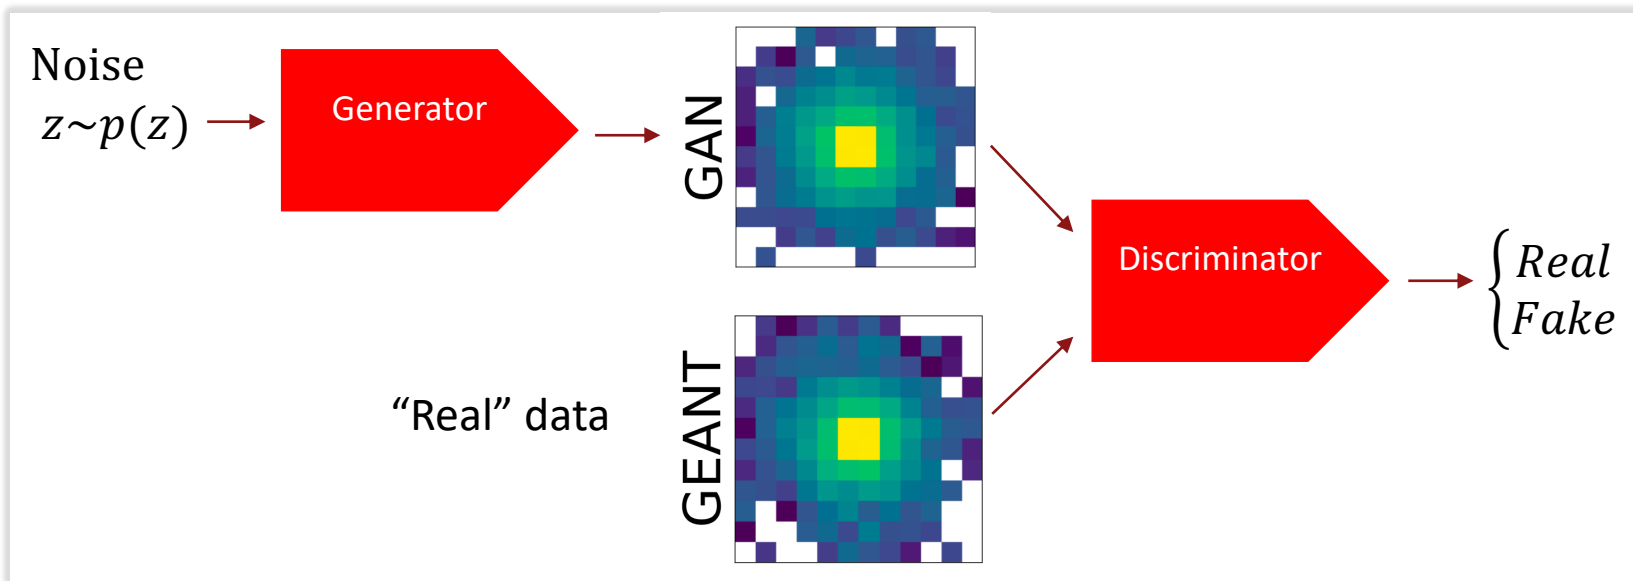


(Karras et al, 2019)

Aleatoric Uncertainty in HEP with Generative Models

Simulators slow / hard to sample from → approximate with Generative Model

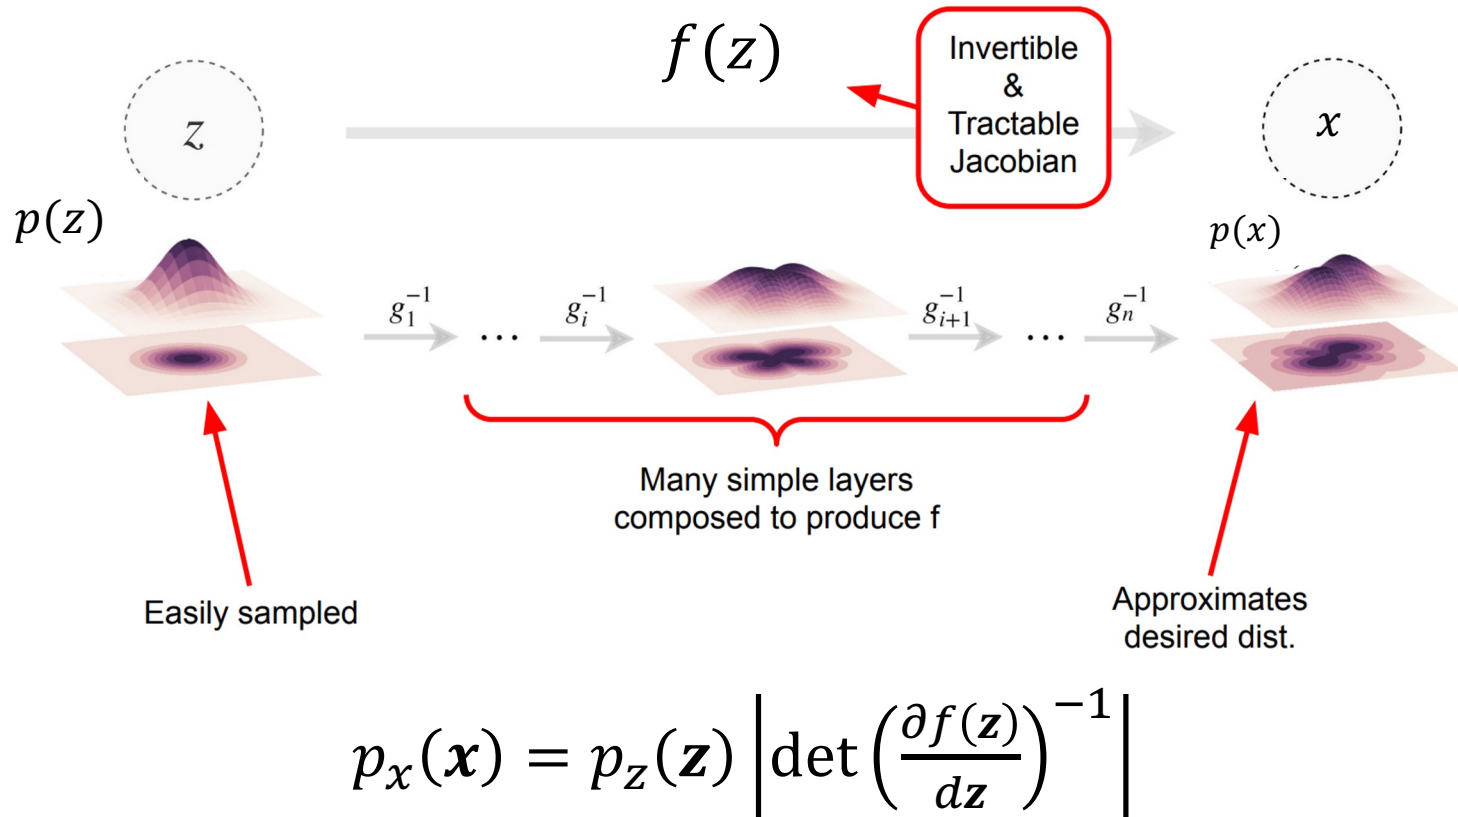
Generative Adversarial Networks:



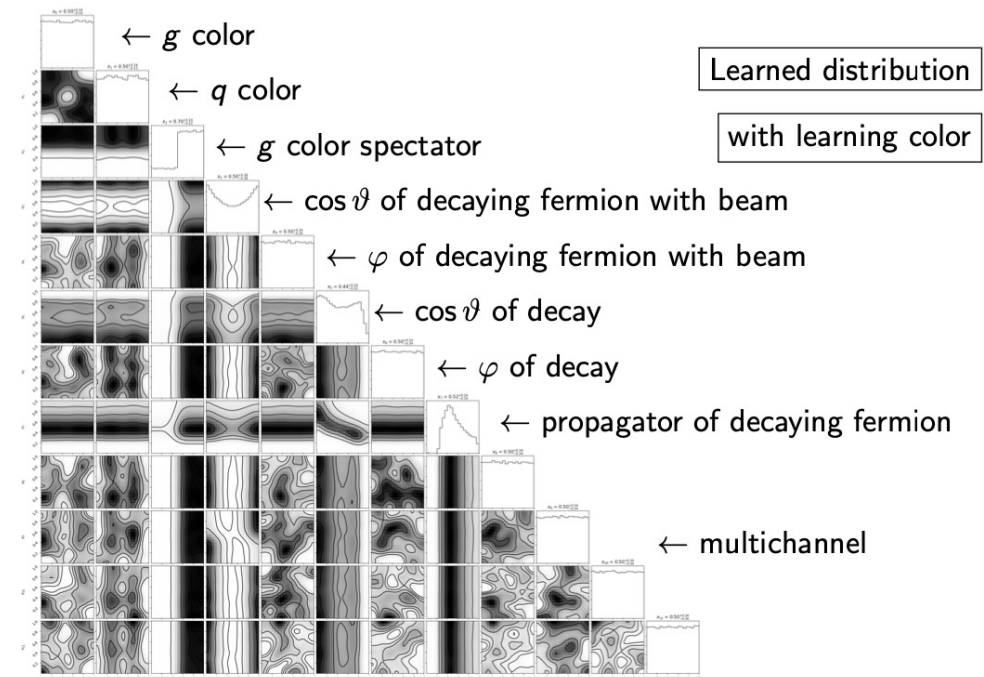
Aleatoric Uncertainty in HEP with Generative Models

Simulators slow / hard to sample from → approximate with Generative Model

Normalizing Flows



Example: Learning $e^+ e^- \rightarrow 3j$ Matrix Elements



Aleatoric Uncertainty in HEP: Neural Unfolding

$$p_{reco}(y) = \int p(y|x)p_{true}(x)$$

Continuous Form



$$y = Rx$$

Discrete Form

Response Matrix in unfolding \rightarrow Aleatoric Uncertainty

Several recent methods using ML to model the response and enable high-dimensional continuous unfolding

- E.g. [2011.05836](#), [2006.06685](#), [1911.09107](#)



What if ML learns the wrong generative model or response?

→ Understanding ML Model / Epistemic Uncertainties

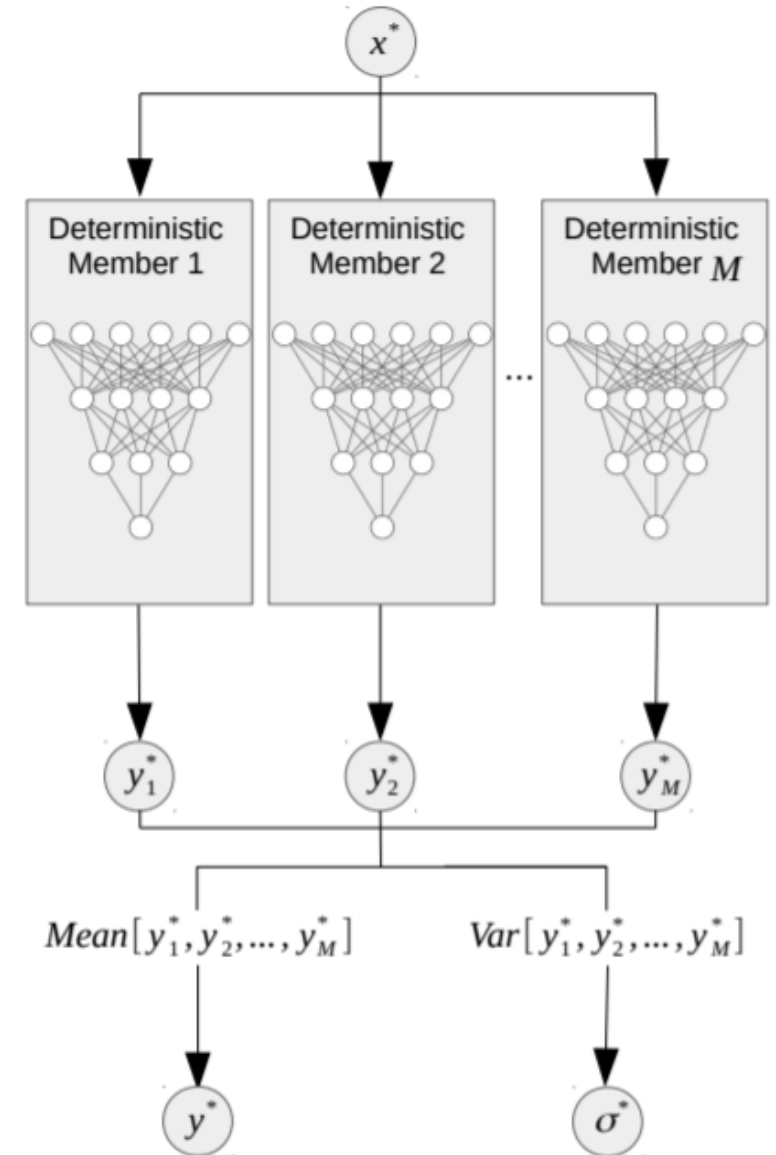
Epistemic Uncertainty with Deep Ensembles

Ensembling:

- Retrain network from multiple initializations

Can be coupled with Bootstrapping


- Randomly sample data, with replacement, to define each model's training set



Model Uncertainty in ML-based Background Estimation




High-Dimensional “ABCD” method with NN’s


- Learn reweighting using classifiers: $w(x) \approx \frac{p_A(x)}{p_B(x)}$
- Estimate background: $\hat{p}_D(x) = w(x)p_C(x)$



What if we didn't learn accurate weights?

- ATLAS $hh \rightarrow 4b$ example: Uncertainties from Deep ensembles & data bootstrap






$$p(y|x, \mathcal{D}) = \int p(y|x, w)p(w|\mathcal{D})dw \approx \frac{1}{N} \sum_{i=1 \dots N} p(y|x, w_i)$$

$w_i \sim p(w|\mathcal{D})$

Aleatoric Uncertainty:
Density Model

Model Uncertainty:
Posterior on weights




$$p(y|x, \mathcal{D}) = \int p(y|x, w)p(w|\mathcal{D})dw \approx \frac{1}{N} \sum_{\substack{i=1 \dots N \\ w_i \sim p(w|\mathcal{D})}} p(y|x, w_i)$$

$$p(w|\mathcal{D}) = \frac{p(\mathcal{D}|w)p(w)}{\int p(\mathcal{D}|w)p(w)dw}$$

Prior on weights ←
Intractable Integral ←


Approximating the Posterior

$p(w|\mathcal{D})$ is multi-modal and complex in NN \rightarrow approximation methods



Local approximations

- Locally, covering one mode well e.g. with a simpler distribution $q(w; \lambda)$
 - Variational inference
 - Laplace approximation







Sampling

- Summarize using samples
 - MCMC
 - Hamiltonian Monte Carlo
 - Stochastic Gradient Langevin Dynamics


Model Uncertainty on ML models for Event Generators

“Bayesian Normalizing Flow”


- Density Model: Normalizing Flow
- Model Uncertainty: Variational Posterior over weights



Monte Carlo Dropout



(a) Standard Neural Net



(b) After applying dropout.


$$f(x) \rightarrow \begin{cases} \text{Mean}[f^1 \dots f^N] \\ \text{Var}[f^1 \dots f^N] \end{cases}$$

Different random Dropouts


Randomly drop connections between neurons, using Bernoulli distribution

Can be viewed as a Variational Approximation

Comparisons




Comparisons with Data Corruptions

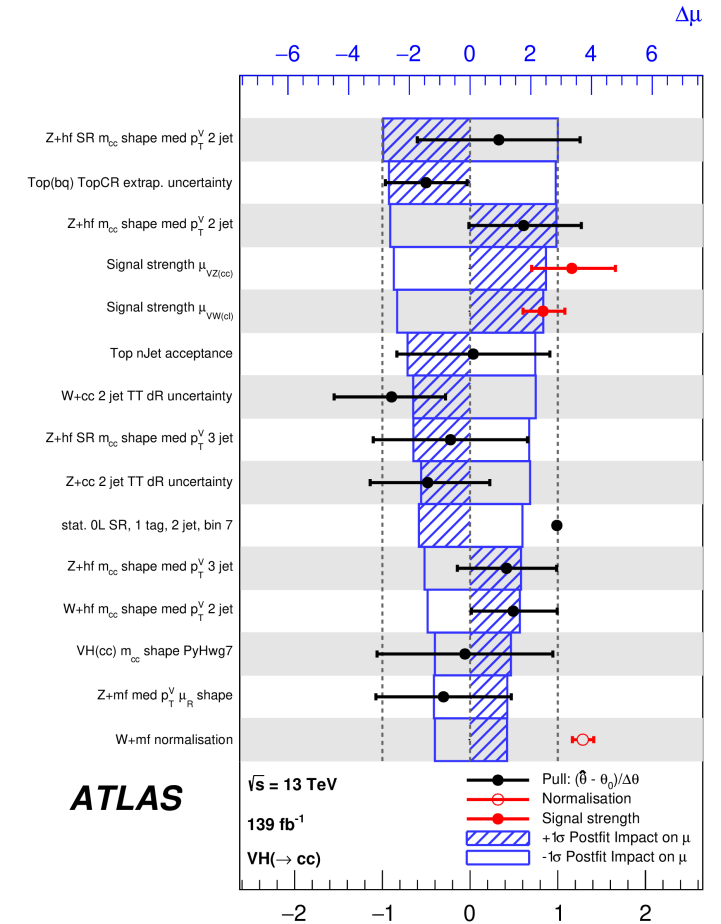


Systematic Uncertainties / Domain Shift in HEP

Systematic Uncertainties



Source of uncertainty	$\mu_{VH(cc)}$
Total	21.5
Statistical	16.2
Systematics	14.0
Statistical uncertainties	
Data statistics only	13.0
Floating normalisations	7.2
Theoretical and modelling uncertainties	
VH($\rightarrow c\bar{c}$)	2.1
Z+jets	7.7
Top-quark	5.6
W+jets	3.4
Diboson	0.8
VH($\rightarrow b\bar{b}$)	0.8
Multi-Jet	1.0
Simulation statistics	
Jets	3.7
Leptons	0.4
E_T^{miss}	0.5
Pile-up and luminosity	0.4
Experimental uncertainties	
Flavour tagging	
c-jets	2.3
b-jets	1.2
light-jets	0.7
τ -jets	0.4
Truth-flavour tagging	
ΔR correction	3.0
Residual non-closure	1.4



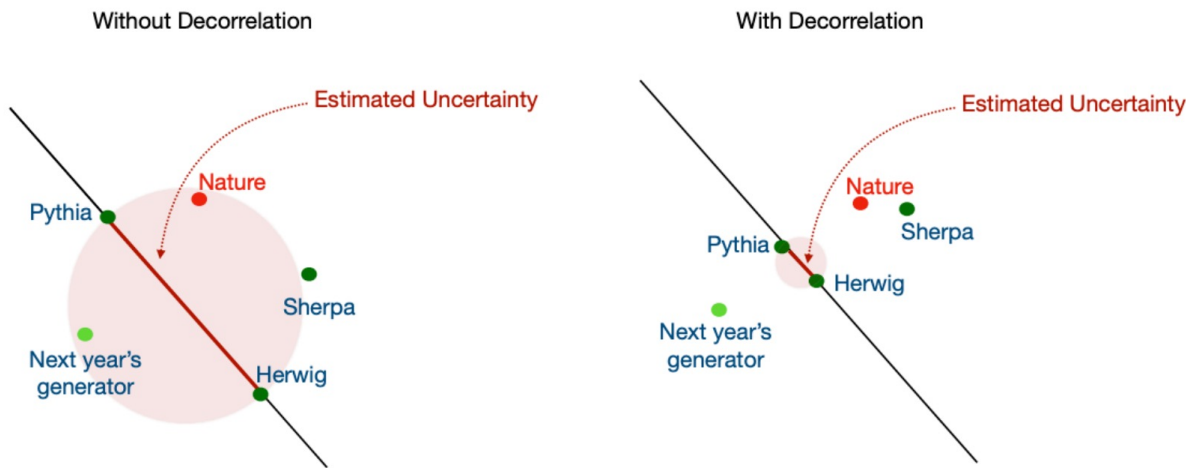
Theory uncertainties? ... Not going to discuss here

See nice recent [PHYSTAT talk](#) from D. Whiteson

See nice recent paper: Ghosh, Nachman, [2109.08159](#)

A Cautionary Tale of Decorrelating Theory Uncertainties

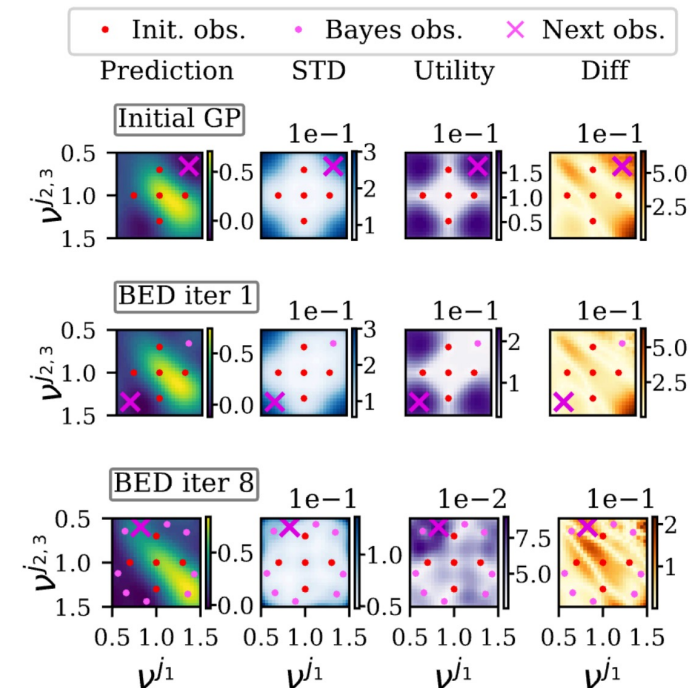
Aishik Ghosh^{a,b} and Benjamin Nachman^{b,c}



Efficient Estimation of Multiple Systematic Uncertainties with Gaussian Processes and Bayesian Experimental Design


Alexis Romero,¹ Kyle Cranmer,² and Daniel Whiteson¹

Work In Progress



Systematic Uncertainty = Domain / Distribution Shift

Unlike ML, we measure / parameterize possible variations over domains



Often can constrain from auxiliary measurements: $p(x_{aux}|v)$
(i.e. from calibrations for reconstructed objects)


- **propagation of errors:** one works with a model $f(x)$ and simply characterizes how uncertainty in the data distribution propagate through the function to the down-stream task irrespective of how it was trained.
- **domain adaptation:** one incorporates knowledge of the distribution for domains (or the parameterized family of distributions $p(x|y, \nu)$) into the training procedure so that the performance of $f(x)$ for the down-stream task is robust or insensitive to the uncertainty in ν .
- **parameterized models:** instead of learning a single function of the data $f(x)$, one learns a family of functions $f(x; \nu)$ that is explicitly parameterized in terms of nuisance parameters and then accounts for the dependence on the nuisance parameters in the down-stream task.
- **data augmentation:** one trains a model $f(x)$ in the usual way using training dataset from multiple domains by sampling from some distribution over ν .

Error Propagation – Standard Approach

42

Train on $\{x_i^0, y_i^0\}$ w/ nominal nuisance $\nu_0 \rightarrow$ learn fixed model $f(x)$

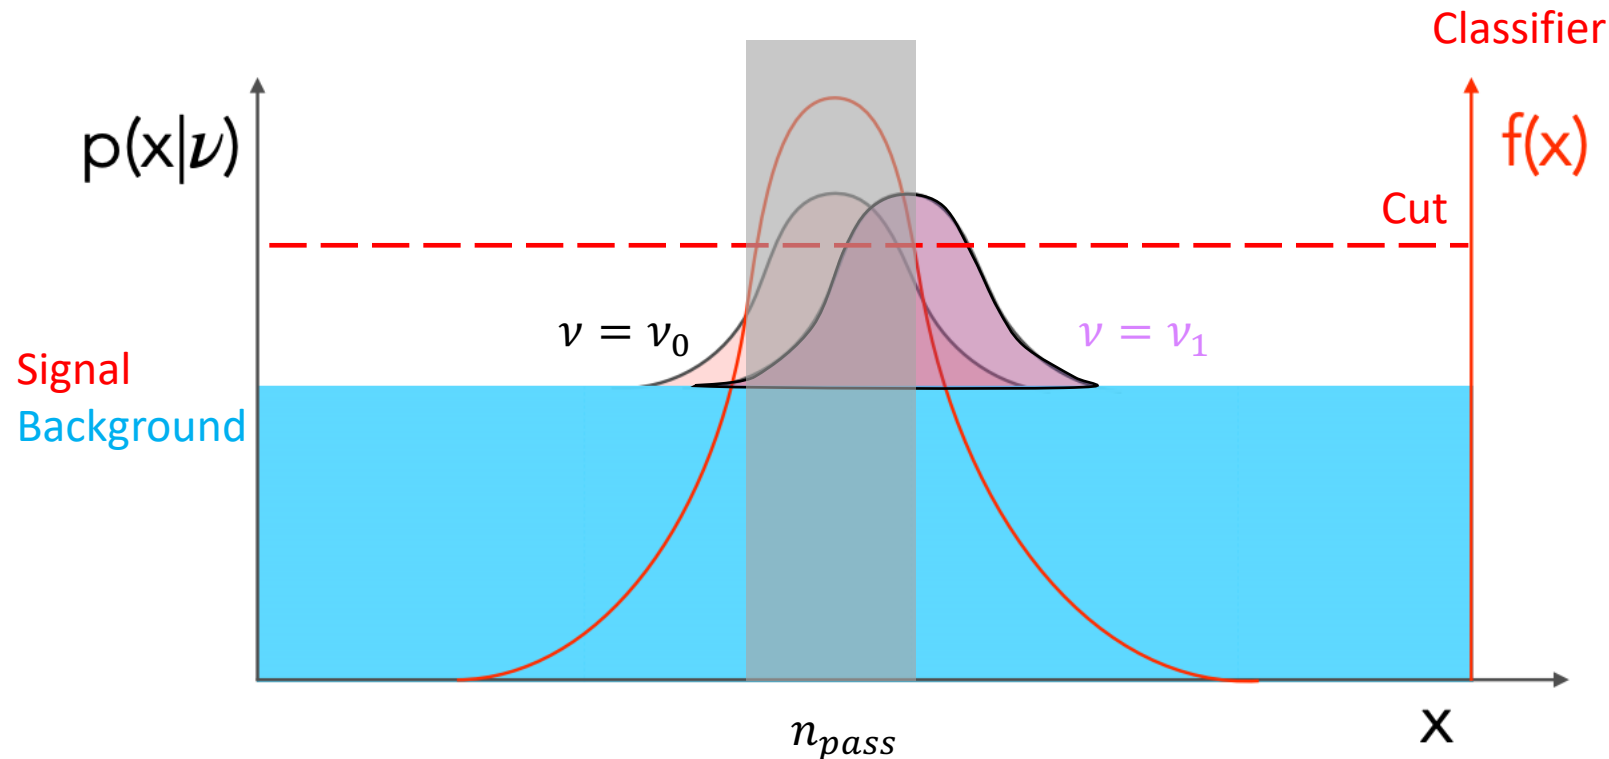
Evaluate on ν variations \rightarrow Observe effects



Error Propagation – Standard Approach

Train on $\{x_i^0, y_i^0\}$ w/ nominal nuisance $\nu_0 \rightarrow$ learn fixed model $f(x)$

Evaluate on ν variations \rightarrow Observe effects




NOT Optimal
Under Variations

Efficiency of Cut

$$\int_{\{x:f(x)>c\}} p(x|\nu)dx$$



$\nu = \nu_0$




$\nu = \nu_1$

Data Augmentation / Marginalization

Training sample includes ν variations: $x \sim \int p(x|y, \nu)p(\nu)d\nu$



Smearred samples \rightarrow “smearred” fixed model $f_{smearred}(x)$



Data Augmentation / Marginalization

Training sample includes ν variations: $x \sim \int p(x|y, \nu) p(\nu) d\nu$

Smearred samples \rightarrow “smearred” fixed model $f_{smearred}(x)$





Related Example:
CMS Boosted Jet Tagging
w/ ParticleNet Graph NN

Training on flat mass
distribution

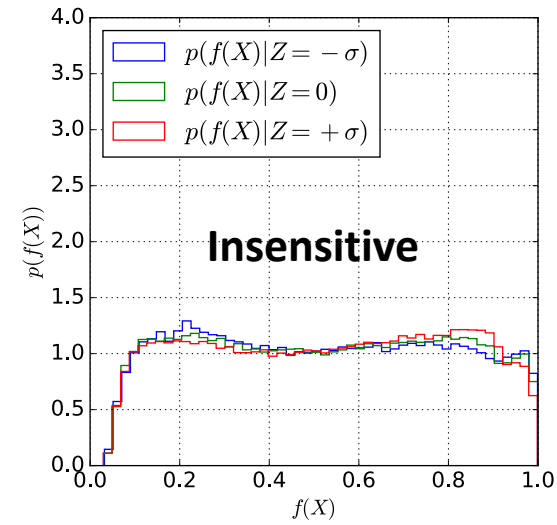

Pivoting / Enforcing Domain Invariance

Want to train model $f(x)$ such that: $p(f|v) = p(f)$ ← f is a *pivotal quantity*

Normal
Training

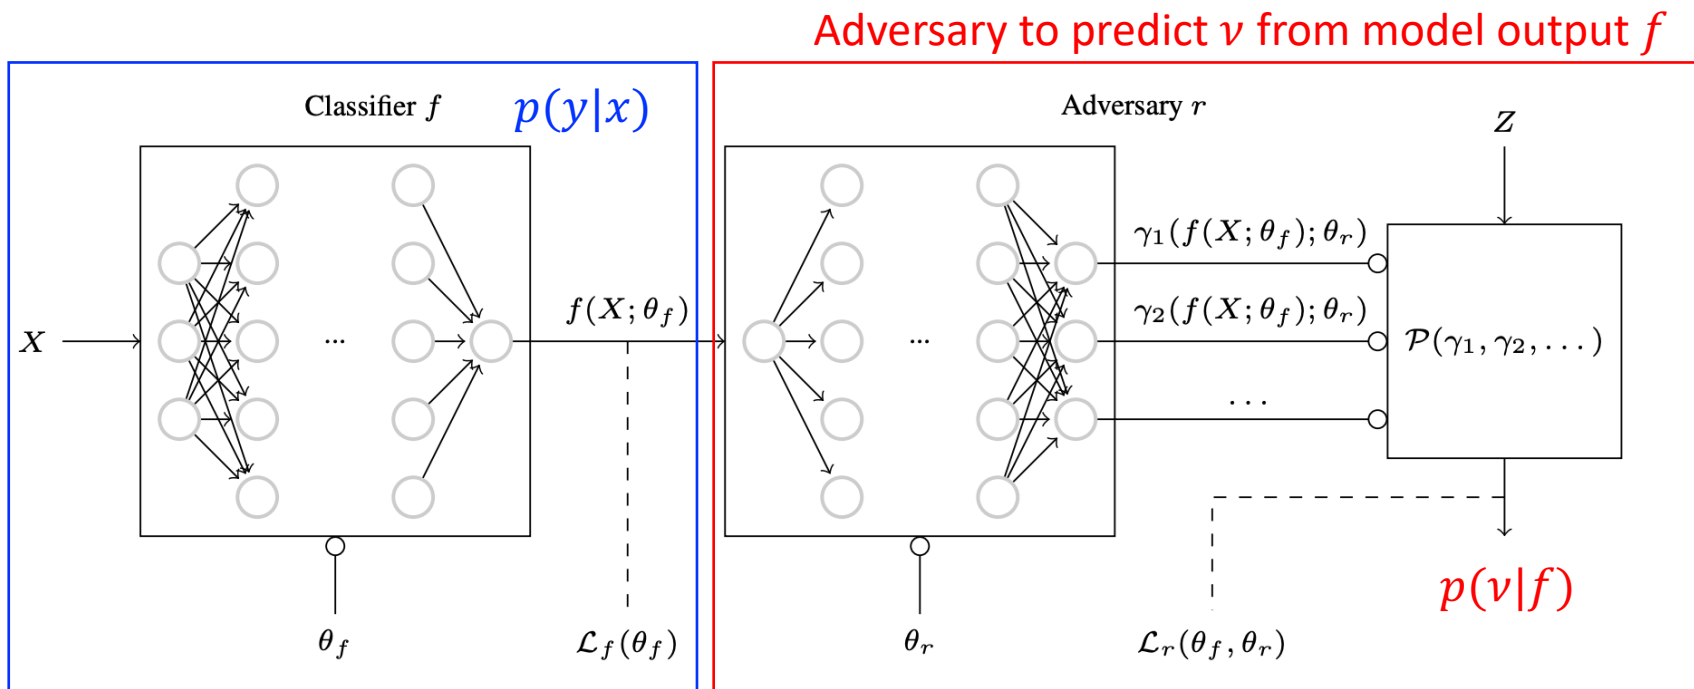


Pivot



Pivoting / Enforcing Domain Invariance


Adversarial Approach:



Min-Max Game: Penalize Classifier when Adversary succeeds

$$\hat{\theta}_f, \hat{\theta}_r = \arg \min_{\theta_f} \max_{\theta_r} E(\theta_f, \theta_r).$$


$$E_\lambda(\theta_f, \theta_r) = \mathcal{L}_f(\theta_f) - \lambda \mathcal{L}_r(\theta_f, \theta_r)$$



“Regularize” training with Adversary

Pivoting / Enforcing Domain Invariance

Adversarial Approach:




Pivoting / Enforcing Domain Invariance

Regularizing Correlations: Non-adversarial approach

Example: *Disco Fever: Robust Networks Through Distance Correlation*



$$L = L_{\text{classifier}}(\vec{y}, \vec{y}_{\text{true}}) + \lambda \text{dCorr}_{y_{\text{true}}=0}^2(\vec{m}, \vec{y})$$

$$\begin{aligned} \text{dCov}^2(X, Y) = & \langle |X - X'| |Y - Y'| \rangle \\ & + \langle |X - X'| \rangle \langle |Y - Y'| \rangle \\ & - 2 \langle |X - X'| |Y - Y''| \rangle \end{aligned}$$





Parameterizing Models

Train with nuisance parameters as input $\{x_i, y_i, v_i\} \rightarrow$ learn model $f(x; v)$



Parameterizing Models

Train with nuisance parameters as input $\{x_i, y_i, v_i\} \rightarrow$ learn model $f(x; v)$





Parameterizing Models



No Systematic Variation

With Systematic Variation

Fixed Model



Parameterized Model



Comparing Approaches


Example:

- Classifier: $h \rightarrow \tau\tau$ vs Bkg
- Uncertainty: τ energy scale


Parameterized Classifier:

$$f(x; \nu)$$

How to choose the ν ?
→ Profile in Likelihood

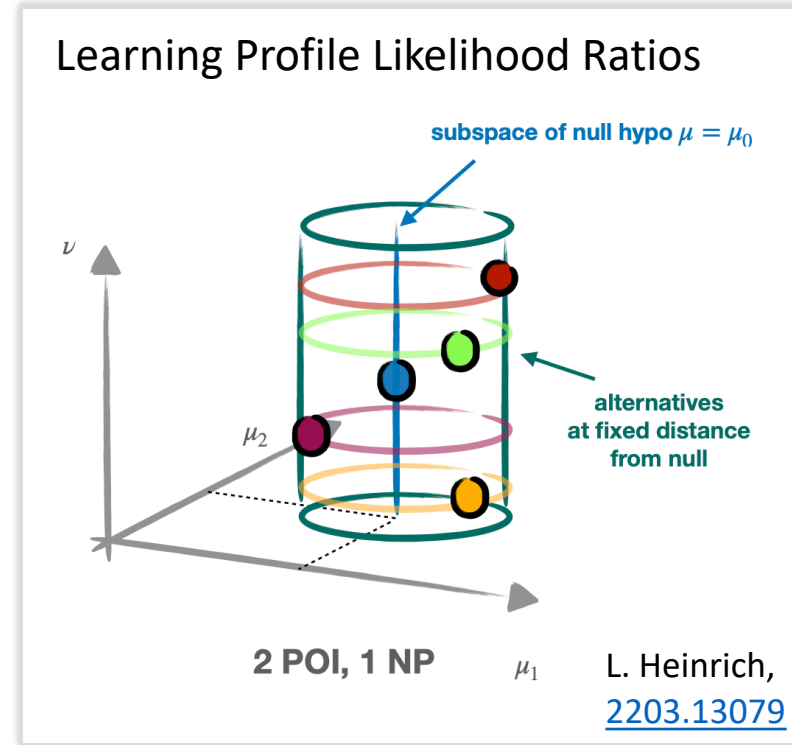
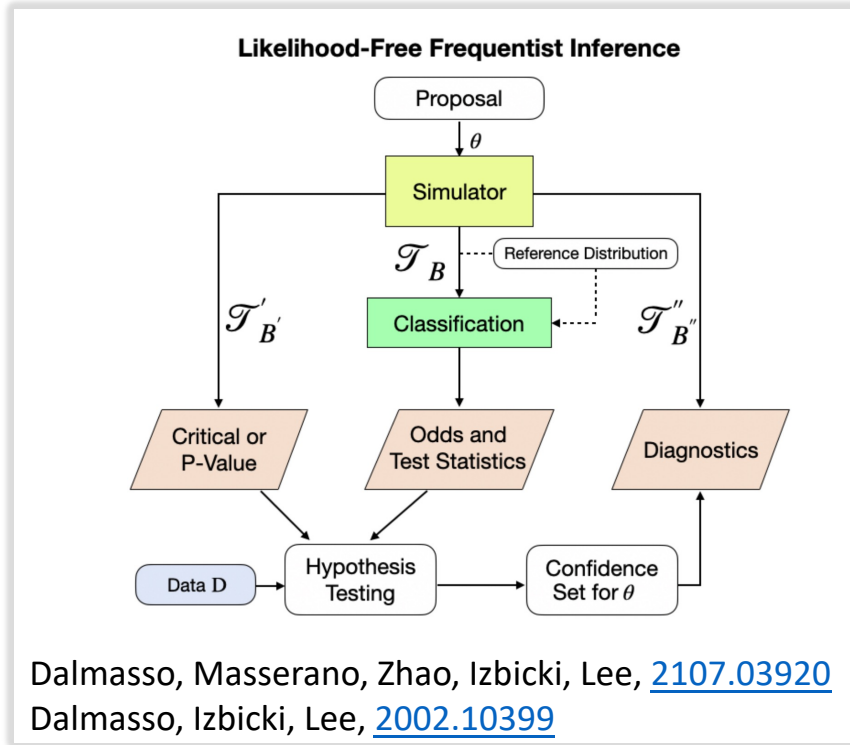
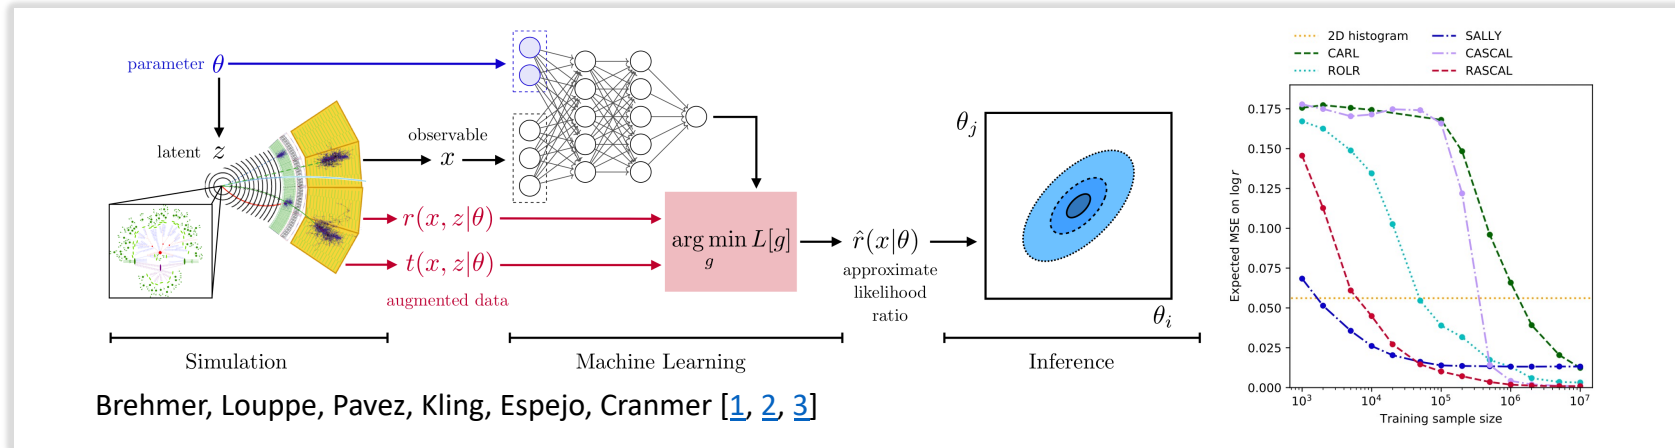


Data with Nominal Nuisance value



Data with Varied Nuisance value

Simulation-Based Inference: Estimating Likelihood Ratios with parameterized Models



Uncertainty when using ML in HEP → How and Where?

- Lots of ML research on estimating Data uncertainty & Model Uncertainty
- Must examine each application & how well calibrated the methods are?

Many areas where Model Uncertainty may be important (not all discussed today)

- ML-based Simulation and Background estimation
- Fast ML in the Trigger – Uncertainty in real-time decision making
- Simulation-based inference – estimating likelihood ratio directly with ML
- Anomaly Detection
- ...

Systematics will always remain a challenge, and understanding how to deal with them in ML models has made progress on several fronts

Backup

Standard HEP Inference

Reconstruction, data selection, event classification enable us to define powerful summary statistics


$$T(x): \mathbb{R}^{10^8} \rightarrow \mathbb{R}$$

Histogram for density estimation, with bin counts:

$$\{t_i\}_{i=1 \dots n_{bins}}$$

Binned Likelihood: $p(t_i | \theta, \nu) = \text{Pois}(t_i | \mu(\theta, \nu))$



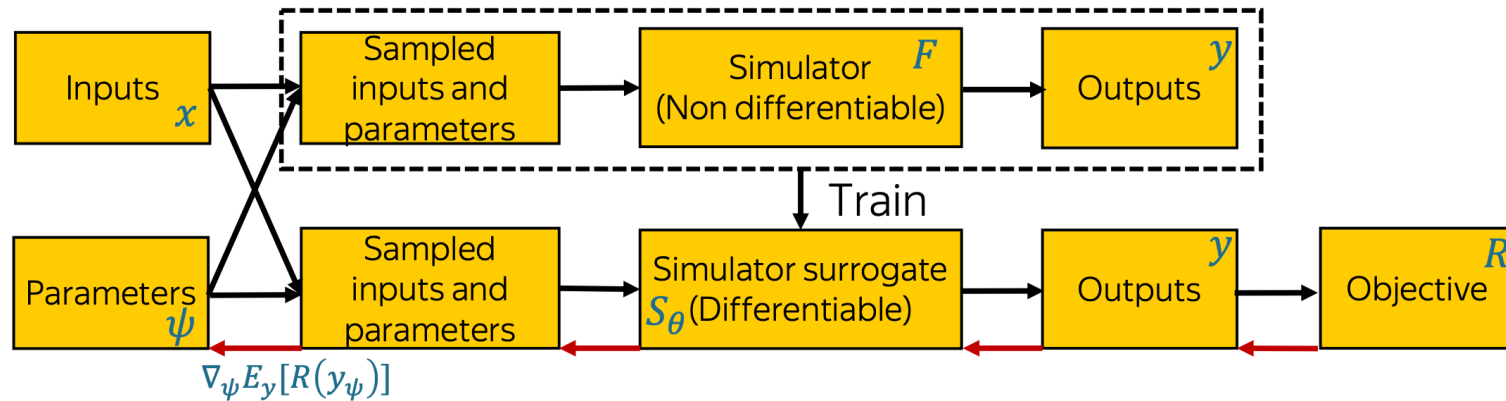
Test Statistic: $\lambda(\theta) = \log \frac{\prod_i p(t_i | \theta, \hat{\nu})}{\prod_i p(t_i | \hat{\theta}, \hat{\nu})}$



$$p(T(x) | \theta)$$



Aleatoric Uncertainty in HEP with Generative Models

Optimizing detector design with Generative Model base Surrogate Simulator



Example: SHiP Magnet Optimization
Reduced length and weight over previous design!

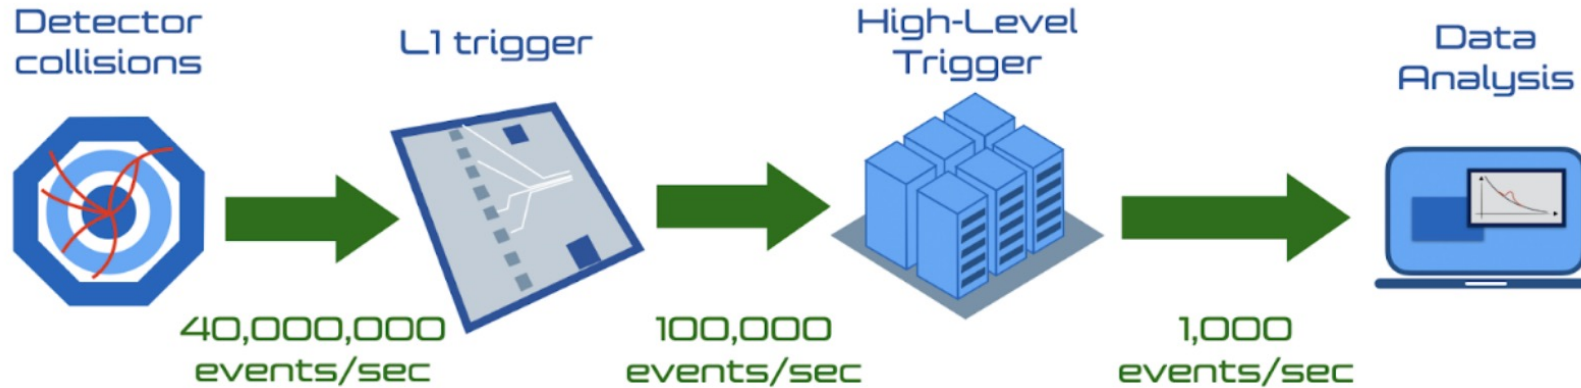
Bayesian Neural Networks for Jet Energy Estimation



Gaussian Variational Posterior over weights
Gaussian Density Network for p_T predictions

Uncertainties for ML in Trigger Systems

60



Decision Theory / Risk Management Problems

- Decisions are irrevocable and constrained by total rate

How certain we are about an ML prediction could change our decision!

Consideration for ML model uncertainties is important here

What if the generative model doesn't perfectly fit data?

Potentially bad description of data! → Case for Epistemic / Model Uncertainty

“Bayesian Normalizing Flow”
with Variational Inference

$$\mathcal{L} = \sum_{n=1}^N \langle \log p_X(x_n | \theta) \rangle_{\theta \sim q_\phi(\theta)} - \text{KL}(q_\phi(\theta), p(\theta))$$

Likelihood:
Normalizing Flow

Weights sampled
from posterior

Variational Posterior:
Gaussian

Simulation Based Inference (SBI)

Start with

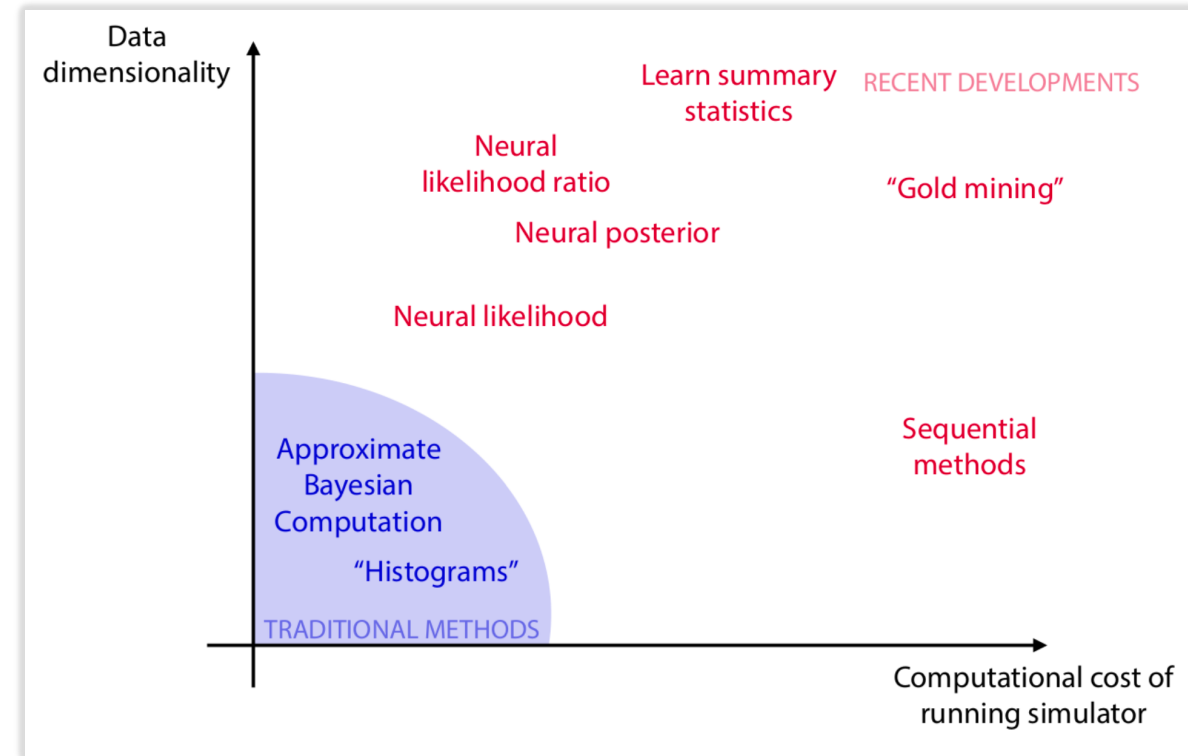

- a simulator that can generate N samples $x_i \sim p(x_i | \theta_i)$,
- a prior model $p(\theta)$,
- observed data $x_{\text{obs}} \sim p(x_{\text{obs}} | \theta_{\text{true}})$.

Then, estimate the posterior

$$p(\theta | x_{\text{obs}}) = \frac{p(x_{\text{obs}} | \theta)p(\theta)}{p(x_{\text{obs}})}$$


Or a likelihood ratio

$$r(\theta) = \frac{p(x_{\text{obs}} | \theta)}{p(x_{\text{obs}} | \theta_0)}$$



Neural Ratio Estimation

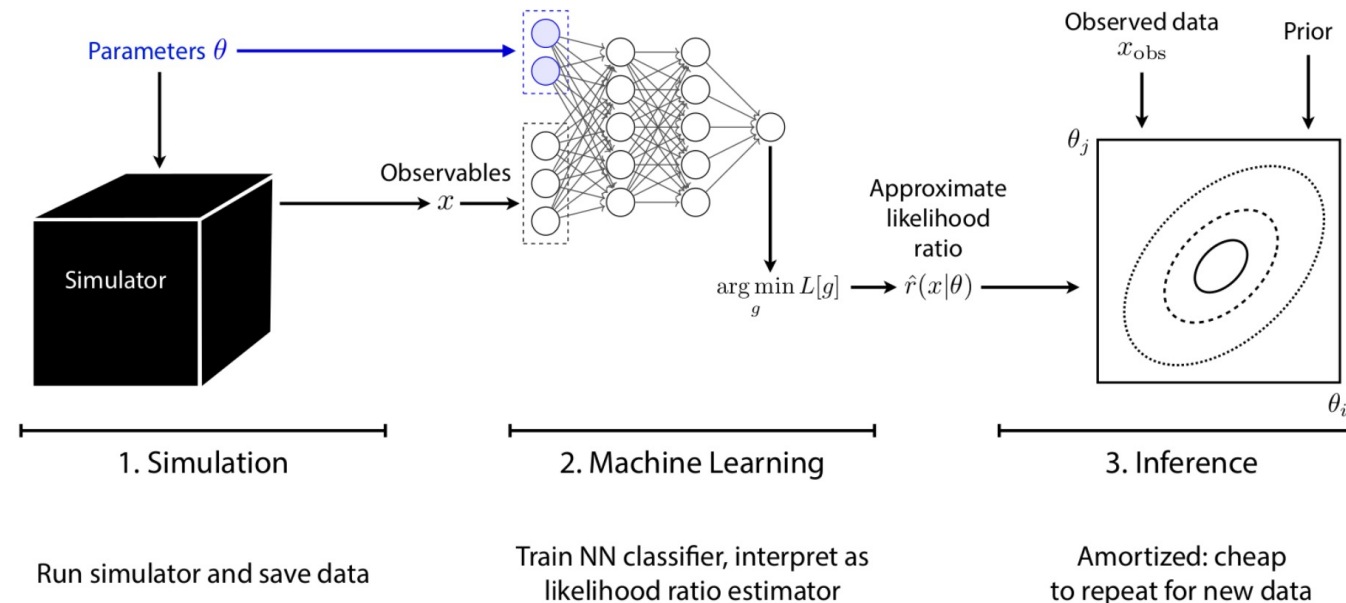
The likelihood-to-evidence $r(x|\theta) = \frac{p(x|\theta)}{p(x)} = \frac{p(x,\theta)}{p(x)p(\theta)}$ ratio can be learned, even if neither the likelihood nor the evidence can be evaluated:



Neural Ratio Estimation

The likelihood-to-evidence $r(x|\theta) = \frac{p(x|\theta)}{p(x)} = \frac{p(x,\theta)}{p(x)p(\theta)}$ ratio can be learned, even if neither the likelihood nor the evidence can be evaluated:

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)} \approx \hat{r}(x|\theta)p(\theta)$$



Can Inference Goal Drive Training?

Train summary statistic $T_w(x)$ to optimize inference goal

Examples: [NEOS](#) and [INFERN0](#)

