The Ecology of Uncertainty : Capturing Epistemic Doubt in Hadronic Structure Analyses

Brandon Kriesten ∘ Argonne National Laboratory

QCD Evolution 2025

Parton Distribution Functions

arXiv: 2503.XXXX for QIS prescription

PDFs describe how the hadron's momentum is carried by the constituents (quarks and gluons), and are defined through factorization theorems.

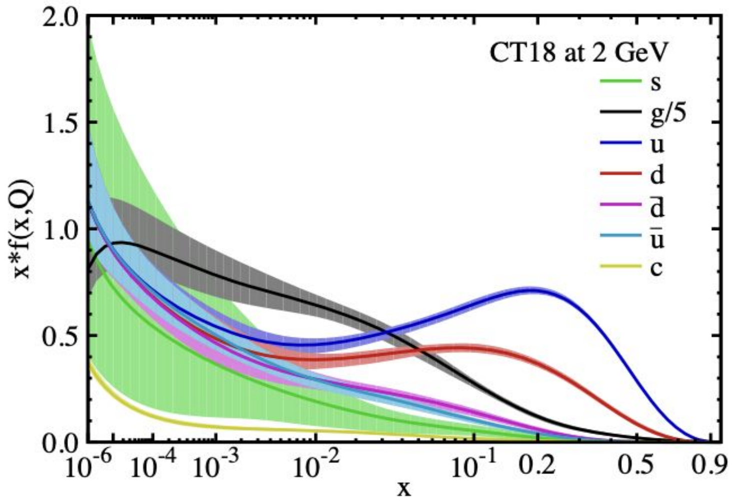

Image credit: PhysRevD.103.014013

PhysRevD.103.014013

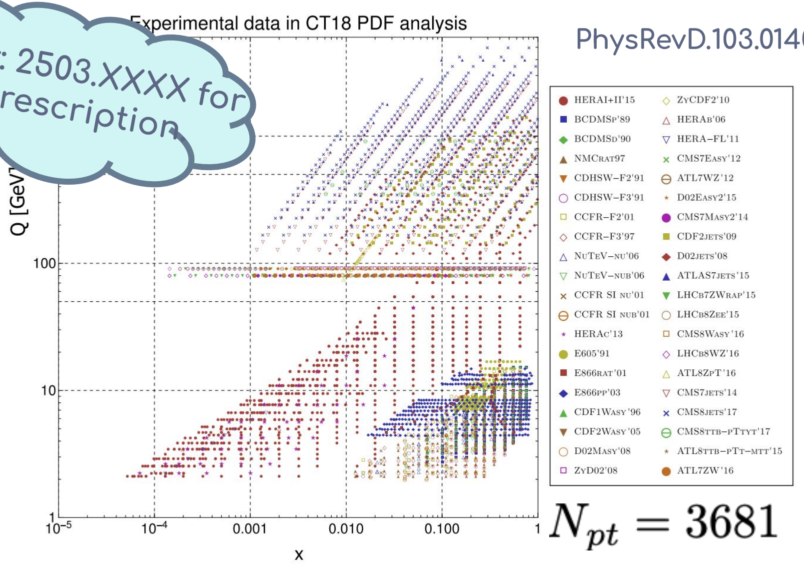

$N_{pt} = 3681$

Global PDF fits unify data across all sectors of the SM, building a robust picture of hadron structure and effectively mediating our knowledge / or lack thereof of fundamental interactions.

PDF uncertainties

PDF uncertainties are an important limitation to discovery reach at colliders.

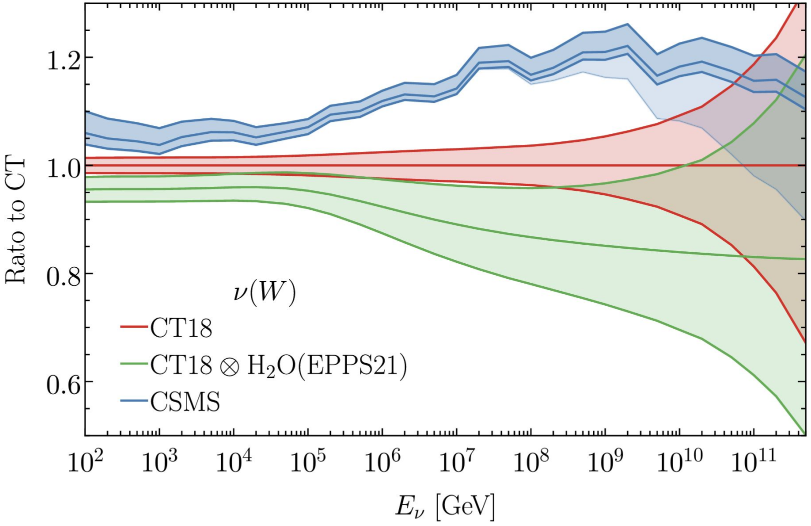

Image credit: PhysRevD.109.113001

Neutrino DIS cross sections to astrophysical scales

Image credit: PhysRevD.103.014013

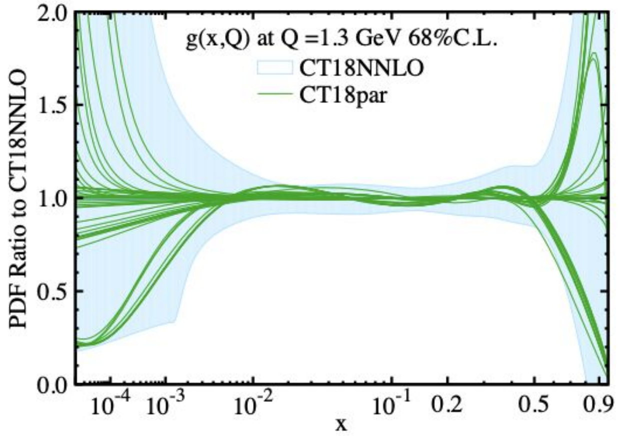

Uncertainties arise from extrapolation, parameterization dependence (assumed prior knowledge of functional form), theory assumptions, and data uncertainties.

This is a multi-faceted problem spanning many energy scales and aspects of QCD!

PDFs (Hadronic QCFs) + SMEFT / BSM (fit contaminations)

BSM + SMEFT contaminations impact uncertainties calculated on QCFs. Alludes to a very relevant question in hadronic structure … Can QCFs absorb signatures of new physics?

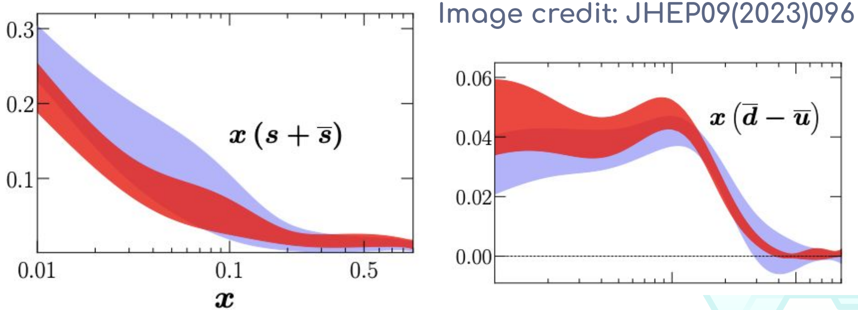

Image credit: JHEP09(2023)096

Image credit: JHEP11(2023)090

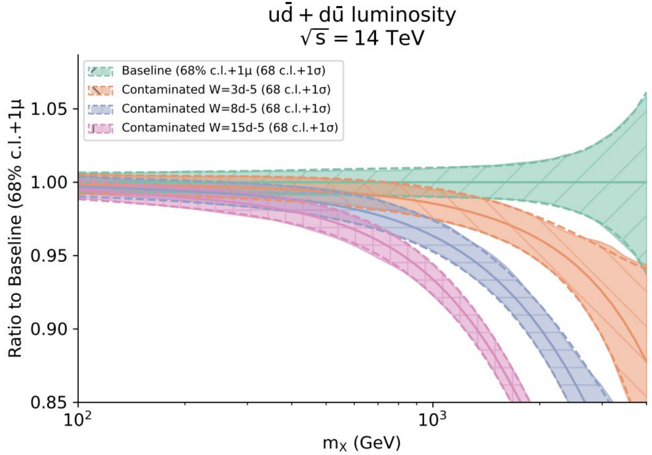

Image credit: PhysRevLett.115.162001

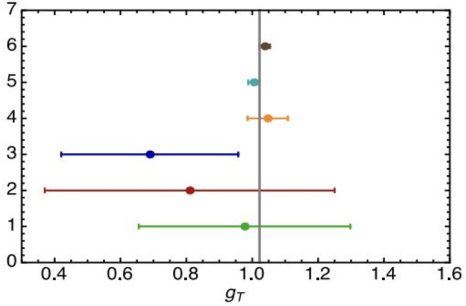

Not just an LHC phenomenon. Very relevant for JLAB and upcoming EIC!

BSM physics is a case study for intersections of Hadronic + High Energy physics. Develop tools to combine analyses to discover NP.

Bridging the Gap with AI/ML: Foundation models for fundamental physics

A foundation model using a shared embedding space can be used to perform specialized downstream tasks: inference, UQ, BSM physics, anomaly detection, emergent phenomena.

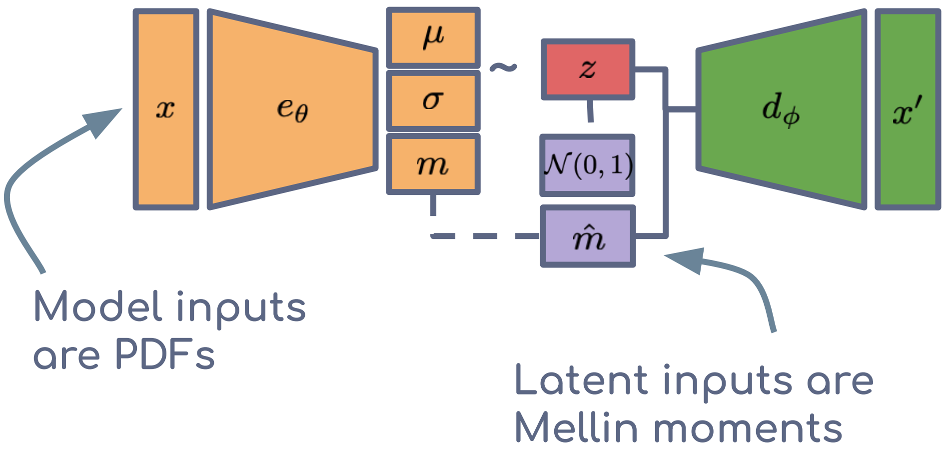

Therefore it is necessary to understand how to <u>create these embedding spaces</u> in an interpretable way with rigorous uncertainty quantification - benchmarking task.

Initial attempts through variational autoencoders inverse mapping to PDFs from LQCD observables.

BK, T.J. Hobbs **PRD 111 (2025) 1**

Interpreting Learned Physics from AI: Inverse Mapping and XAI

<u>Goal</u>: to go from measured data / LQCD observables to PDFs in a high dimensional space.

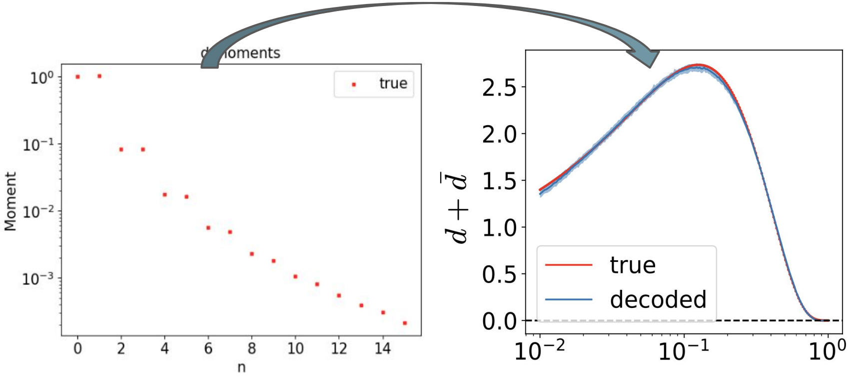

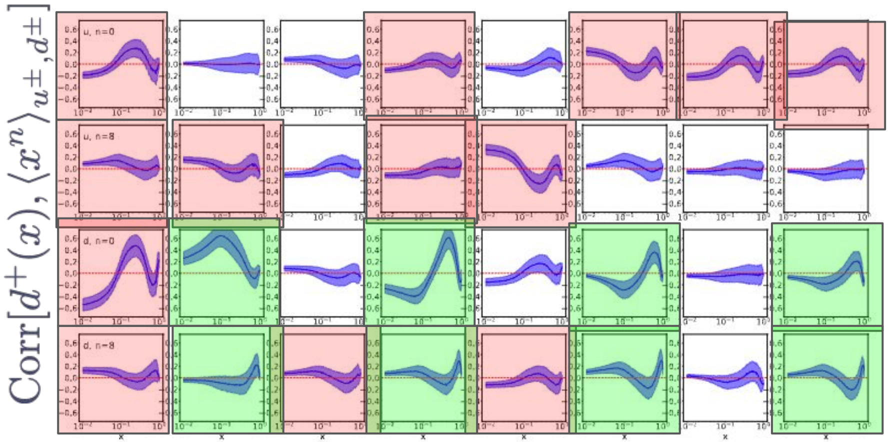

Spurious correlations in red - not physics.

BK, T.J. Hobbs **PRD 111 (2025) 1**

Generative models can hide spurious correlations, we want to see which features the model is looking at while learning physics.

Interpreting Learned Physics from AI: Inverse Mapping and XAI

Goal: to go from measured data / LQCD observables to PDFs in a high dimensional space.

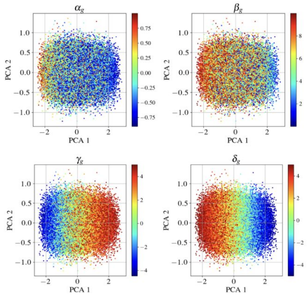

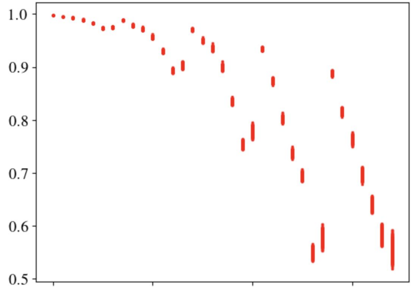

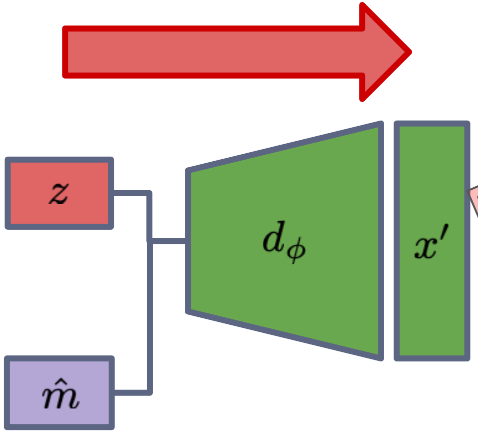

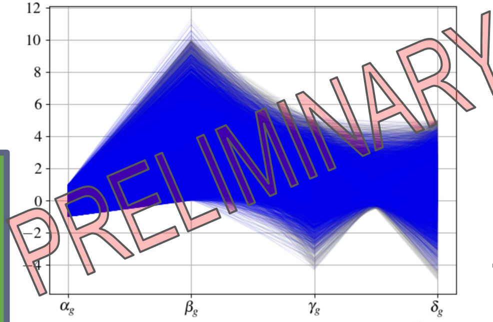

Generating gluon PDF from lattice QCD Ioffe-time distributions with uncertainty quantification.

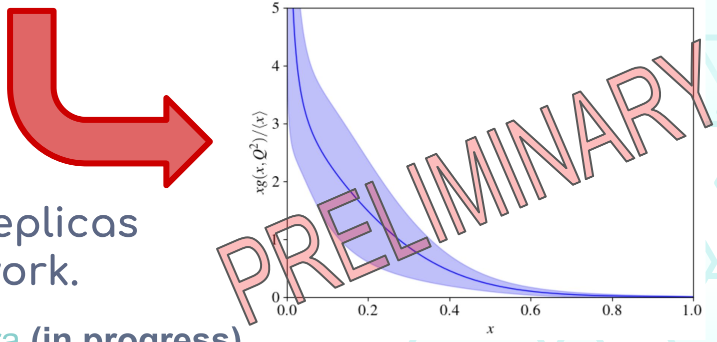

Jack-knife samples to MC replicas through end-to-end framework.

W Goode, T.J. Hobbs, BK, H-W Lin, A Niemiera **(in progress)**

Interpreting Learned Physics from AI: Inverse Mapping and XAI

Can we trust how the ML model organizes physics in embedding spaces? Techniques such as guided backprop can be used.

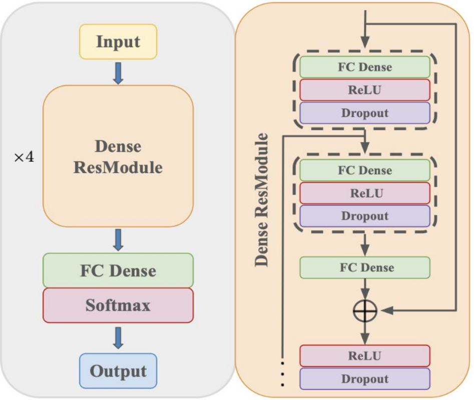

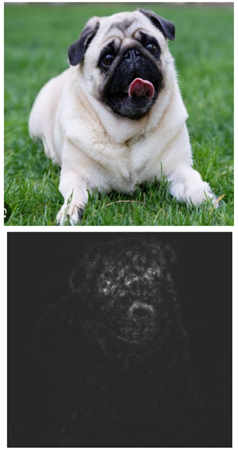

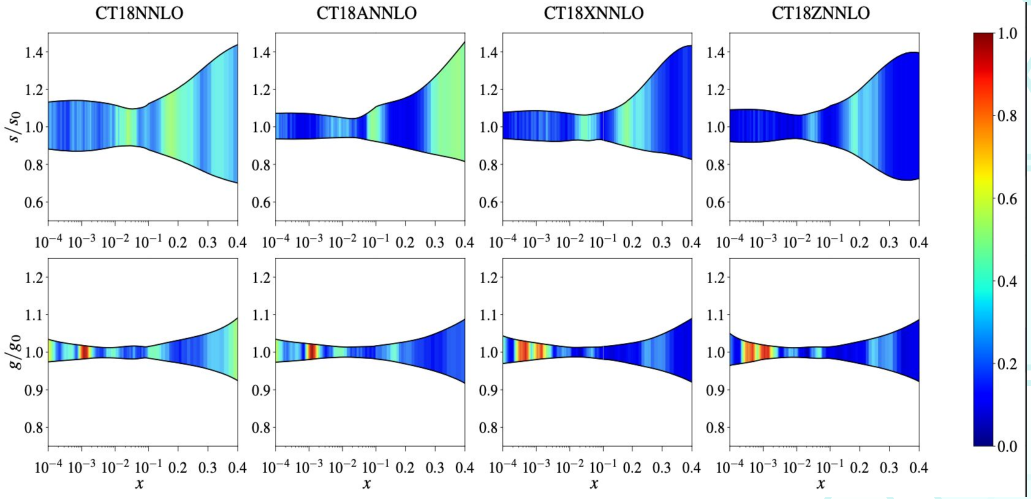

How do we quantify our uncertainty from these models?

BK, J. Gomprecht, T.J. Hobbs **JHEP 11 (2024) 007**

Probabilistic AI / ML

Probabilistic AI / ML is a mathematical paradigm of defining the outputs of algorithms in the language of probabilistic distributions. As an example, for classification algorithms the outputs are categorical distributions.

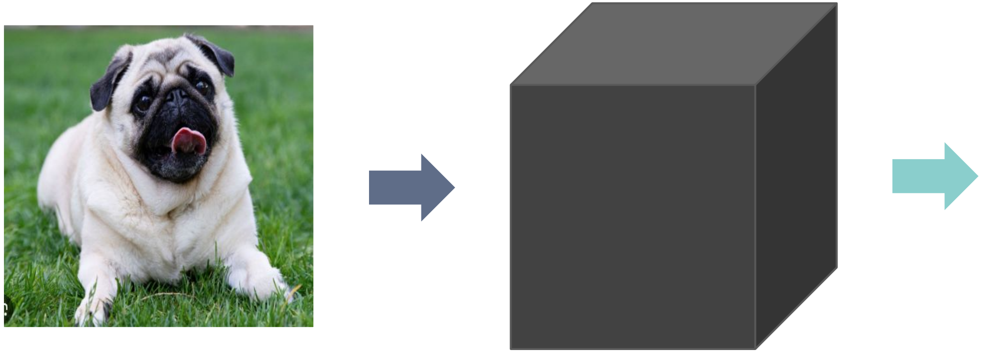

Dog: 98 %
Cat: 1.8 %
Bird: 0.2 %

Uncertainty Quantification for Machine Learning

AI / ML algorithms can not only be wrong, but also be really confidently wrong!

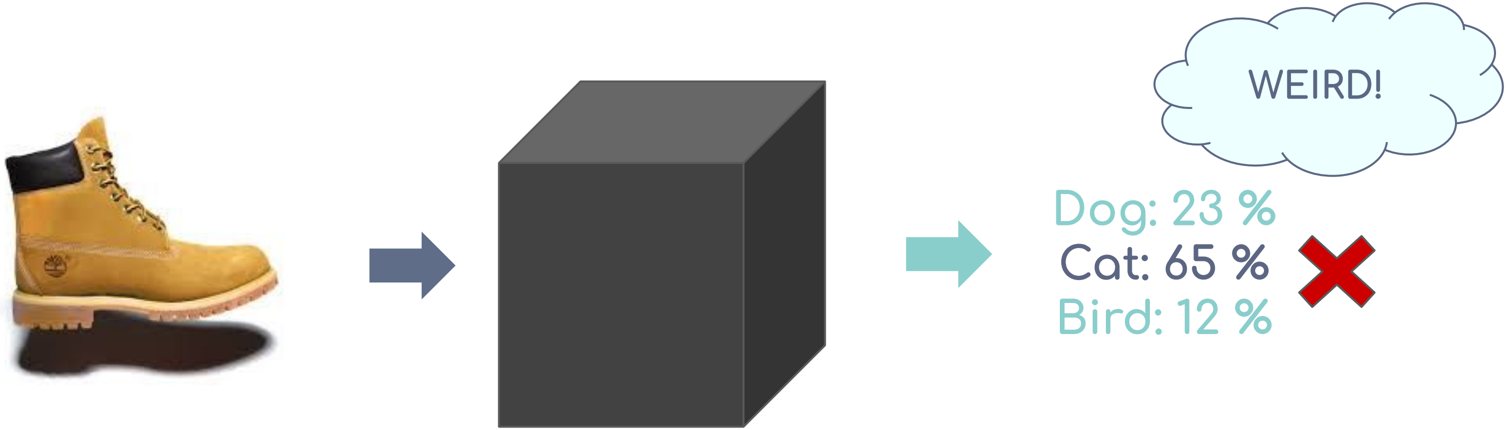

Uncertainty Quantification for Machine Learning

What we want is something more like this, but how do we teach an ML algorithm to say 'I don't know?'

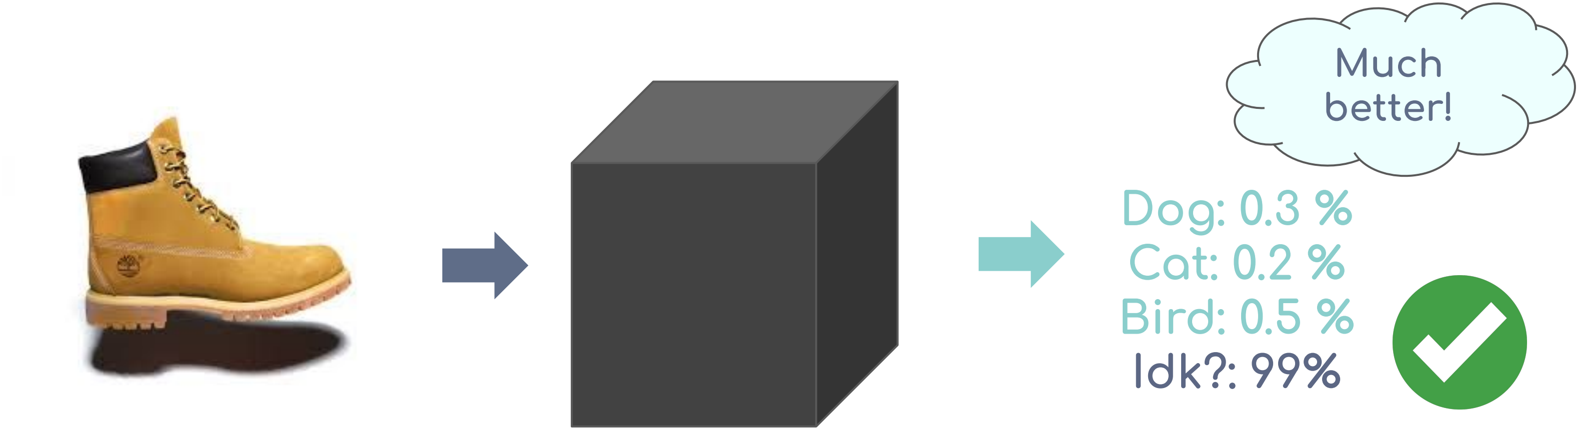

How is this put into practice for classification models?

Argonne NATIONAL LABORATORY

Uncertainty Quantification for ML-based Classification

Bayesian Neural Networks

$$p(w_c|x^*, \mathcal{D}) = \int p(w_c|x^*, \theta)p(\theta|\mathcal{D})$$

whereby Monte Carlo sampling the model parameters, we create an ensemble of categorical distributions

$$\left\{p(w_c|x^*, \theta^{(i)}\right\}_{i=1}^{M}$$

Computationally expensive!

What if instead we predicted the parameters of the ensemble - the conjugate prior to the categorical distribution … a Dirichlet!

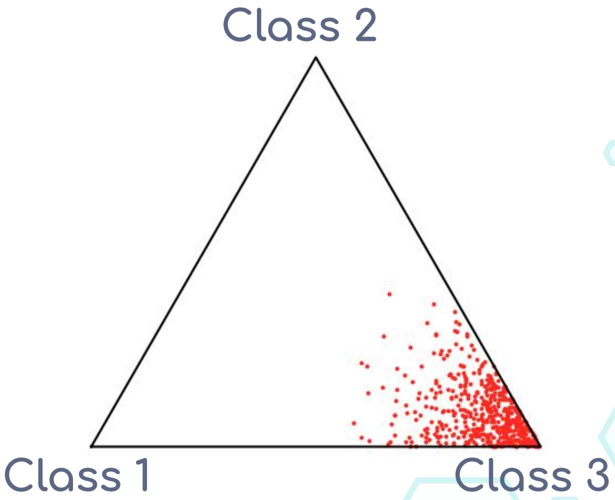

Dirichlet Prior Networks

A categorical distribution models individual probabilities of specific categories $(p_1, p_2, \ldots, p_k)$. The Dirichlet distribution models the prior beliefs over each category's probabilities, designated by parameters $(\alpha_1, \alpha_2, \ldots, \alpha_k)$.

$$p(\mu; \alpha) = \frac{\Gamma(\alpha_o)}{\prod_{k=1}^{K} \Gamma(\alpha_k)} \prod_{k=1}^{K} p_k^{\alpha_k - 1}$$

So now we predict the alpha parameters and sample the Dirichlet to get as many categorical distributions as we want.

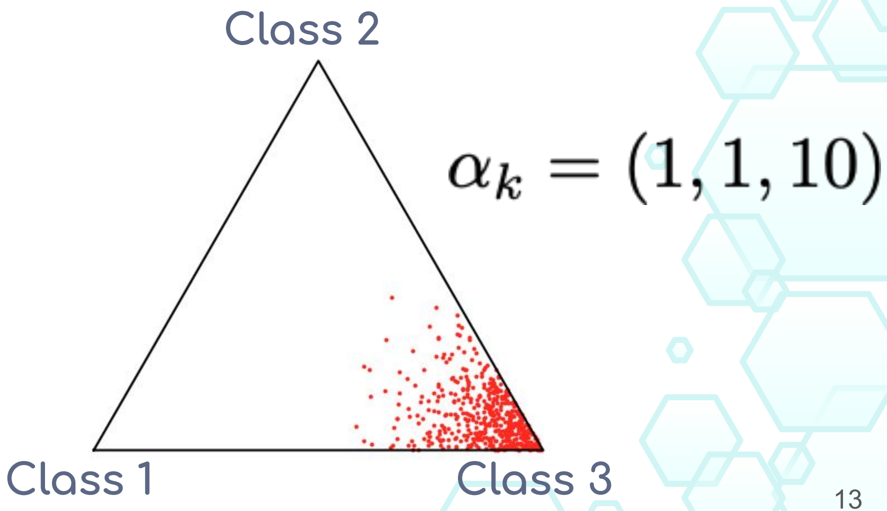

Defining Aleatoric and Epistemic Uncertainty

Through training with maximum likelihood estimation (MLE) we can rework the expression to factorize into two distinct components. One is epistemic - the KL divergence term; and the other is aleatoric - the entropy term.

$$\mathbb{E}_{p_{true}(x,y)}\left[\mathcal{L}^{NLL}(y,x,\theta)\right] = \mathbb{E}_{p_{true}(x)}\left[\boxed{D_{KL}\left(p_{true}(y|x) \middle\| p(y|x,\theta)\right)} - \boxed{\mathbb{H}\left(p_{true}(y|x)\right)}\right]$$

Epistemic

Epistemic uncertainty is reducible by training procedures or algorithmic development. Also can be influenced by choice of distribution to model the data.

Aleatoric

Aleatoric uncertainty is not reducible because it is related to the true underlying data distribution.

Malinin and Gales arXiv:1802.10501

A new uncertainty emerges … distributional uncertainty

The epistemic uncertainty seems to naturally separate as well:

- a contribution arising from modelling by the AI / ML algorithm
- a contribution arising from the choice of underlying distribution to describe the data.

$$p(w_c|x^*, \mathcal{D}) = \int\int p(w_c|\mu)\boxed{p(\mu|x^*, \theta)p(\theta|\mathcal{D})}d\theta d\mu$$

In theory, we can model all three at the same time (aleatoric, distributional, epistemic) - creating an ecology of uncertainties!

Malinin and Gales arXiv:1802.10501

Possible translation to physics-based observables: PRD 109(2024)7

An ecology of uncertainties and metrics

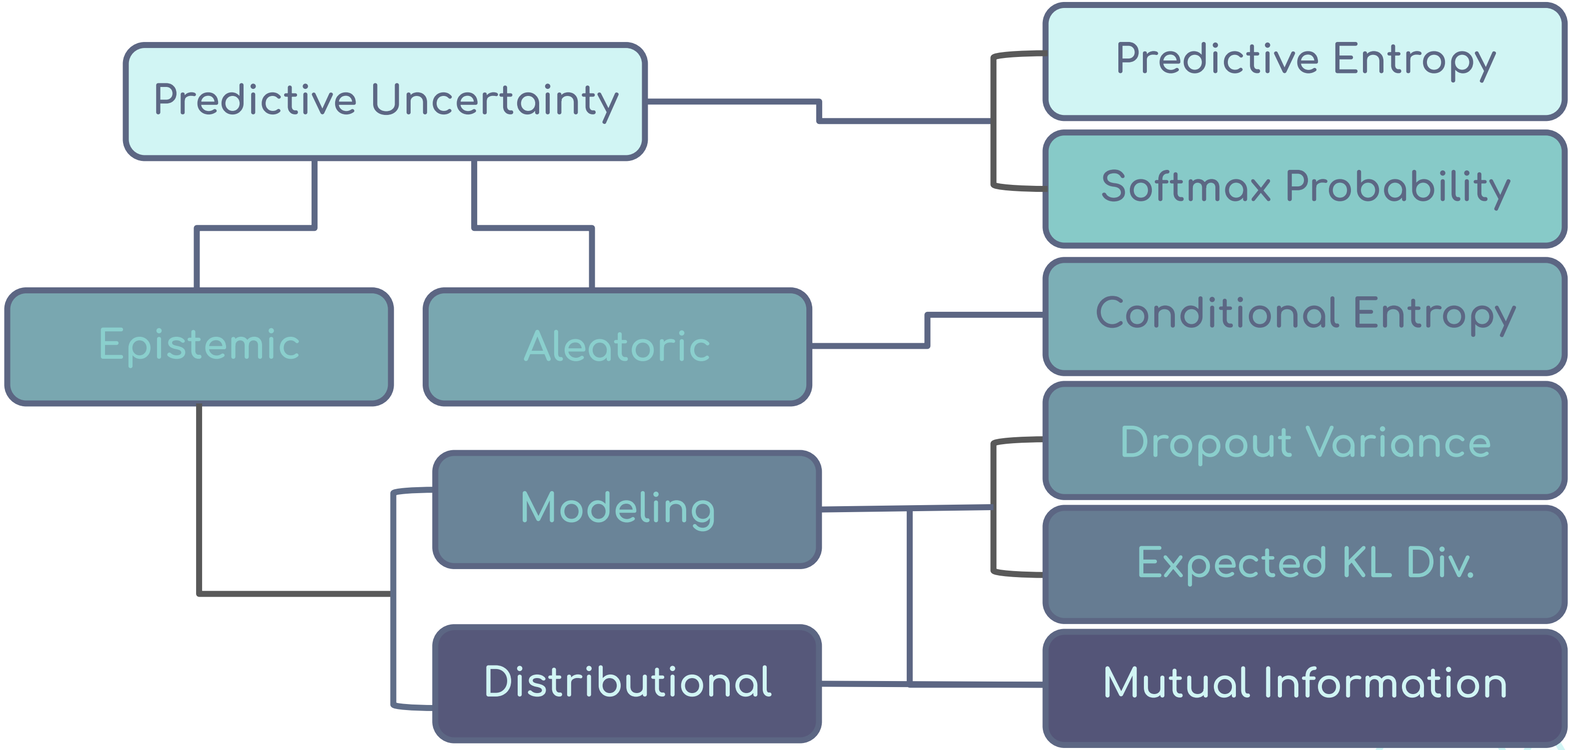

Dirichlet Prior Networks - an example

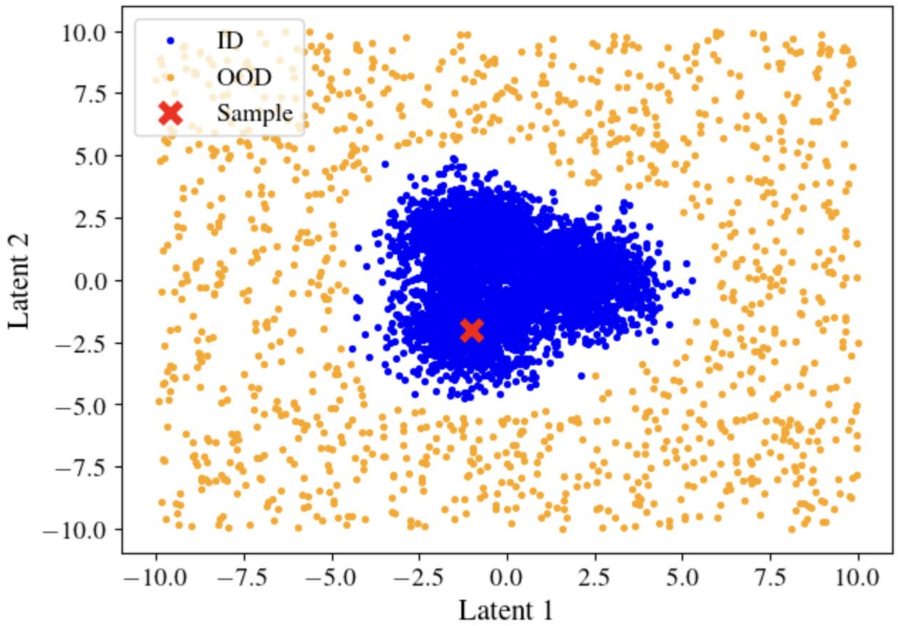

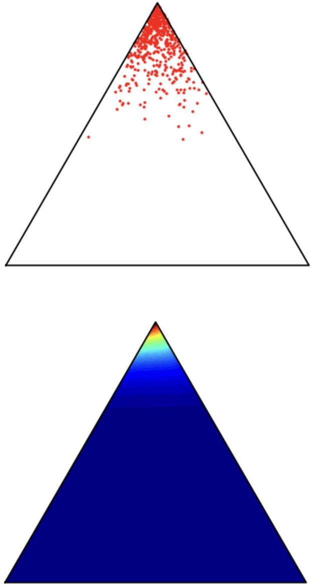

$$\alpha_k = (1, 10, 1)$$

In-distribution sampling with low data uncertainty (samples are located at a specific corner) and low knowledge uncertainty (high sample density).

Dirichlet Prior Networks - an example

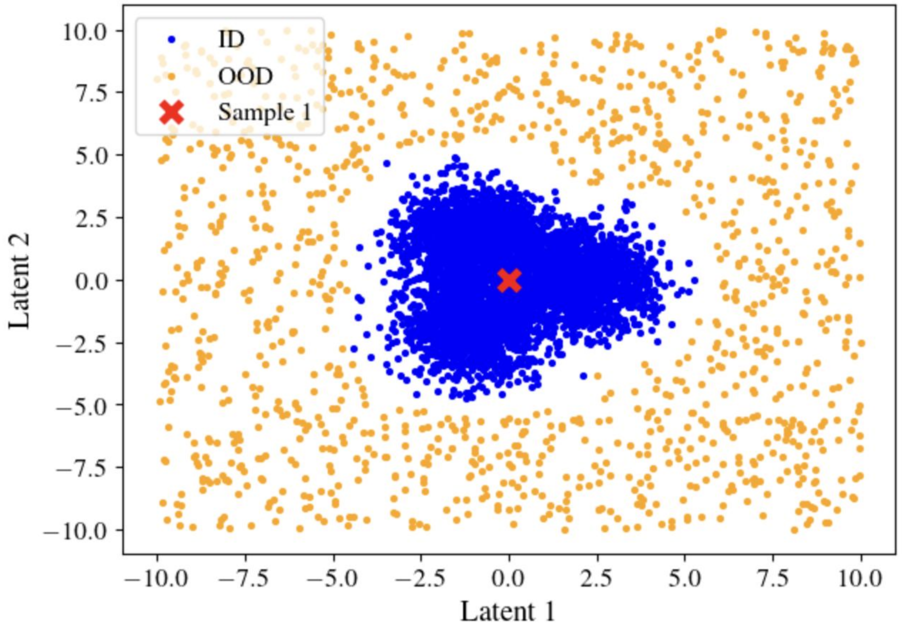

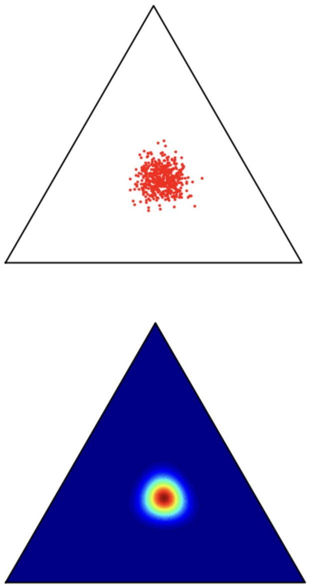

$$\alpha_k = (8, 8, 8)$$

In-distribution sampling with high data uncertainty (samples are squeezed to the center of the simplex) with low knowledge uncertainty (high sample density).

Dirichlet Prior Networks - an example

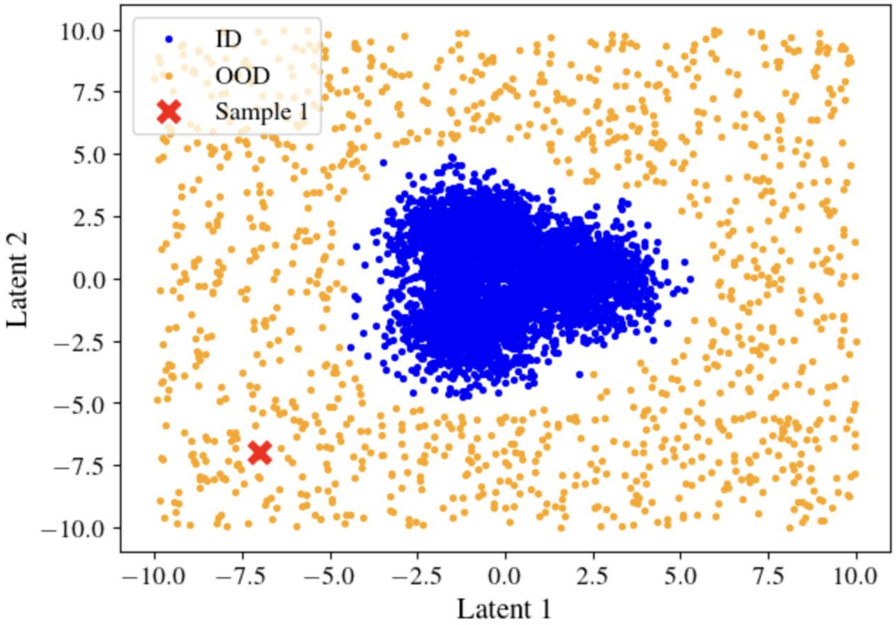

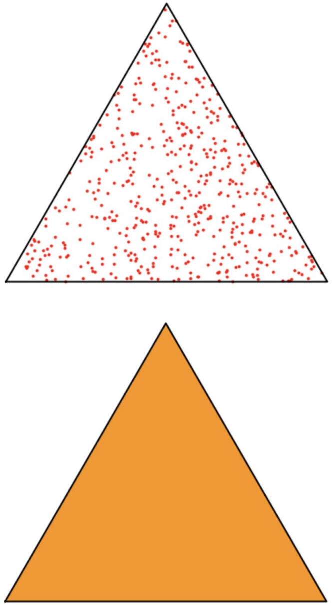

$$\alpha_k = (1, 1, 1)$$

Out-of-distribution sampling with high data uncertainty (samples are not located at a specific corner) and high knowledge uncertainty (samples are diffuse).

Information Theory-based Quantitative Metrics of Uncertainty

We can separate the total classification predictive uncertainty into contributions from aleatoric and epistemic through analytic expressions of the Dirichlet parameters.

$$\mathbb{E}_{p(\mu|\mathbf{x}^*,\hat{\theta})}\left[\mathbb{H}[p(y|\mu)]\right] = \sum_{k=1}^{K} -\frac{\alpha_k}{\alpha_o}\left(\psi(\alpha_k+1) - \psi(\alpha_o+1)\right)$$

$$\mathcal{I}[y,\mu|x^*,\mathcal{D}] = \mathbb{H}\left[\mathbb{E}_{p(\mu|x^*,\mathcal{D})}[p(y|\mu)]\right] - \mathbb{E}_{p(\mu|\mathbf{x}^*,\mathcal{D})}\left[\mathbb{H}[p(y|\mu)]\right]$$

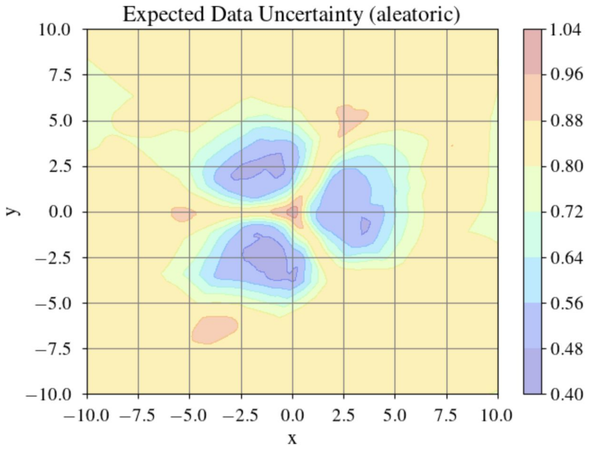

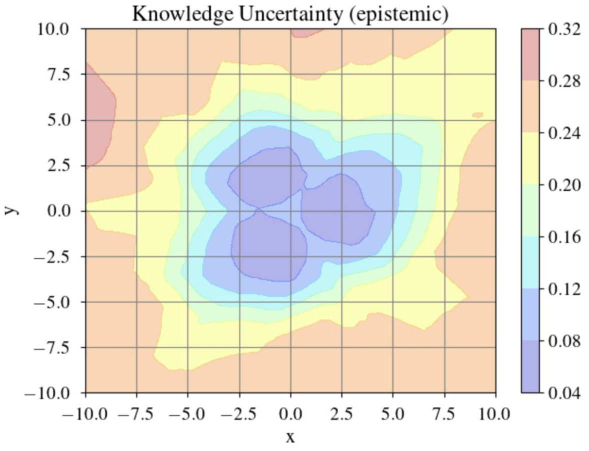

Quantitatively discriminating between models - a BSM scenario

Consider a situation in νDIS of disentangling SM physics from BSM physics realized in AEWI (anomalous electroweak interactions) where the SM EW parameters - the CKM matrix elements - are shifted within the discovery potential of ~ few σ while maintaining a ~1% shift at the level of the nucleon structure functions.

Is it possible to disentangle this physics using experimental observables such as the NNLO F2 structure function? (e.g. slide 4)

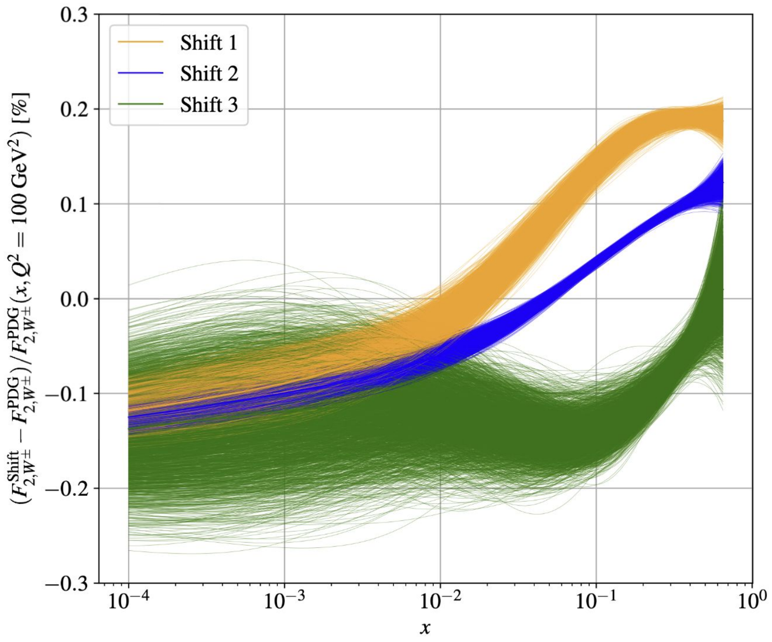

BK, T.J. Hobbs **arXiv: 2412.16286**

Quantitatively discriminating between models - a BSM scenario

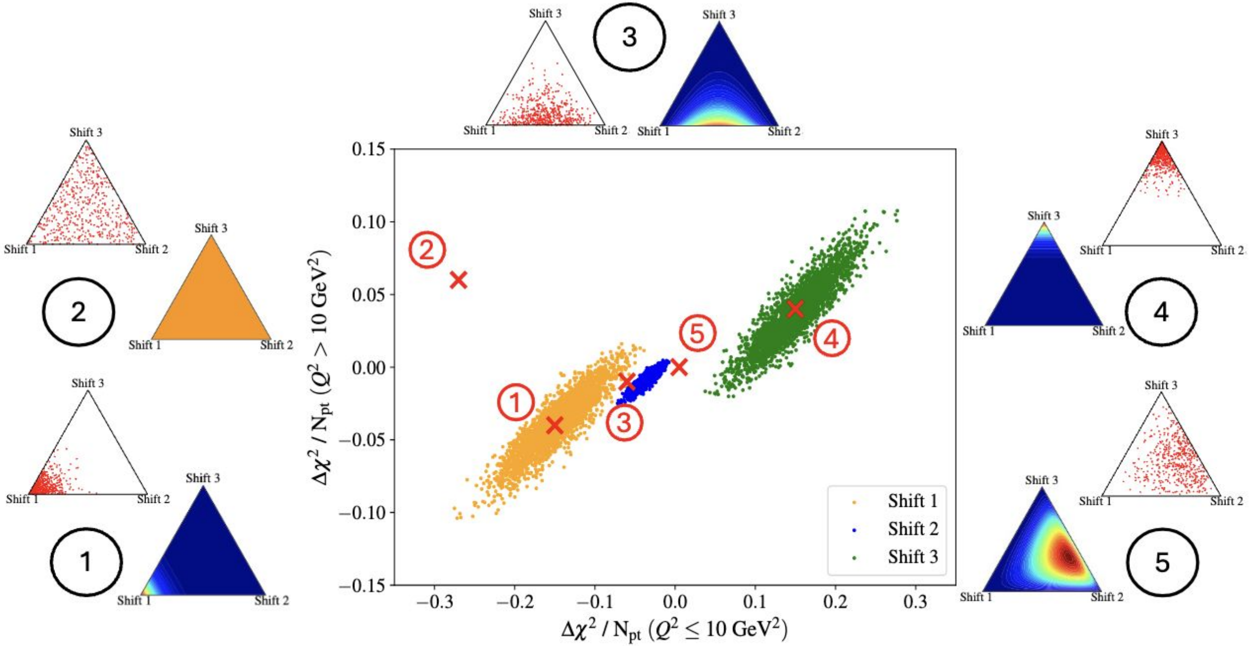

Dimensionally reduce to a calculated $\Delta\chi^2/N_{pt}$ statistic at $Q^2 > 10$ GeV^2 and $Q^2 \leq 10$ GeV^2 on the CDHSW νDIS dataset in the various shifted AEWI scenarios. We can then use the EDL framework to quantify classification uncertainty in this space.

BK, T.J. Hobbs **arXiv: 2412.16286**

Quantitatively discriminating between models - a BSM scenario

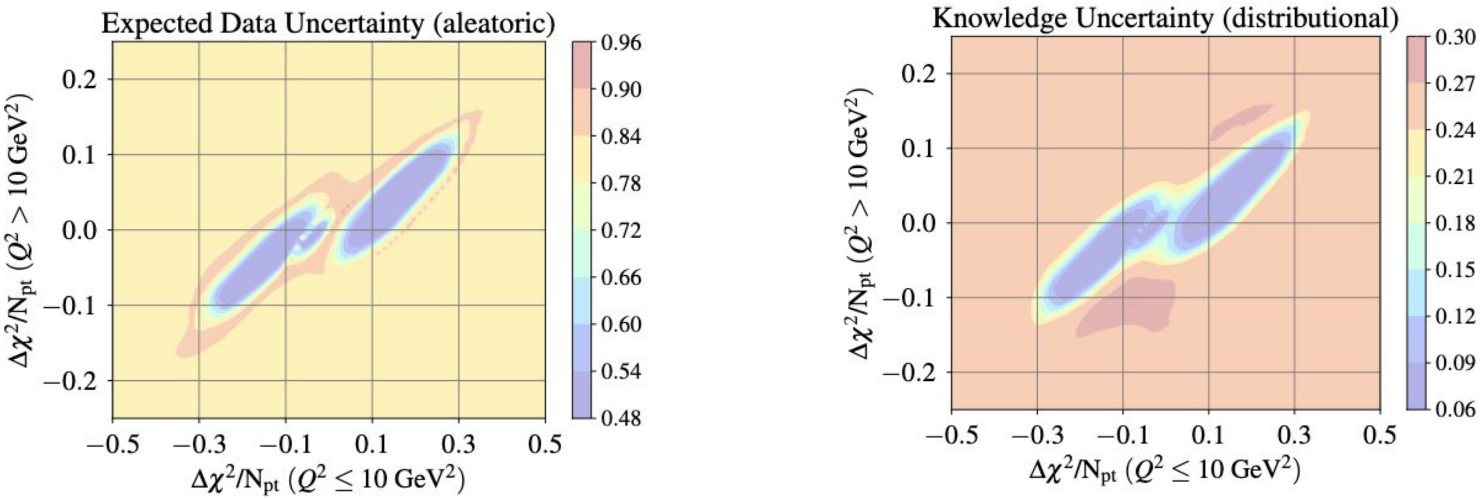

The information theory metrics can be used for high dimensional inputs where it may not be obvious where such overlaps occur.

BK, T.J. Hobbs **arXiv: 2412.16286**

Inverse problems in BSM model selection

We consider a BSM scenario - this methodology is relevant for any model selection - of CP - odd observables in SMEFT and 2HDM.

$4: X^2H^2$	
Q_{HG}	$H^\dagger H\, G^A_{\mu\nu} G^{A\mu\nu}$
$Q_{H\widetilde{G}}$	$H^\dagger H\, \widetilde{G}^A_{\mu\nu} G^{A\mu\nu}$
Q_{HW}	$H^\dagger H\, W^I_{\mu\nu} W^{I\mu\nu}$
$Q_{H\widetilde{W}}$	$H^\dagger H\, \widetilde{W}^I_{\mu\nu} W^{I\mu\nu}$
Q_{HB}	$H^\dagger H\, B_{\mu\nu} B^{\mu\nu}$
$Q_{H\widetilde{B}}$	$H^\dagger H\, \widetilde{B}_{\mu\nu} B^{\mu\nu}$
Q_{HWB}	$H^\dagger \tau^I H\, W^I_{\mu\nu} B^{\mu\nu}$
$Q_{H\widetilde{W}B}$	$H^\dagger \tau^I H\, \widetilde{W}^I_{\mu\nu} B^{\mu\nu}$

arXiv:1312.2014v5

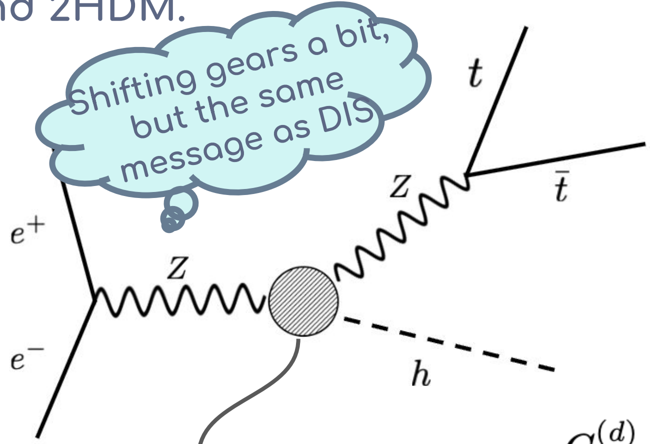

$$\mathcal{L}_{\mathrm{SMEFT}} = \mathcal{L}_{\mathrm{SM}} + \sum_{d>4} \sum_i \frac{C_i^{(d)}}{\Lambda^{d-4}} \mathcal{O}_i^{(d)}$$

Question: If we can generate CP violating contributions in the Higgs sector through a variety of NP mechanisms, how do we know which model (and subsequently which new physics) is responsible in experimental observables?

S. Das Bakshi, T.J. Hobbs, BK **(in progress)**

Embedding BSM scenarios in an embedding space

We take a set of experimental observables and project 18 BSM models to a single latent space to separate and create the basis of a foundation model to perform several downstream tasks.

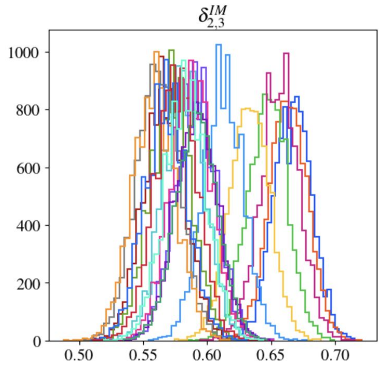

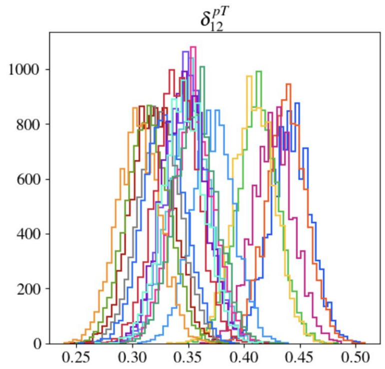

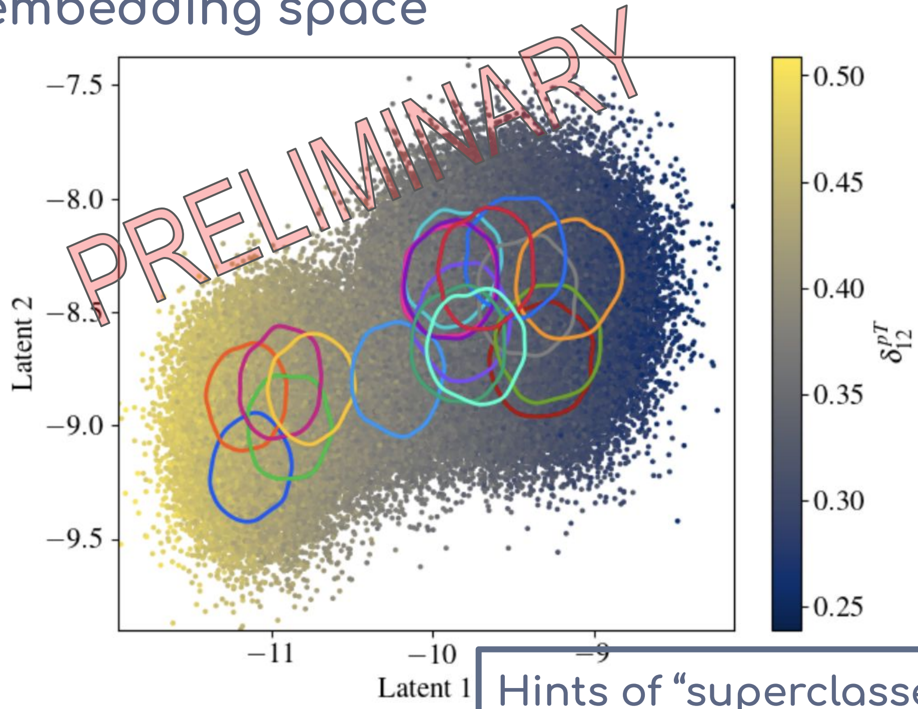

Hints of “superclasses”

S. Das Bakshi, T.J. Hobbs, BK (in progress)

Future Work: Hyper Opinions and uncertainty on the ground truth

Consider a scenario where there exists an uncertainty on the ground truth label itself. How do we rule out classification labels?

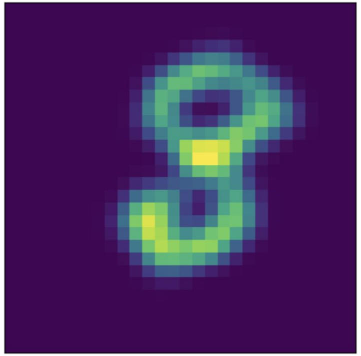

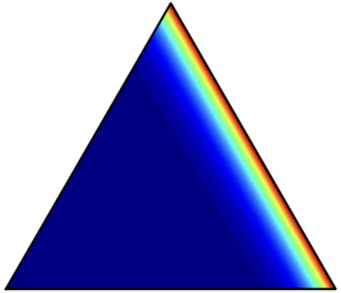

Example: the above illustration with MNIST. It is uncertain whether the ground truth is 3 or 8 but it certainly isn't 0.

We can use a Grouped Dirichlet distribution to model not only the distribution over the priors of the categorical, but also the distribution over composite labels.

$$\mathrm{GDD}(p|\alpha, c) = Z^{-1} \prod_{k=1}^{K} p_k^{\alpha_k - 1} \prod_{j=1}^{\eta} \left(\sum_{l \in S_j} p_l \right)^{c_j}$$

Extends physics use case to rule out models from classification.

BK, T.J. Hobbs **(in progress)**

Future Work: Towards a Hadronic physics + BSM foundation model

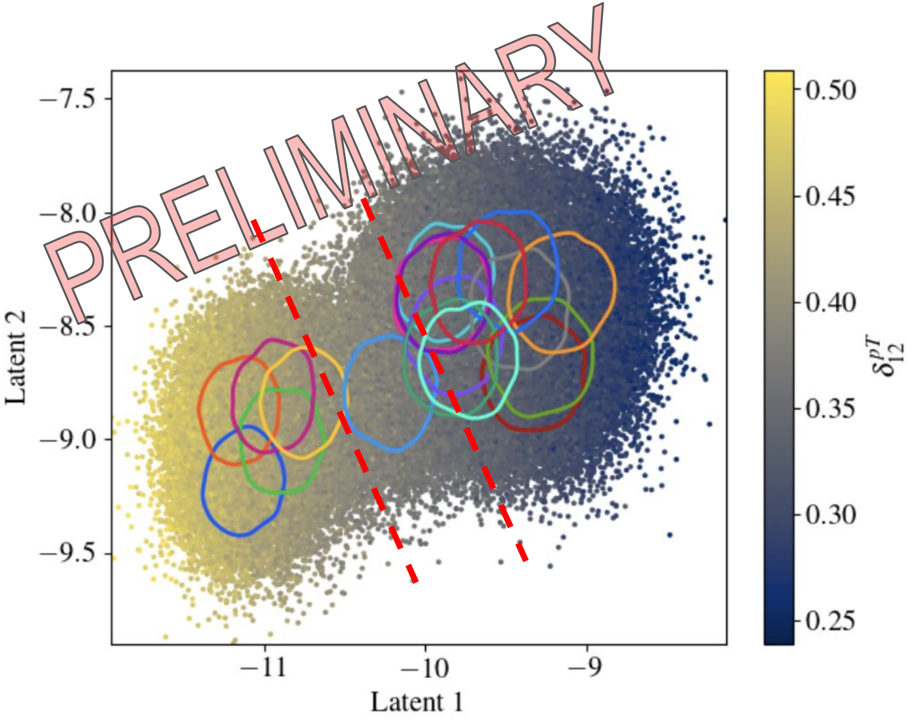

Distinguishing BSM "superclasses" with information theory metrics and grouping models based on a perceived uncertainty on the ground truth label.

pp/ep scattering for EIC + LHC, incorporation of PDF uncertainties.

Entropy in the embedding space translates to uncertainty in the predictions of observables?

Everything here can be directly translated to GPDs / TMDs in which PDFs play an important constraint.

S. Das Bakshi, T.J. Hobbs, BK **(in progress)**

Conclusions / Outlooks

- A comprehensive inverse mapping framework for PDFs using AI/ML tools unifies PDF extractions across energy scales with different inputs including LQCD.
- EDL for UQ decomposes uncertainty into aleatoric, epistemic, and distributional components. Possible translation to hadronic structure uncertainties leading to targeted improvements of errors.
- We can use AI/ML to unfold BSM physics from hadronic structure global fits and isolate new-physics signals from SM backgrounds.
- As the next generation of particle physics experiments come online (EIC, HL-LHC, etc), foundation models offer a route for robust, data-driven predictions for new physics observables - if we can benchmark these tools against known physics.

This work at Argonne National Laboratory was supported by the U.S. Department of Energy under contract DE-AC02-06CH11357.

U.S. DEPARTMENT of ENERGY

Thank you for your attention!

This work at Argonne National Laboratory was supported by the U.S. Department of Energy under contract DE-AC02-06CH11357.