Bayesian NN for UQ in neutrino classification

Aaron Higuera

NPML 2025 University of Tokyo

Classification in LArTPC

- Classification
  - Neutrino flavor
  - Particle ID
  - Hit-level ID (Semantic segmentation)
- Regression
  - Momentum
  - Angle
- Tools:
  - Convolution Neural Networks
  - Sparse U-Net
  - Graphical Convolution Networks
  - Vision-Language Models
  - Quanvolutional NN

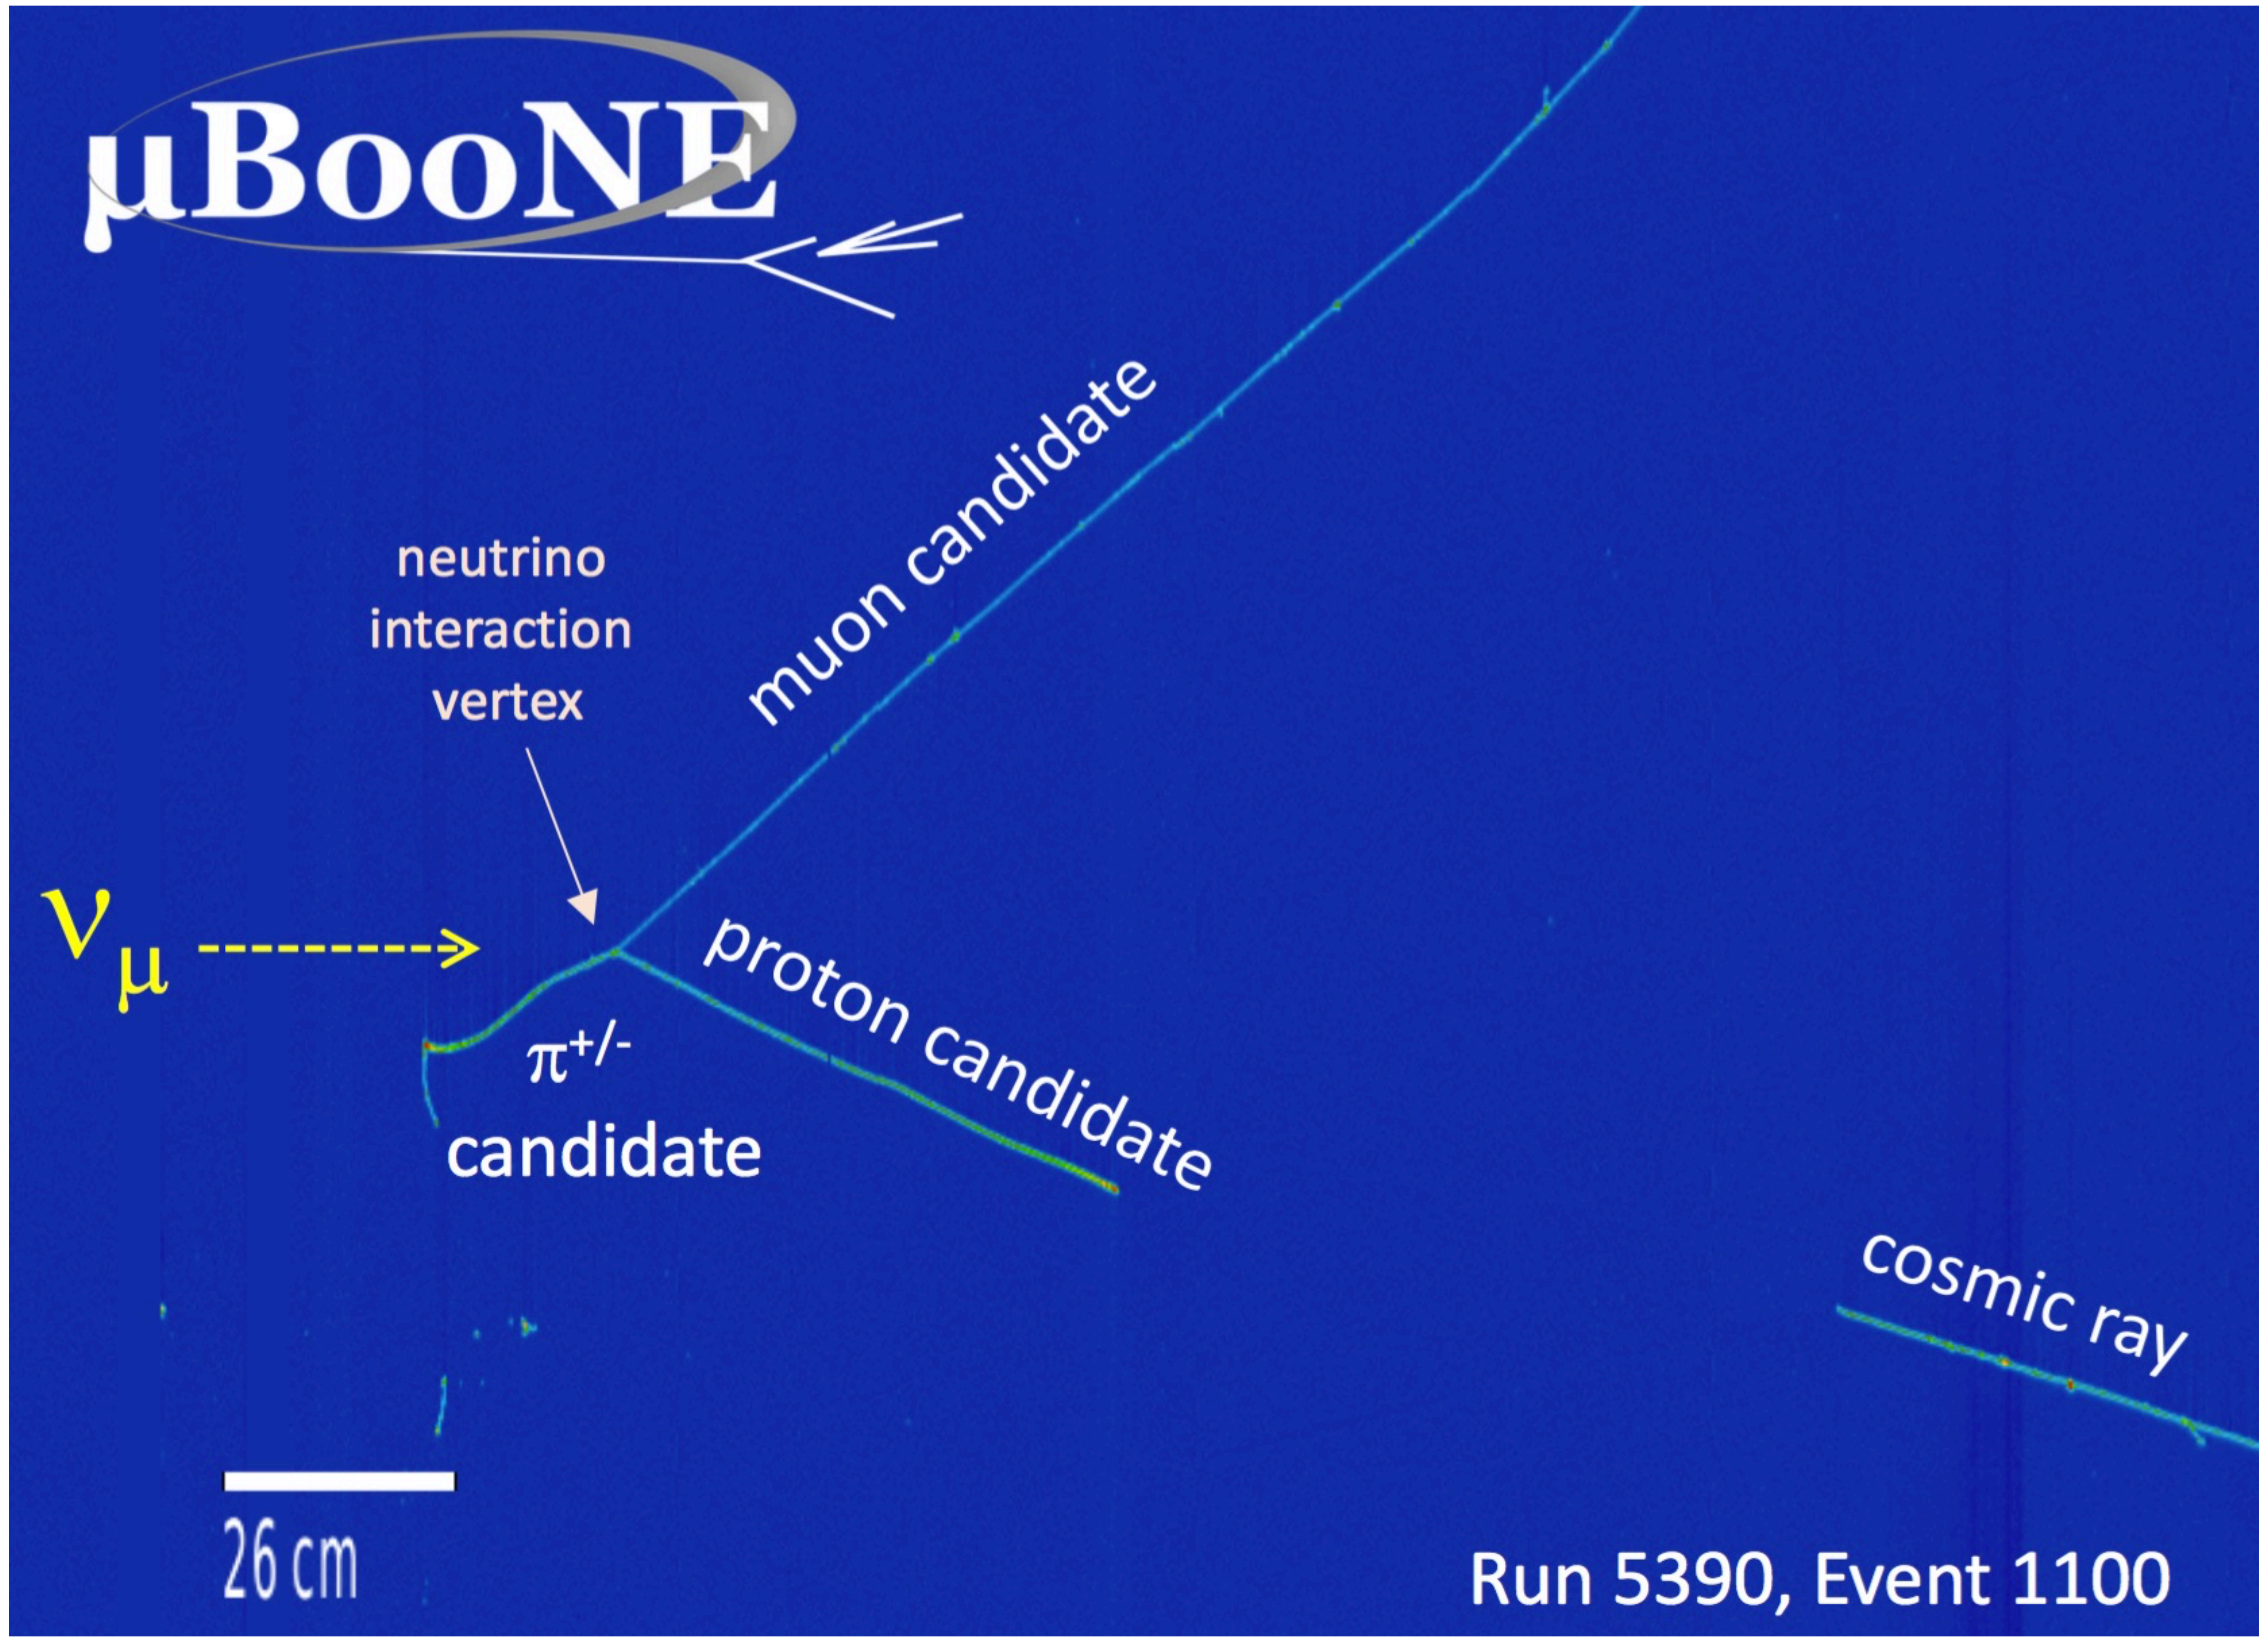

Classification

Classification

- Neutrino flavor
- Particle ID
- Hit-level ID

How robust is the method?

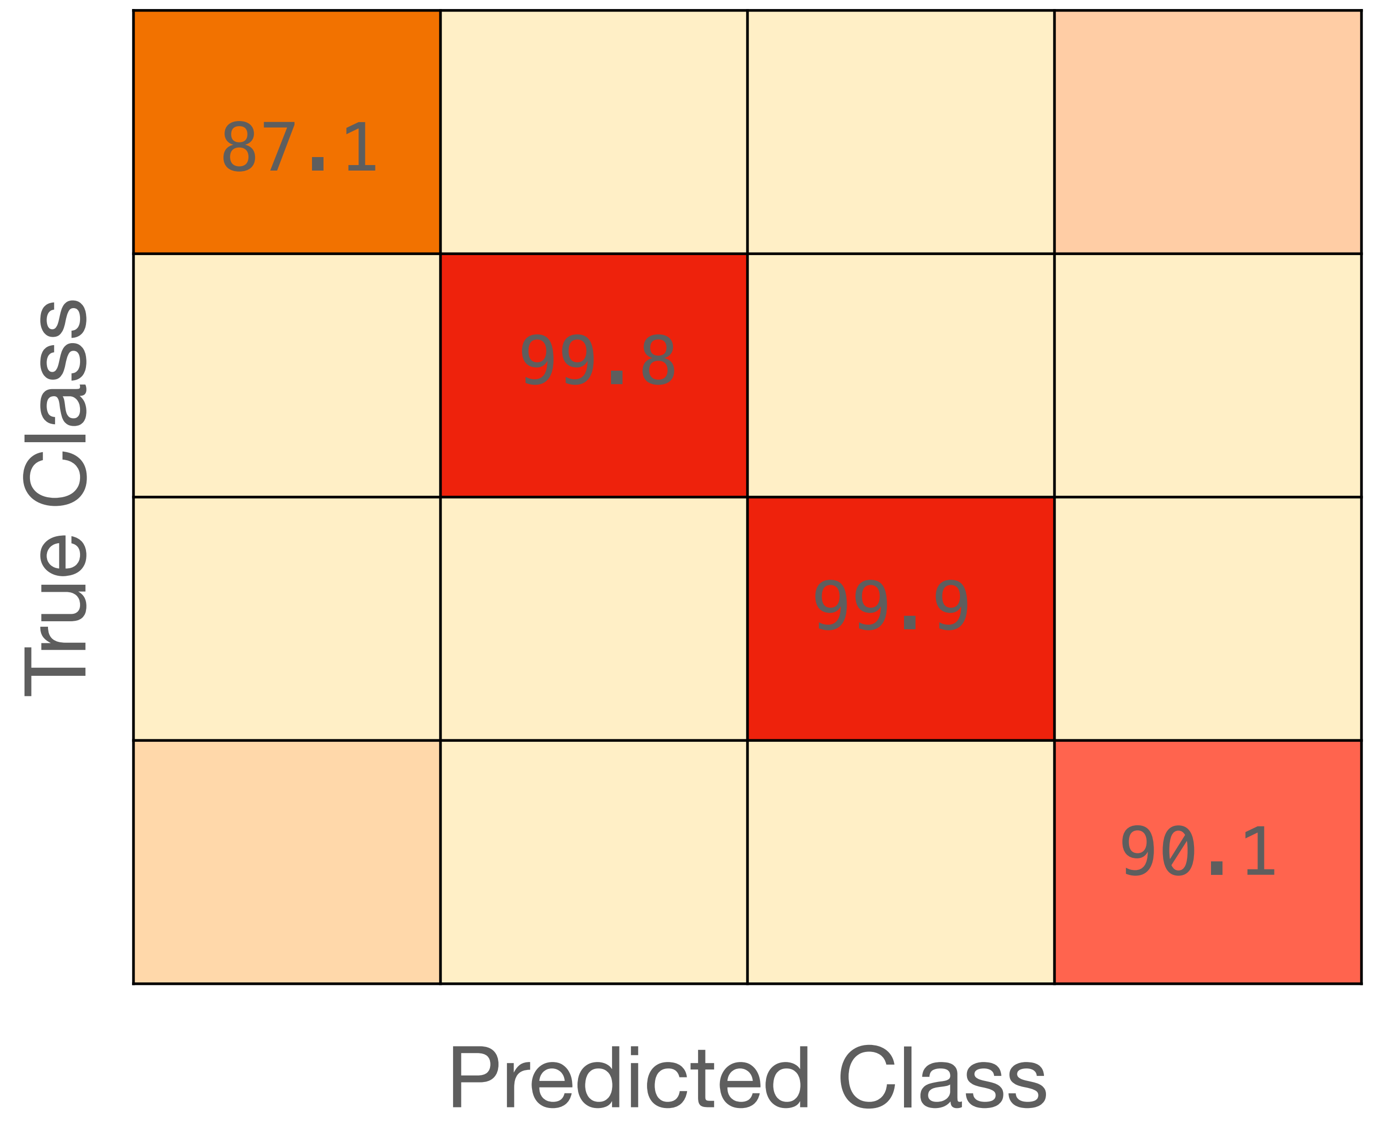

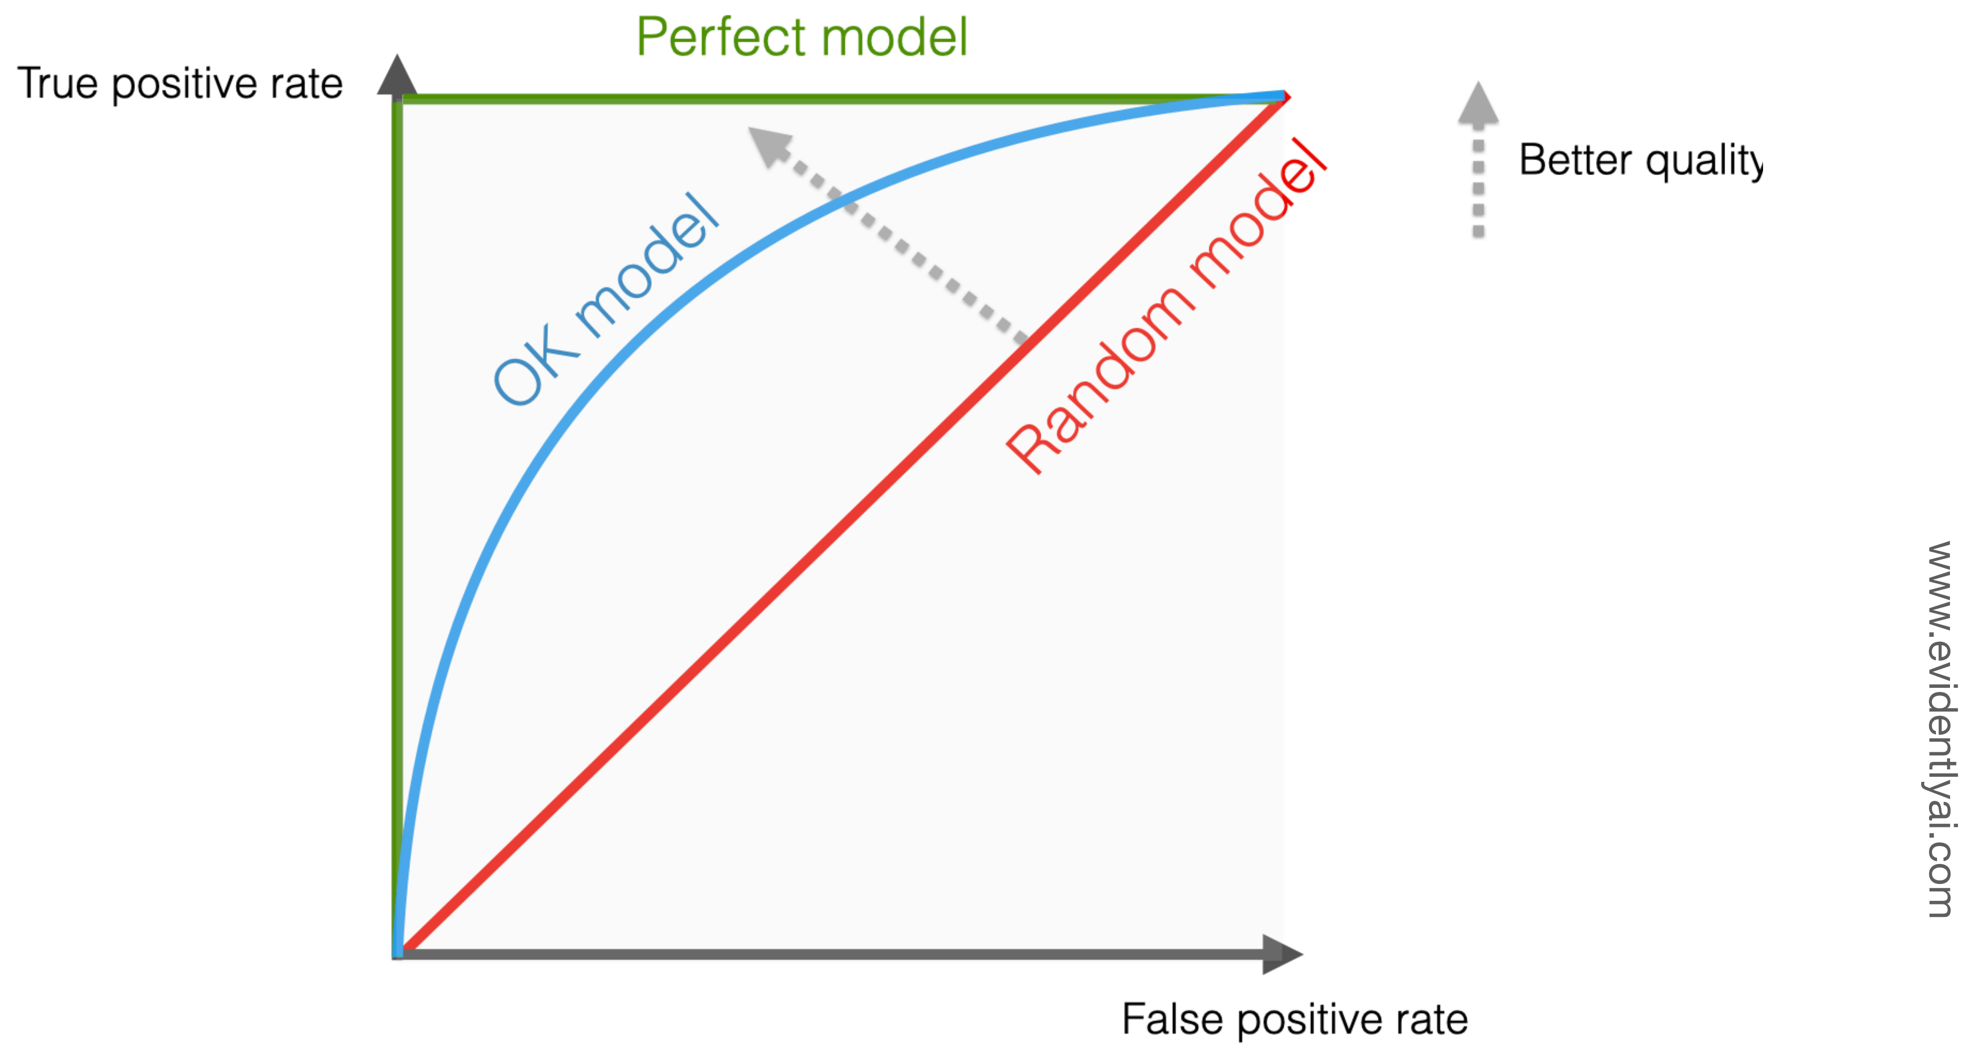

A. Higuera, NPML 2025 University of Tokyo

Classification

Classification

- Neutrino flavor
- Particle ID
- Hit-level ID

How robust is the method?

Robustness typically refers to a model's ability to maintain performance under various forms of perturbation or distribution shifts

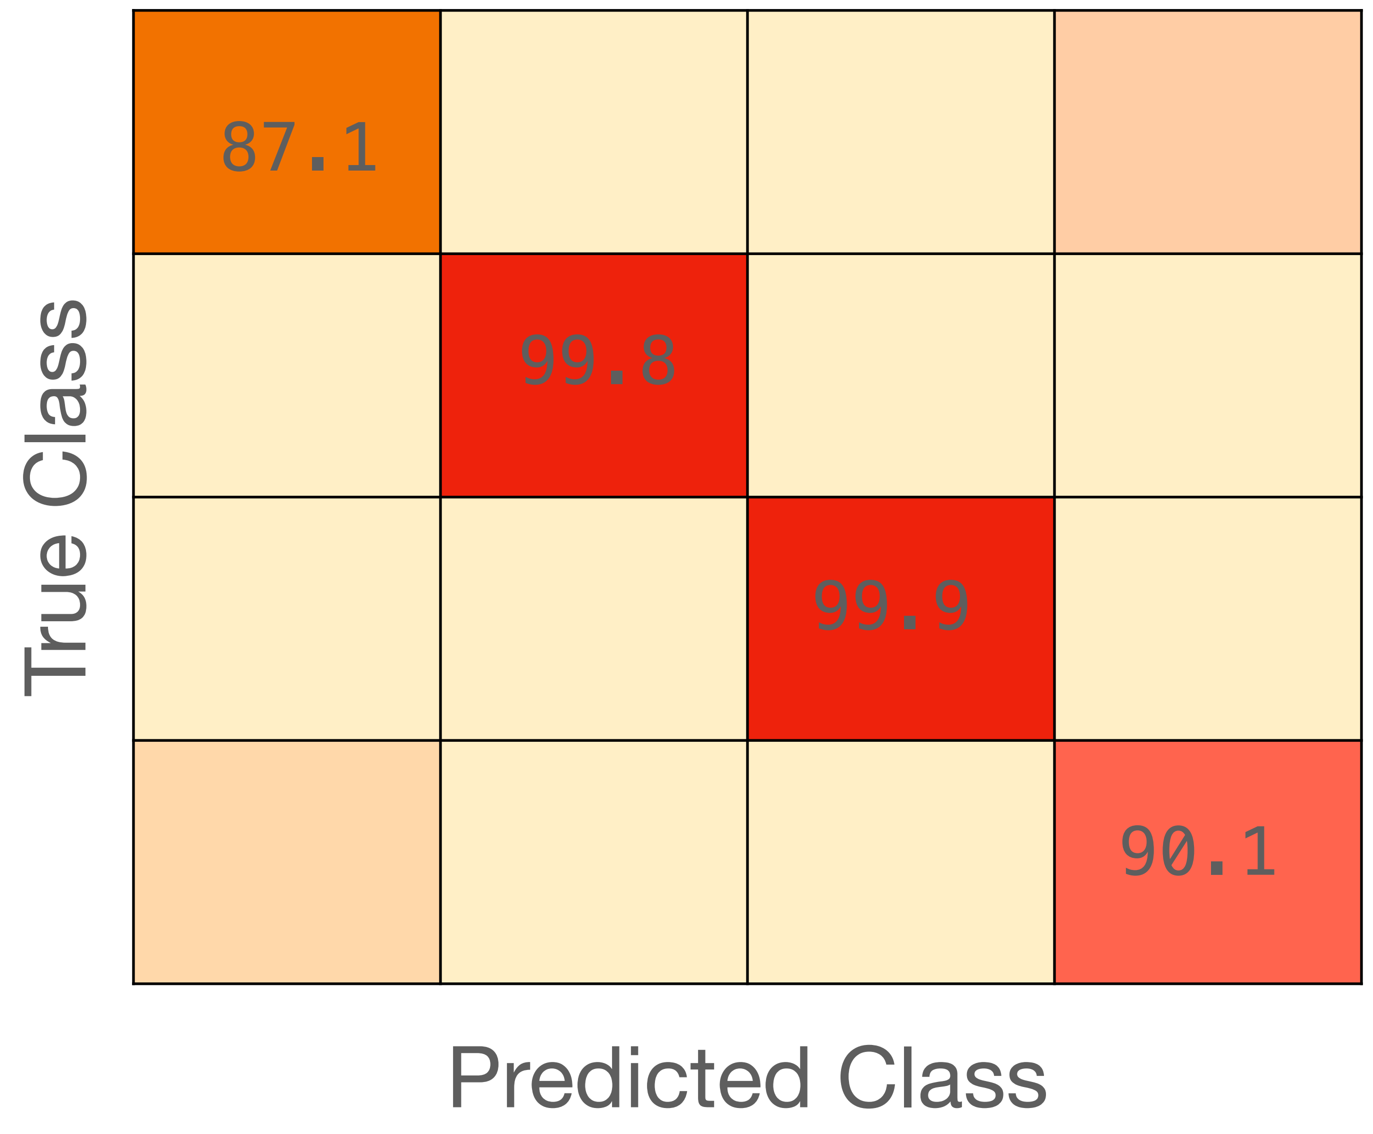

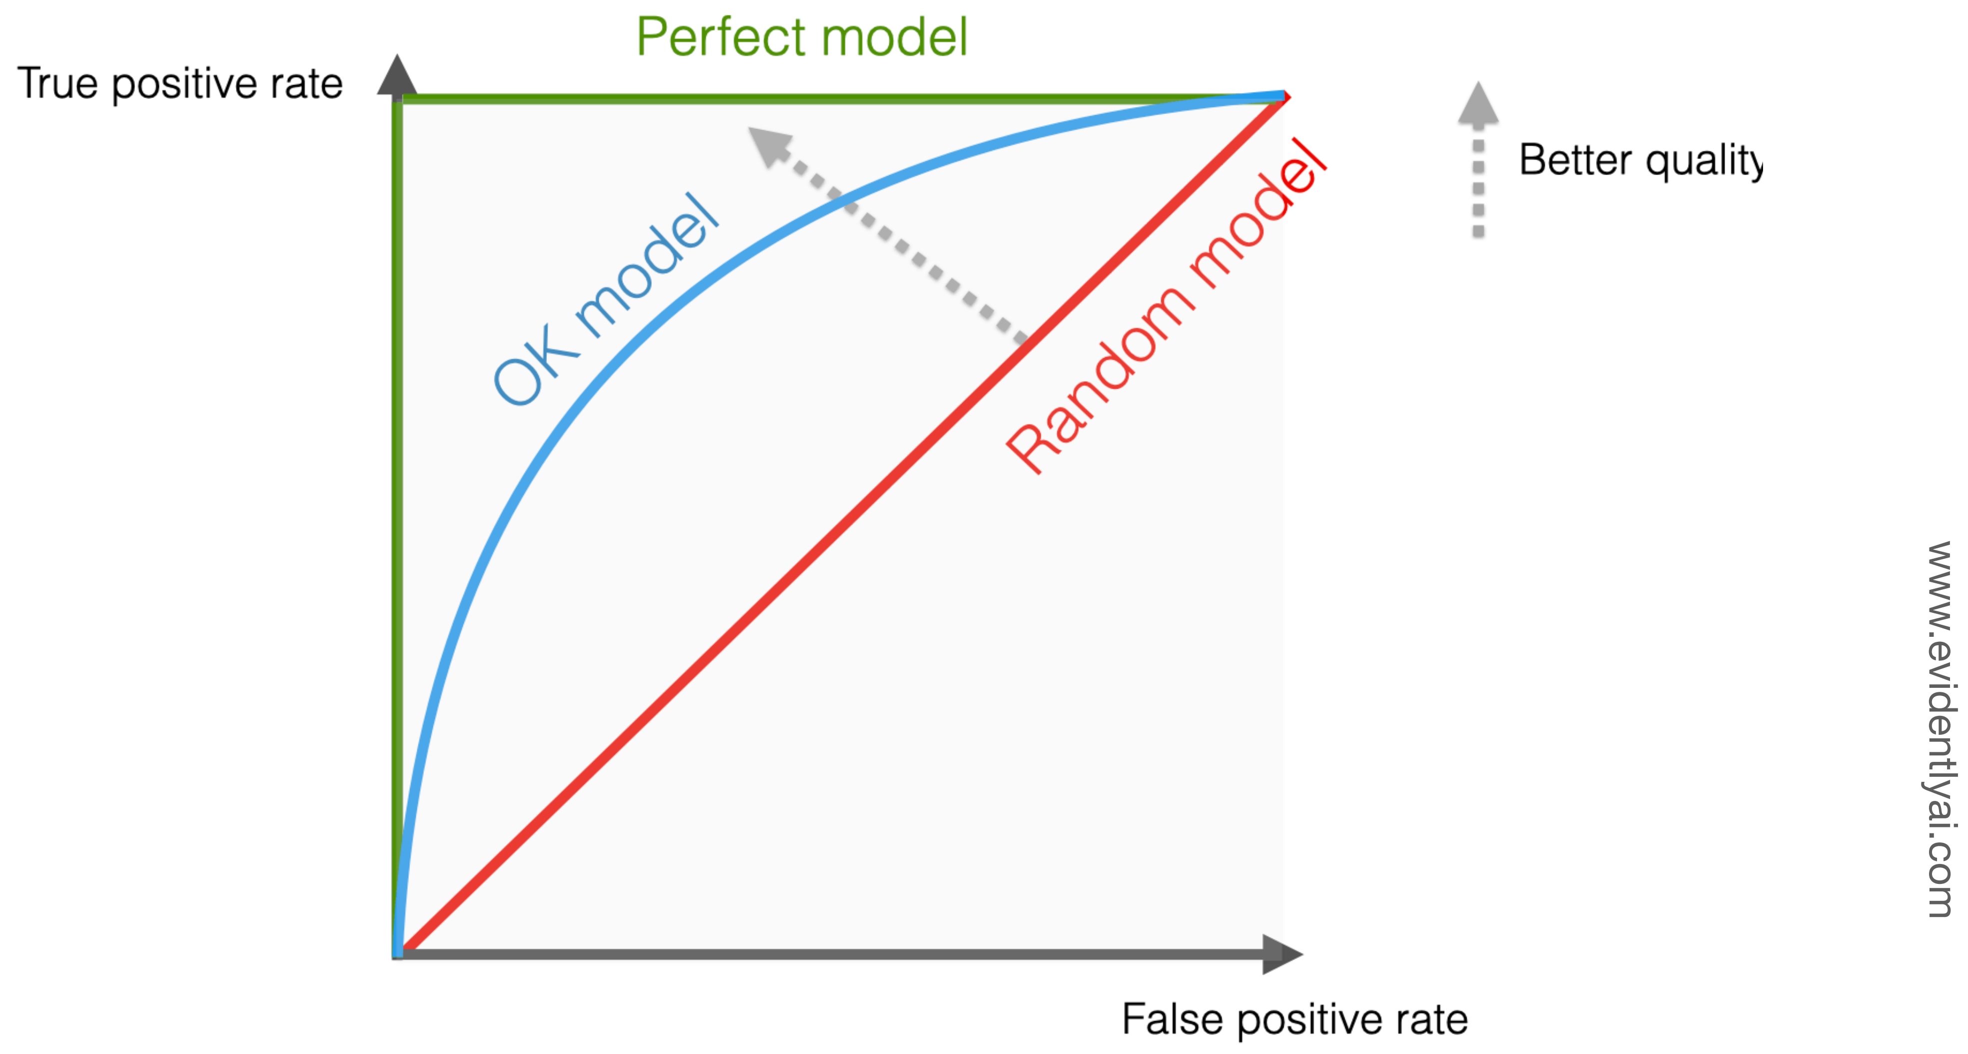

Classification

Classification

- Neutrino flavor
- Particle ID
- Hit-level ID

How robust is the method?

Robustness typically refers to a model's ability to maintain performance under various forms of perturbation or distribution shifts

Confidence measures how certain a model is about its predictions. Does the model know what it knows?

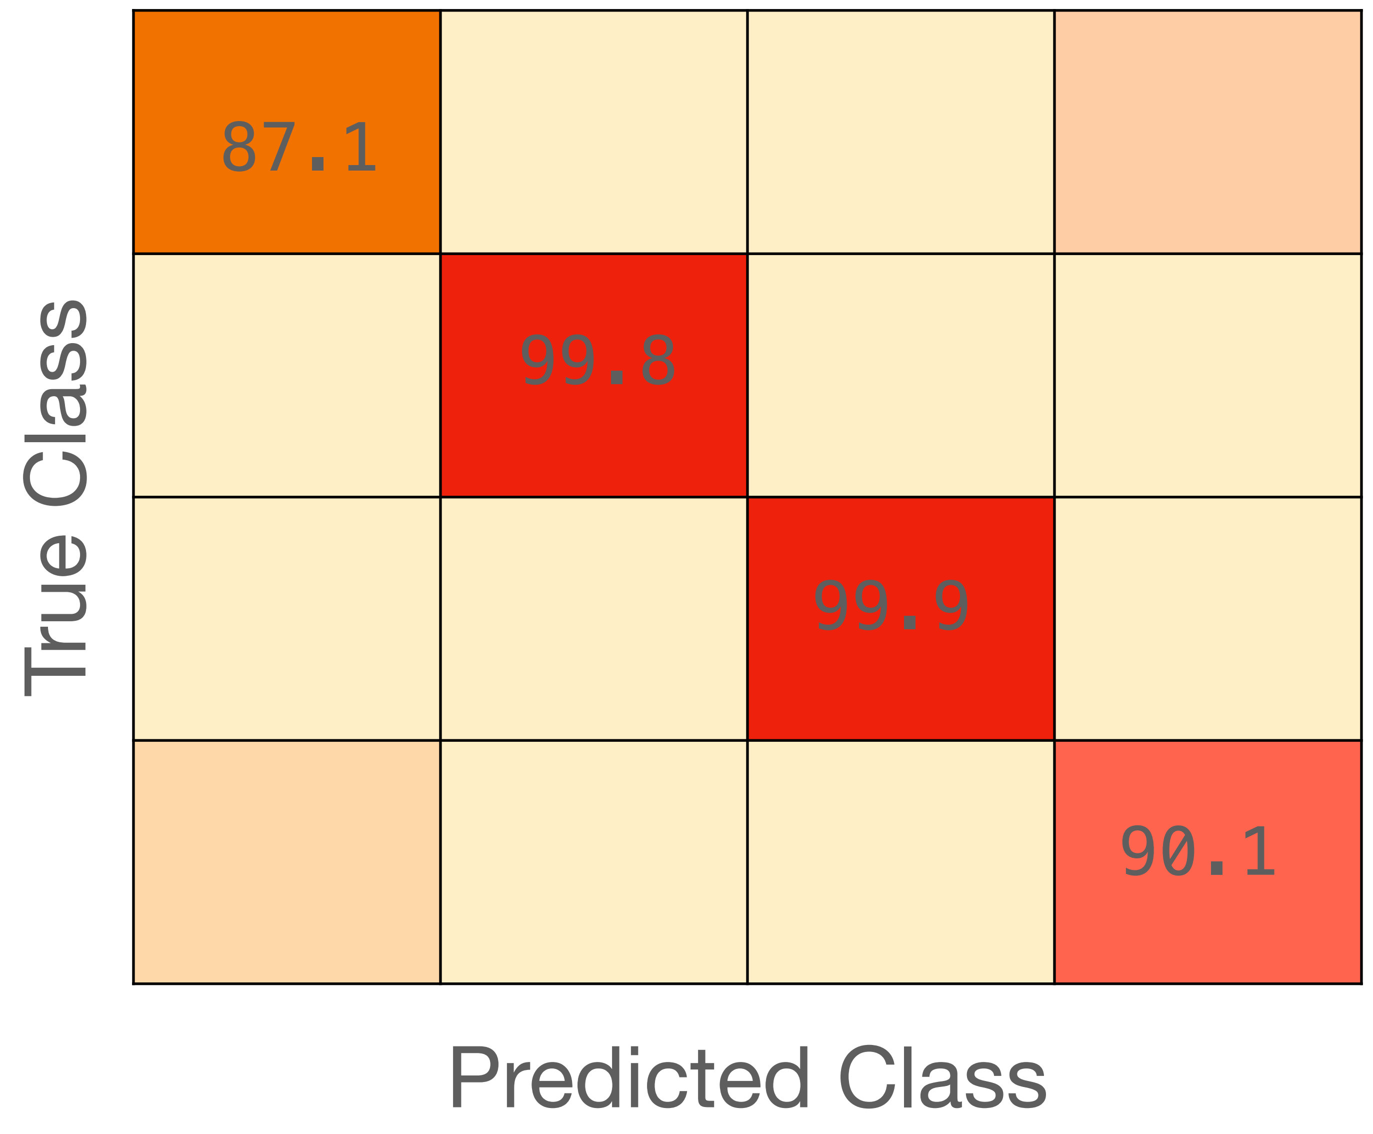

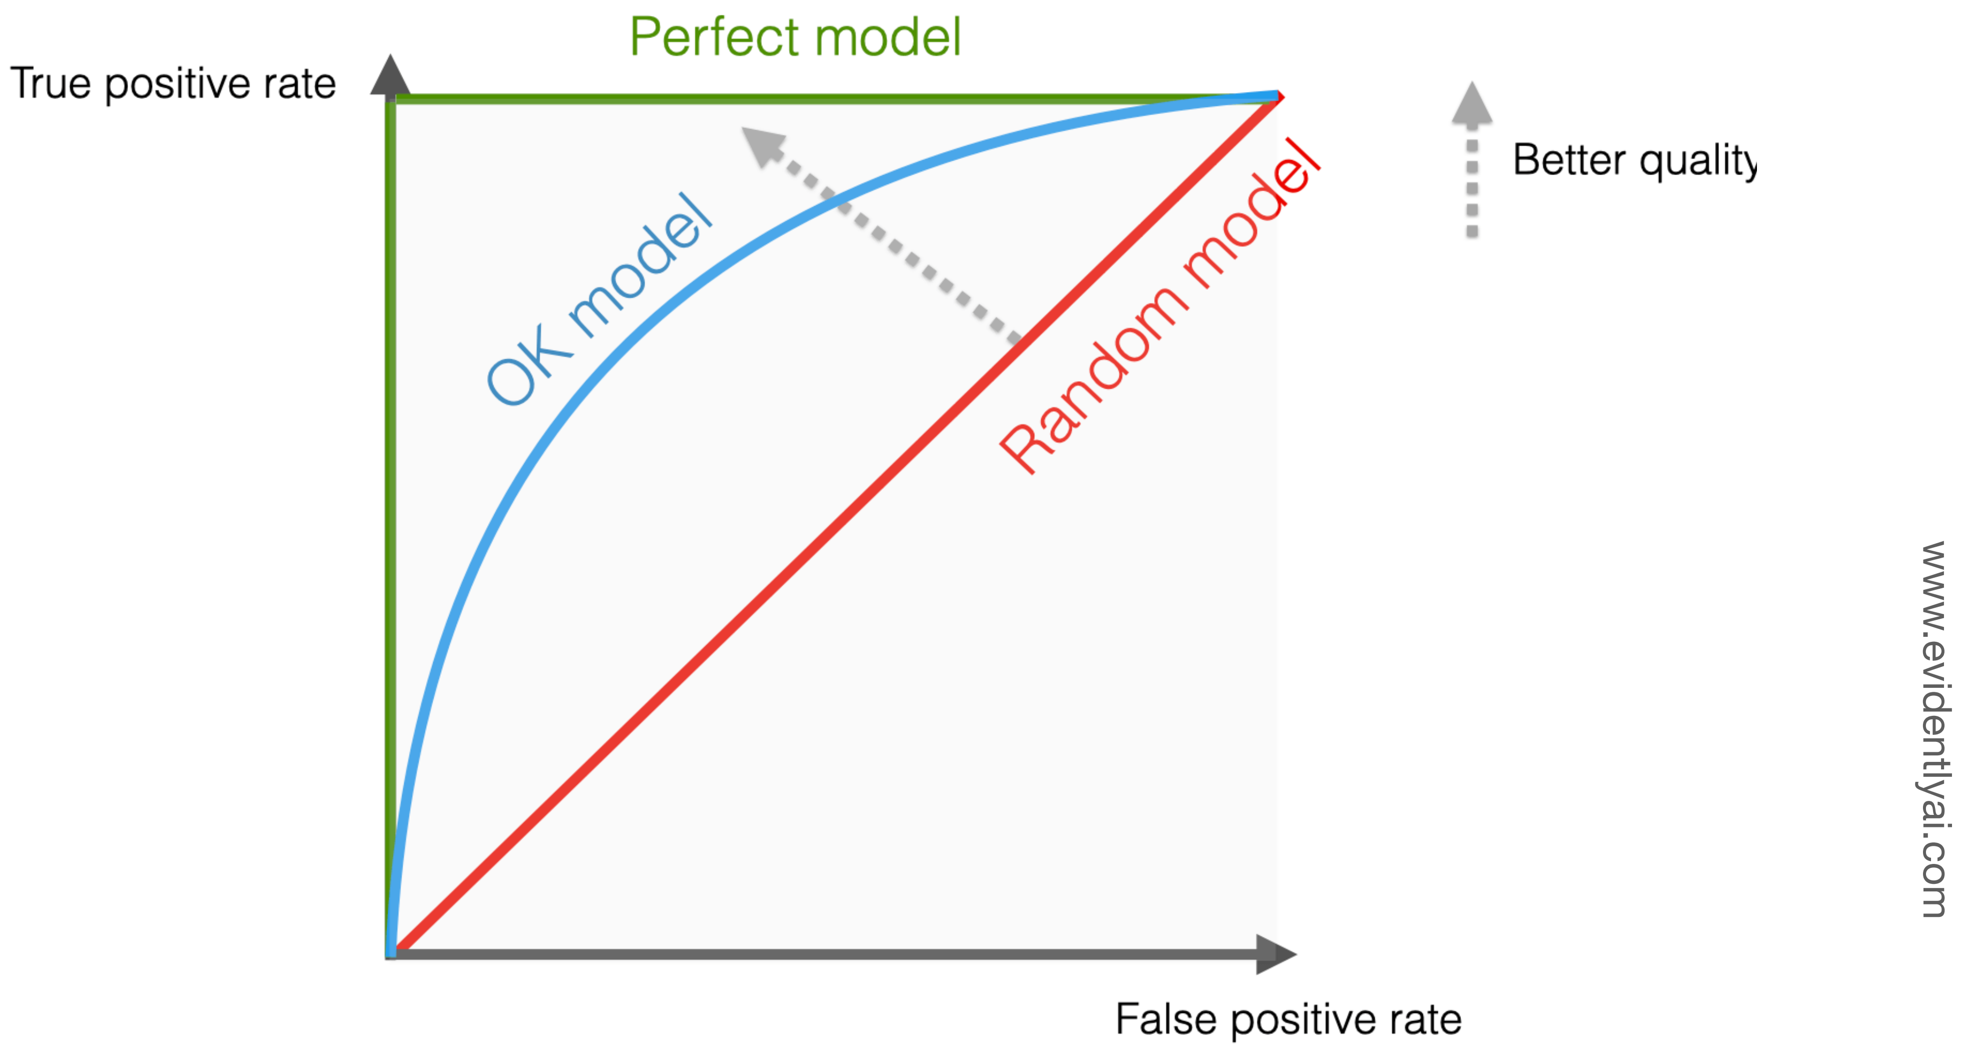

A. Higuera, NPML 2025 University of Tokyo

Classification

How robust is the method?

Robustness typically refers to a model's ability to maintain performance under various forms of perturbation or distribution shifts

Confidence measures how certain a model is about its predictions. Does the model know what it knows?

How confident should we be about this prediction, and how do we measure that confidence? = **Uncertainty quantification** (UQ)

UQ is formal framework for measuring and expressing confidence

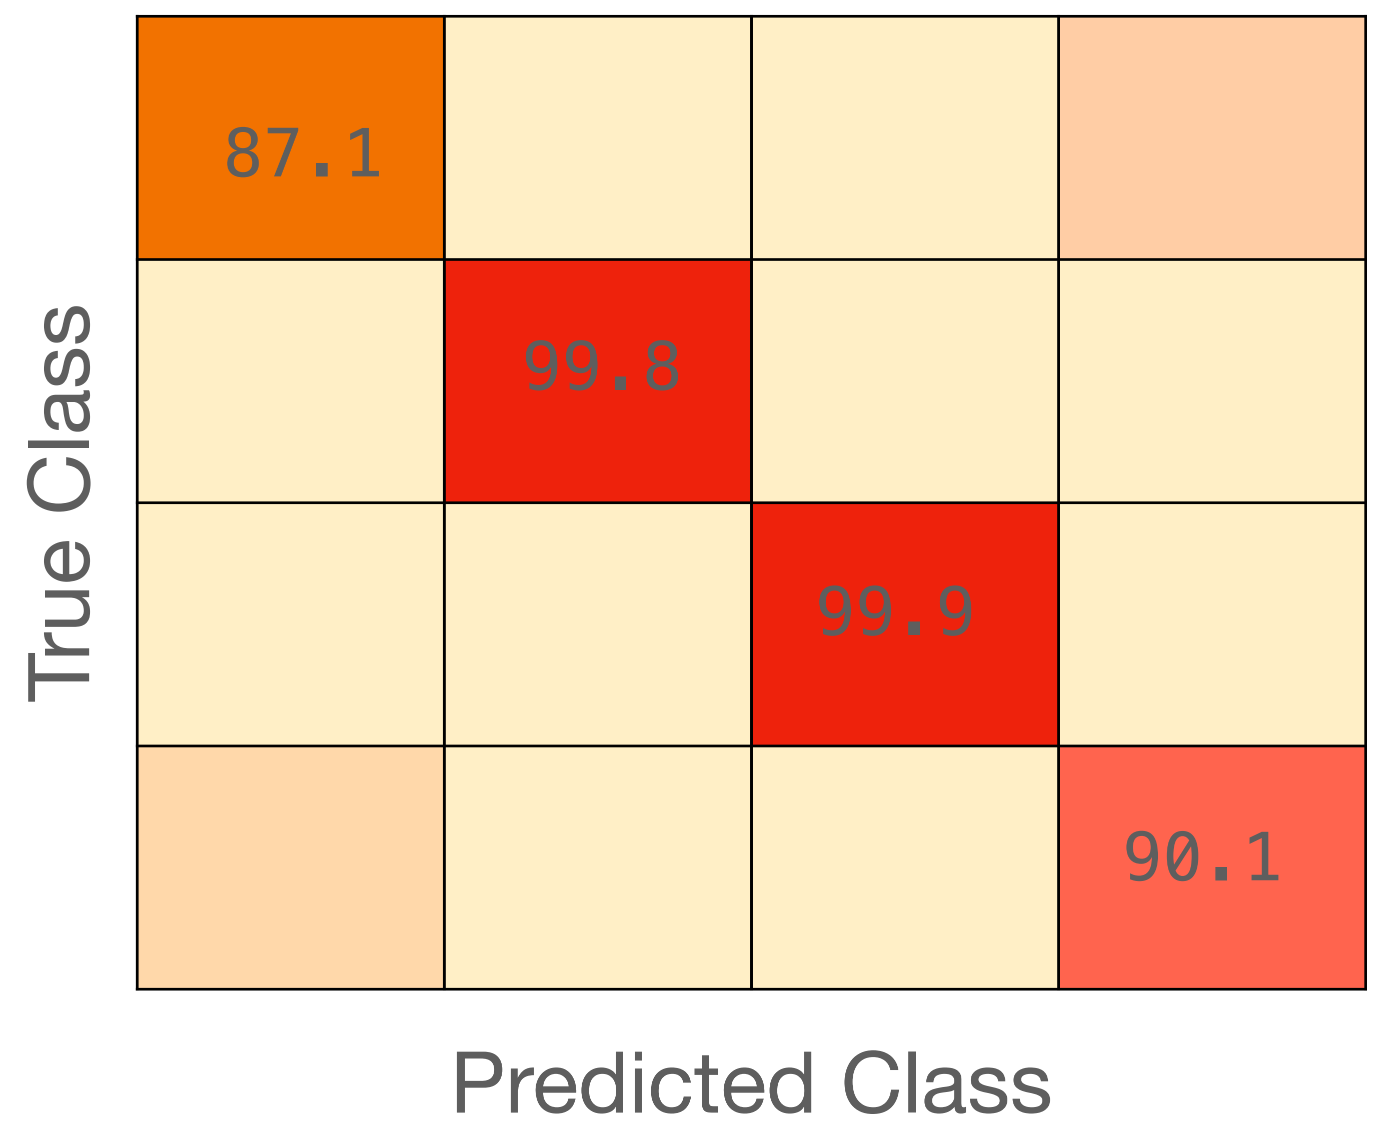

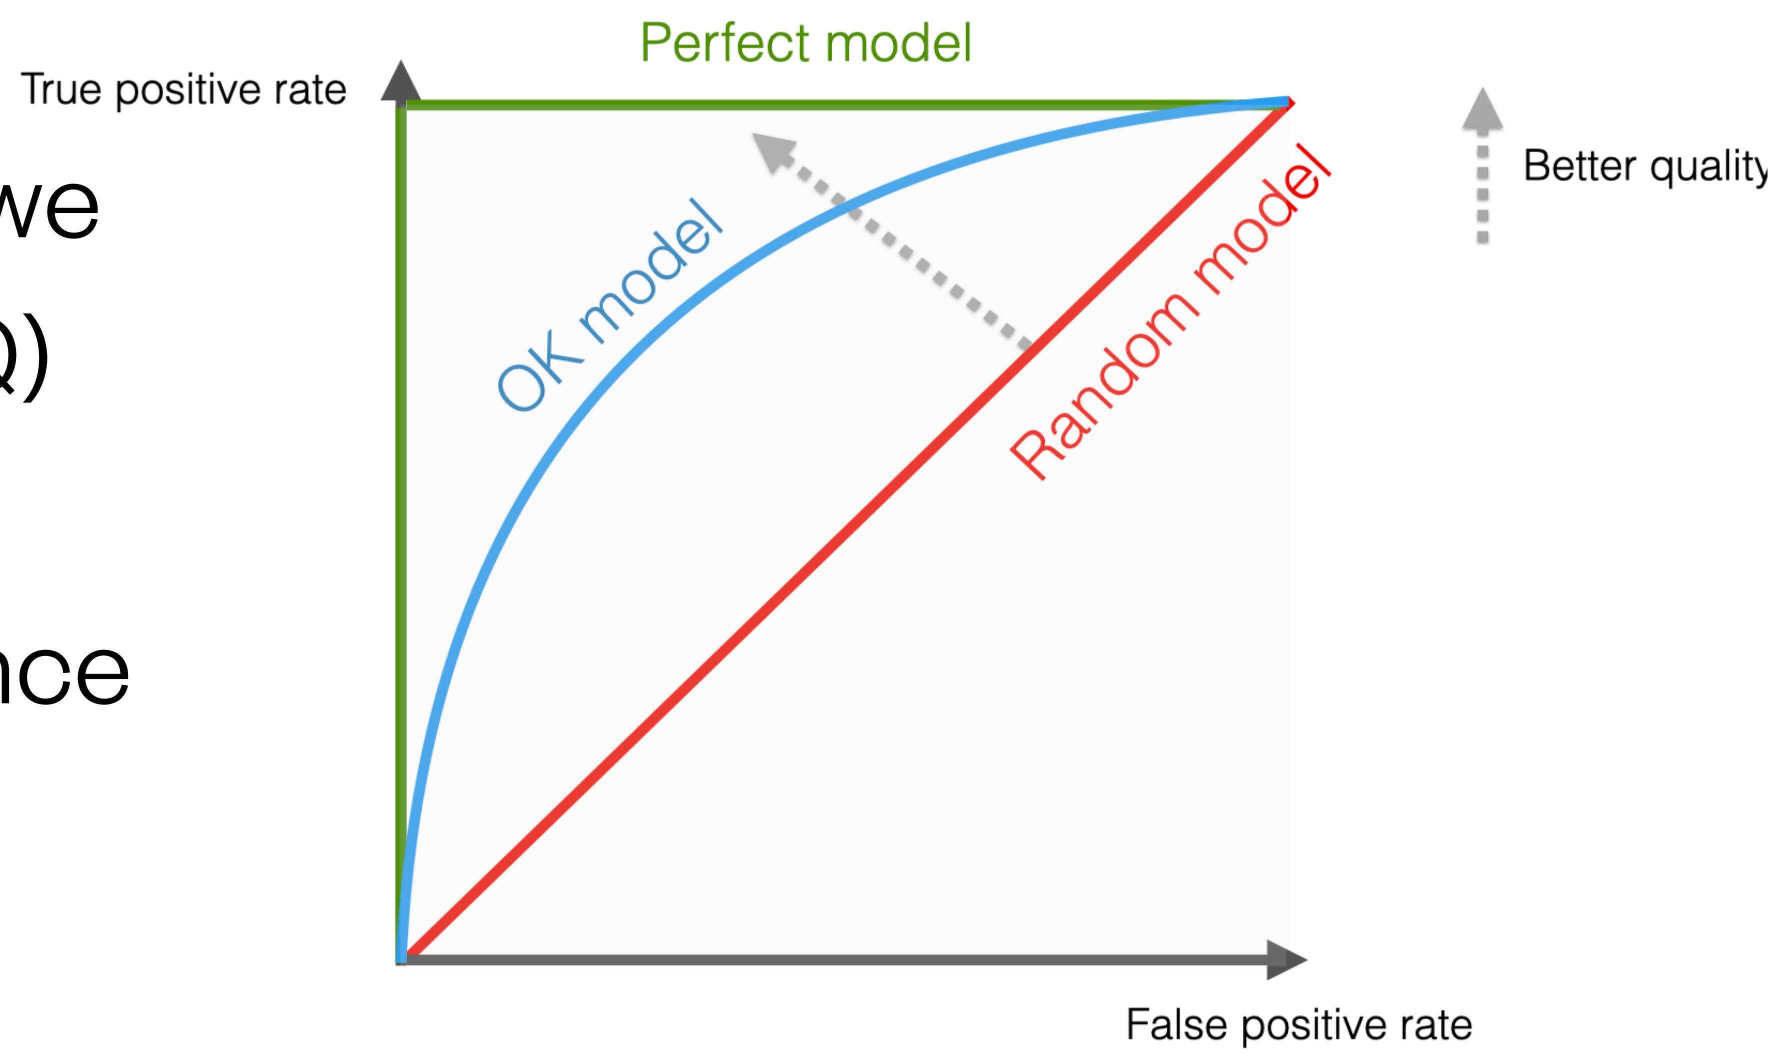

UQ in ML

Why UQ is often overlooked in the HEP domain when using ML

Metrics-driven focus:

Most ML efforts emphasize accuracy and precision

Computational overhead:

Even when using approximate methods (MC dropout, variational inference), they can multiply computing training and inference costs

Uncertainties

Aleatoric: irreducible uncertainty or stochastic uncertainty (data noise, irreducible randomness)

Epistemic: Model/parameter uncertainty (weights errors)

UQ in ML

Why UQ is often overlooked in the HEP domain when using ML

Metrics-driven focus:

Most ML efforts emphasize accuracy, precision, and recall

Computational overhead:

Even when using approximate methods (MC dropout, variational inference), they can multiply computing training and inference costs

Deep neural network uncertainty quantification for LArTPC reconstruction (D. Koh et al 2023 JINST 18 P12013)

UQ in ML

Deep neural network uncertainty quantification for LArTPC reconstruction (D. Koh et al 2023 JINST 18 P12013)

- Task:
  - Single particle classification, semantic segmentation and multi-particle classification
- Methods:
  - Ensemble Methods
  - Monte Carlo Dropout (MCD)
  - Evidential Deep Learning (EDL)

UQ in ML

Deep neural network uncertainty quantification for LArTPC reconstruction (D. Koh et al 2023 JINST 18 P12013)

- Task:
  - Single particle classification, semantic segmentation and multi-particle classification
- Methods:
  - **Ensemble Methods**
    - Train multiple identical networks with different initializations
    - Average predictions across ensemble members
    - Simple but computationally expensive
  - Monte Carlo Dropout (MCD)
  - Evidential Deep Learning (EDL)

UQ in ML

Deep neural network uncertainty quantification for LArTPC reconstruction (D. Koh et al 2023 JINST 18 P12013)

- Task:
  - Single particle classification, semantic segmentation and multi-particle classification
- Methods:
  - Ensemble Methods
  - **Monte Carlo Dropout (MCD)**
    - Use dropout during inference to create stochastic predictions
    - Multiple forward passes with different neurons dropped
    - Moderate computational cost
  - Evidential Deep Learning (EDL)

UQ in ML

Deep neural network uncertainty quantification for LArTPC reconstruction (D. Koh et al 2023 JINST 18 P12013)

- Task:
  - Single particle classification, semantic segmentation and multi-particle classification
- Methods:
  - Ensemble Methods
  - Monte Carlo Dropout (MCD)
  - **Evidential Deep Learning (EDL)**
    - Models predictions using Dirichlet distributions
    - Provides full posterior in single forward pass
    - Three loss variants tested: MLL, Bayes Risk, Brier Score

UQ in ML

Deep neural network uncertainty quantification for LArTPC reconstruction ([D. Koh et al 2023 JINST 18 P12013])

- Task:
  - Single particle classification, semantic segmentation and multi-particle classification
- Methods:
  - Ensemble Methods
  - Monte Carlo Dropout (MCD)
  - Evidential Deep Learning (EDL)
- Outputs:
  - **Predictive Entropy**: For each individual prediction, the model outputs a probability distribution over classes. The entropy of this distribution serves as an uncertainty measure for that specific event

Bayesian and Deterministic NNs

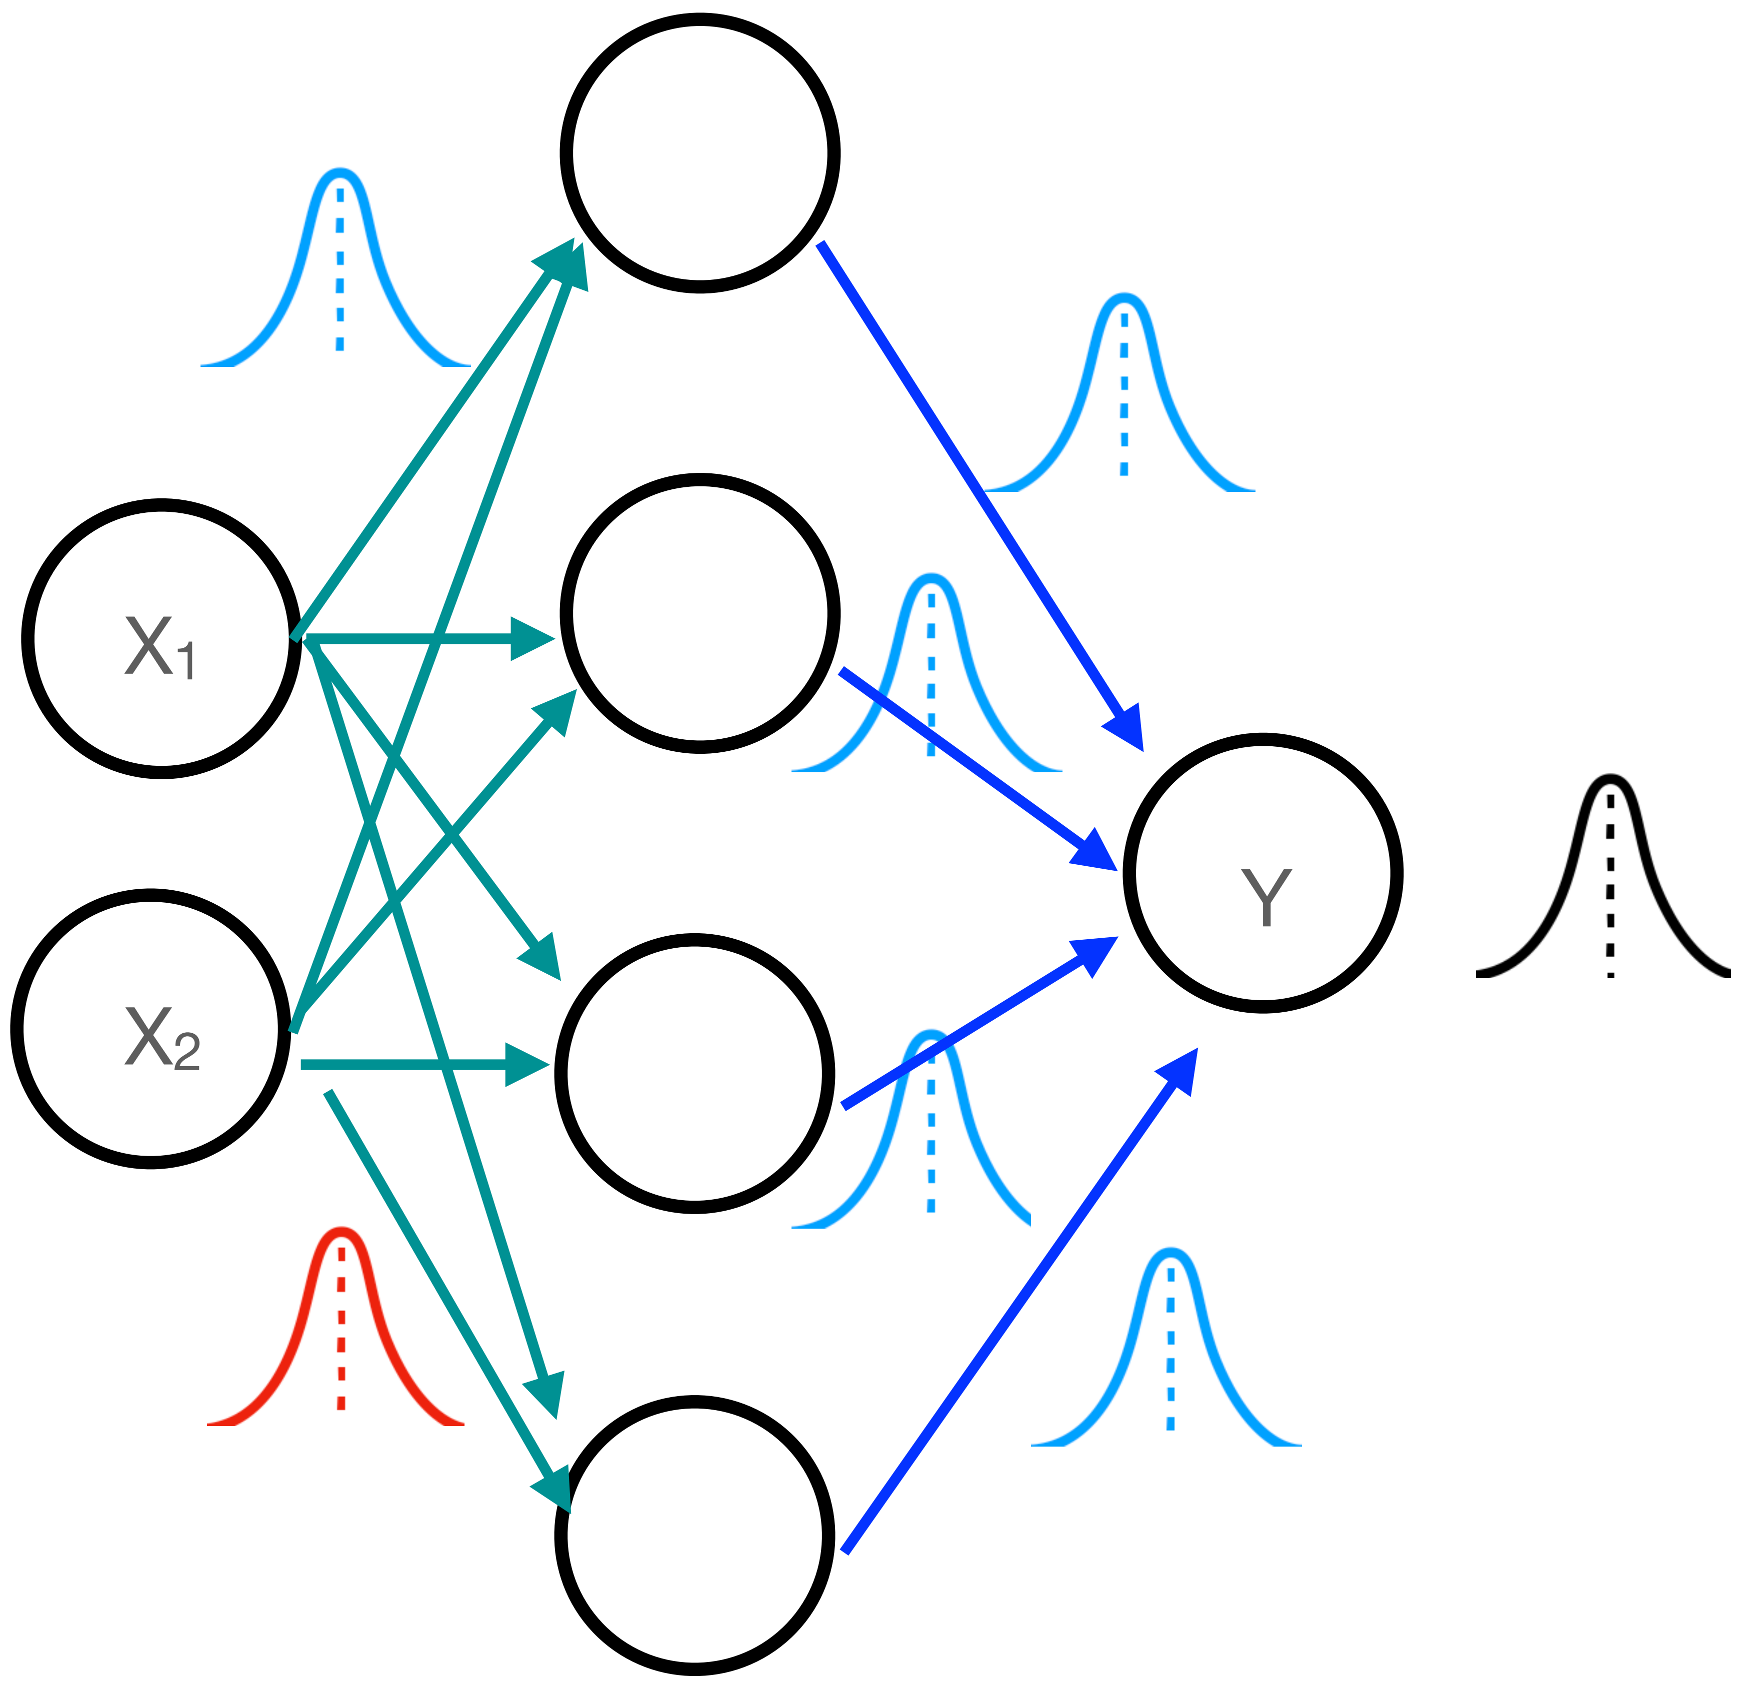

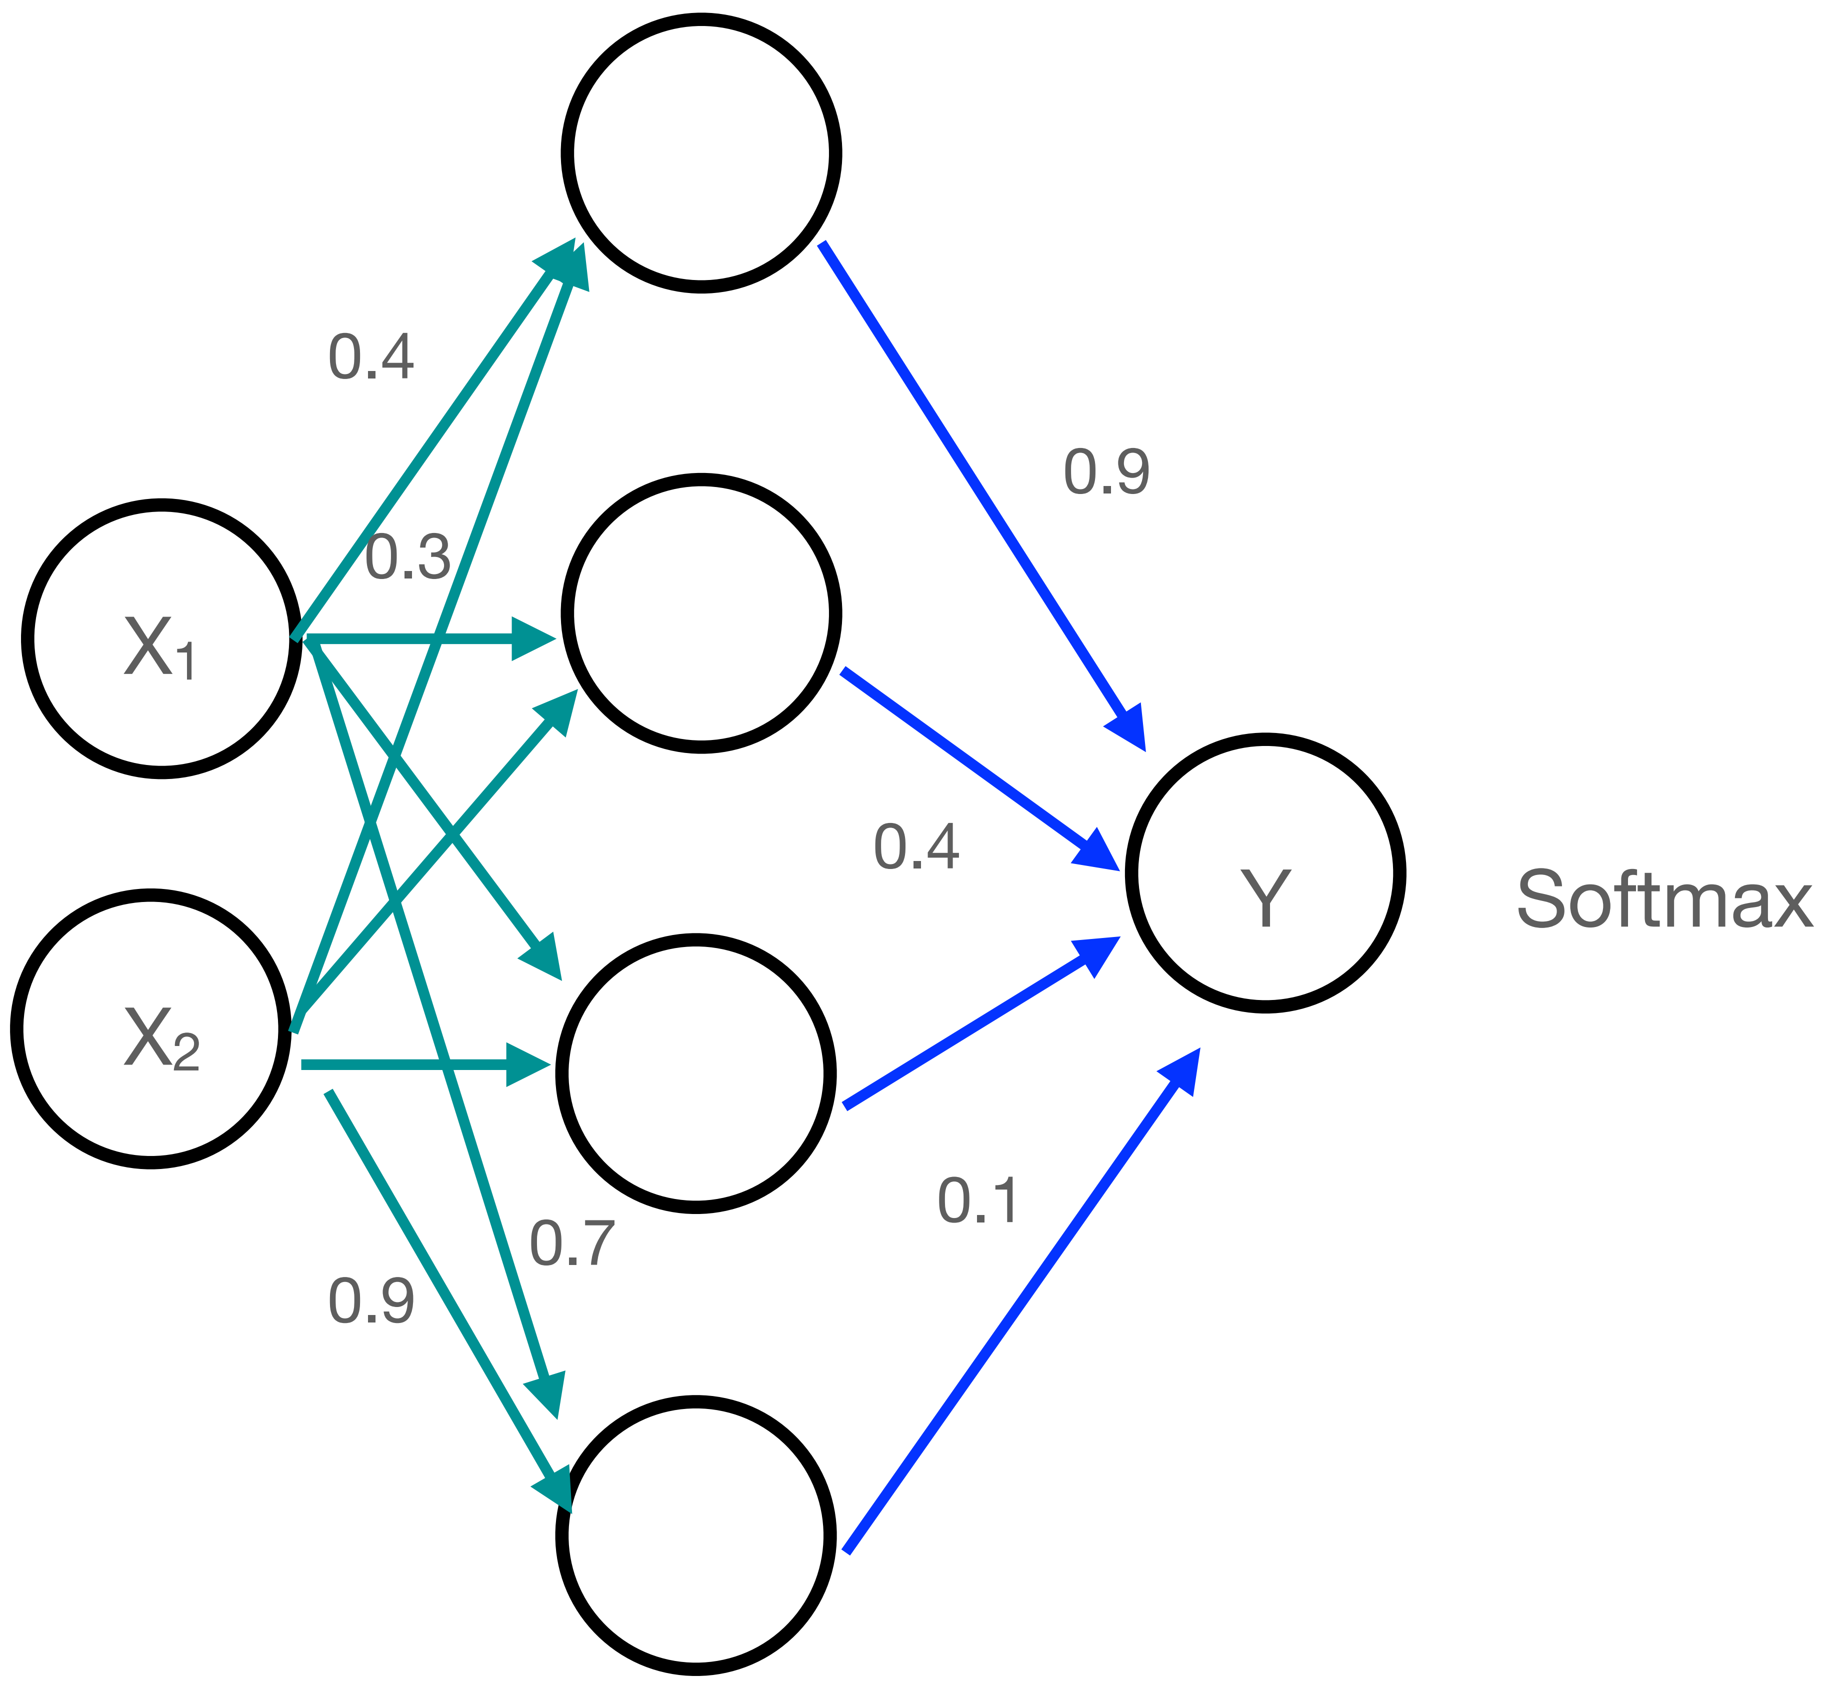

Bayesian NNs

	Ensamble Method	BNNs	MC Dropout
Training	Multi models	Single model w/ probabilistic weights	Single model
Inferences	Average over predictions from all models	Sample from weight posterior distributions	Multiple forward passes with dropout active
Advantage	Simple	Formal framework	Fast
Uncertainty Type	Epistemic	Epistemic+Aleatoric	Epistemic

$$nll = -Y_{pred} log(Y_{true})$$

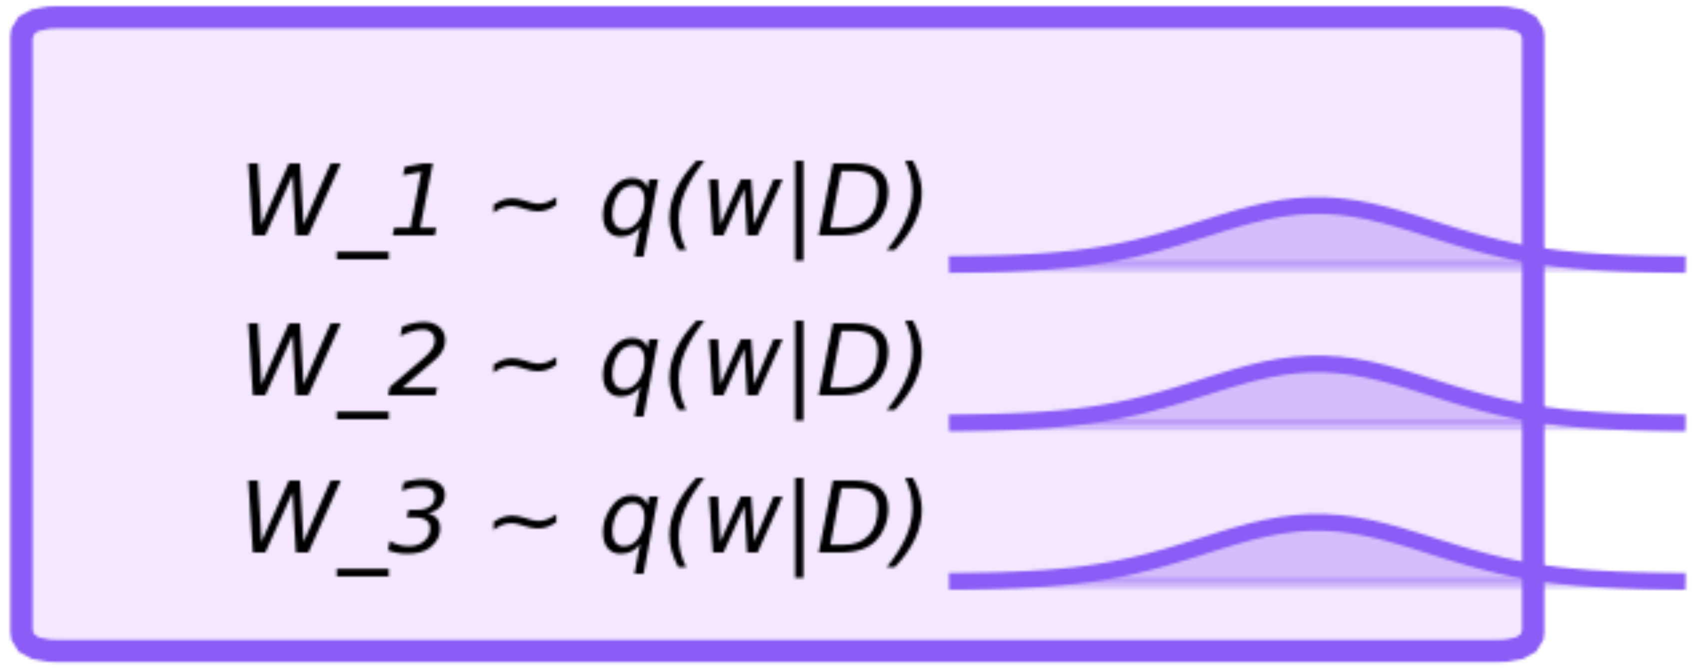

Single model with probabilistic weights

Example

Using MicroBooNE open datasets

Task: Neutrino flavor classification NC, CCν_e CC ν_μ

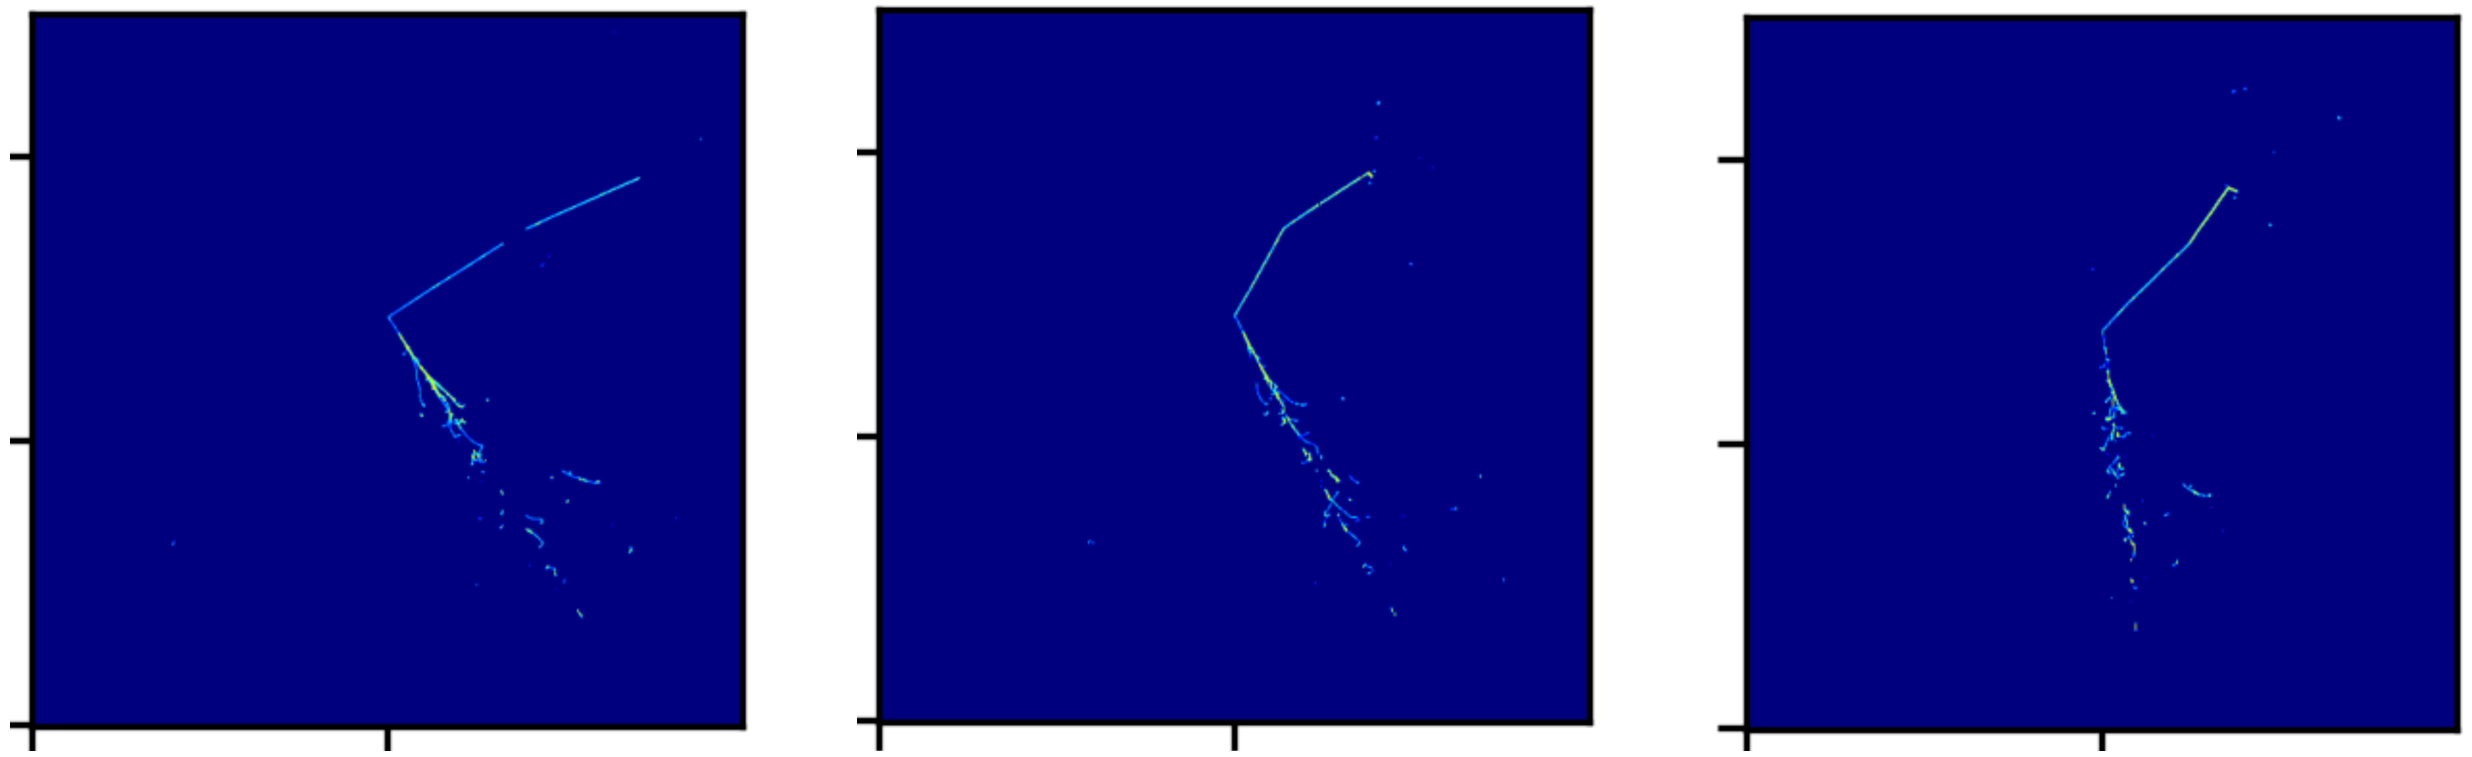

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

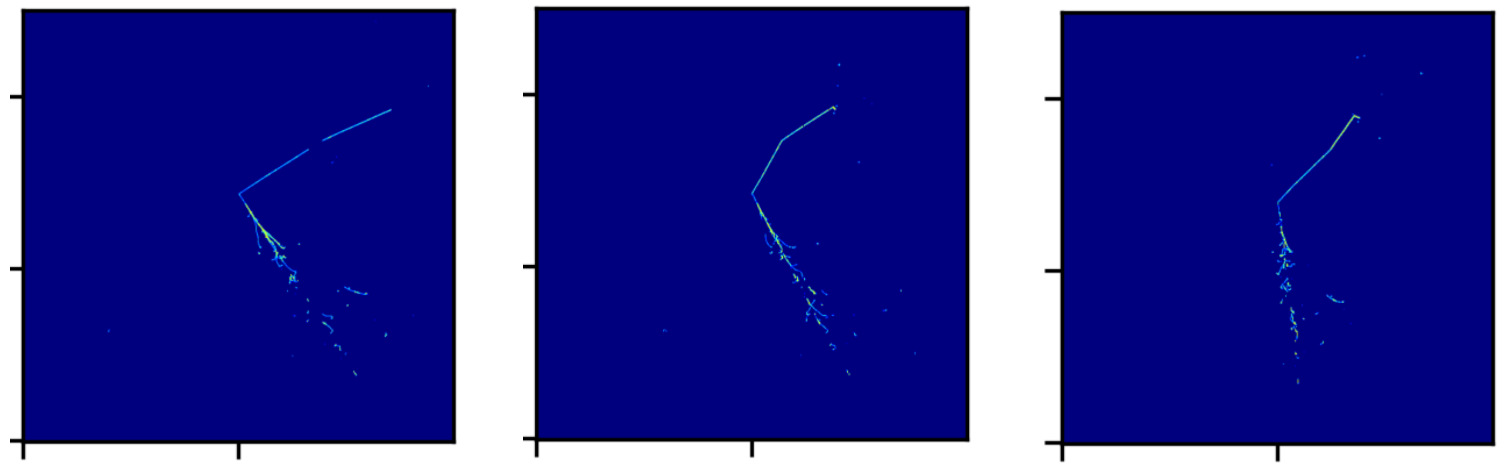

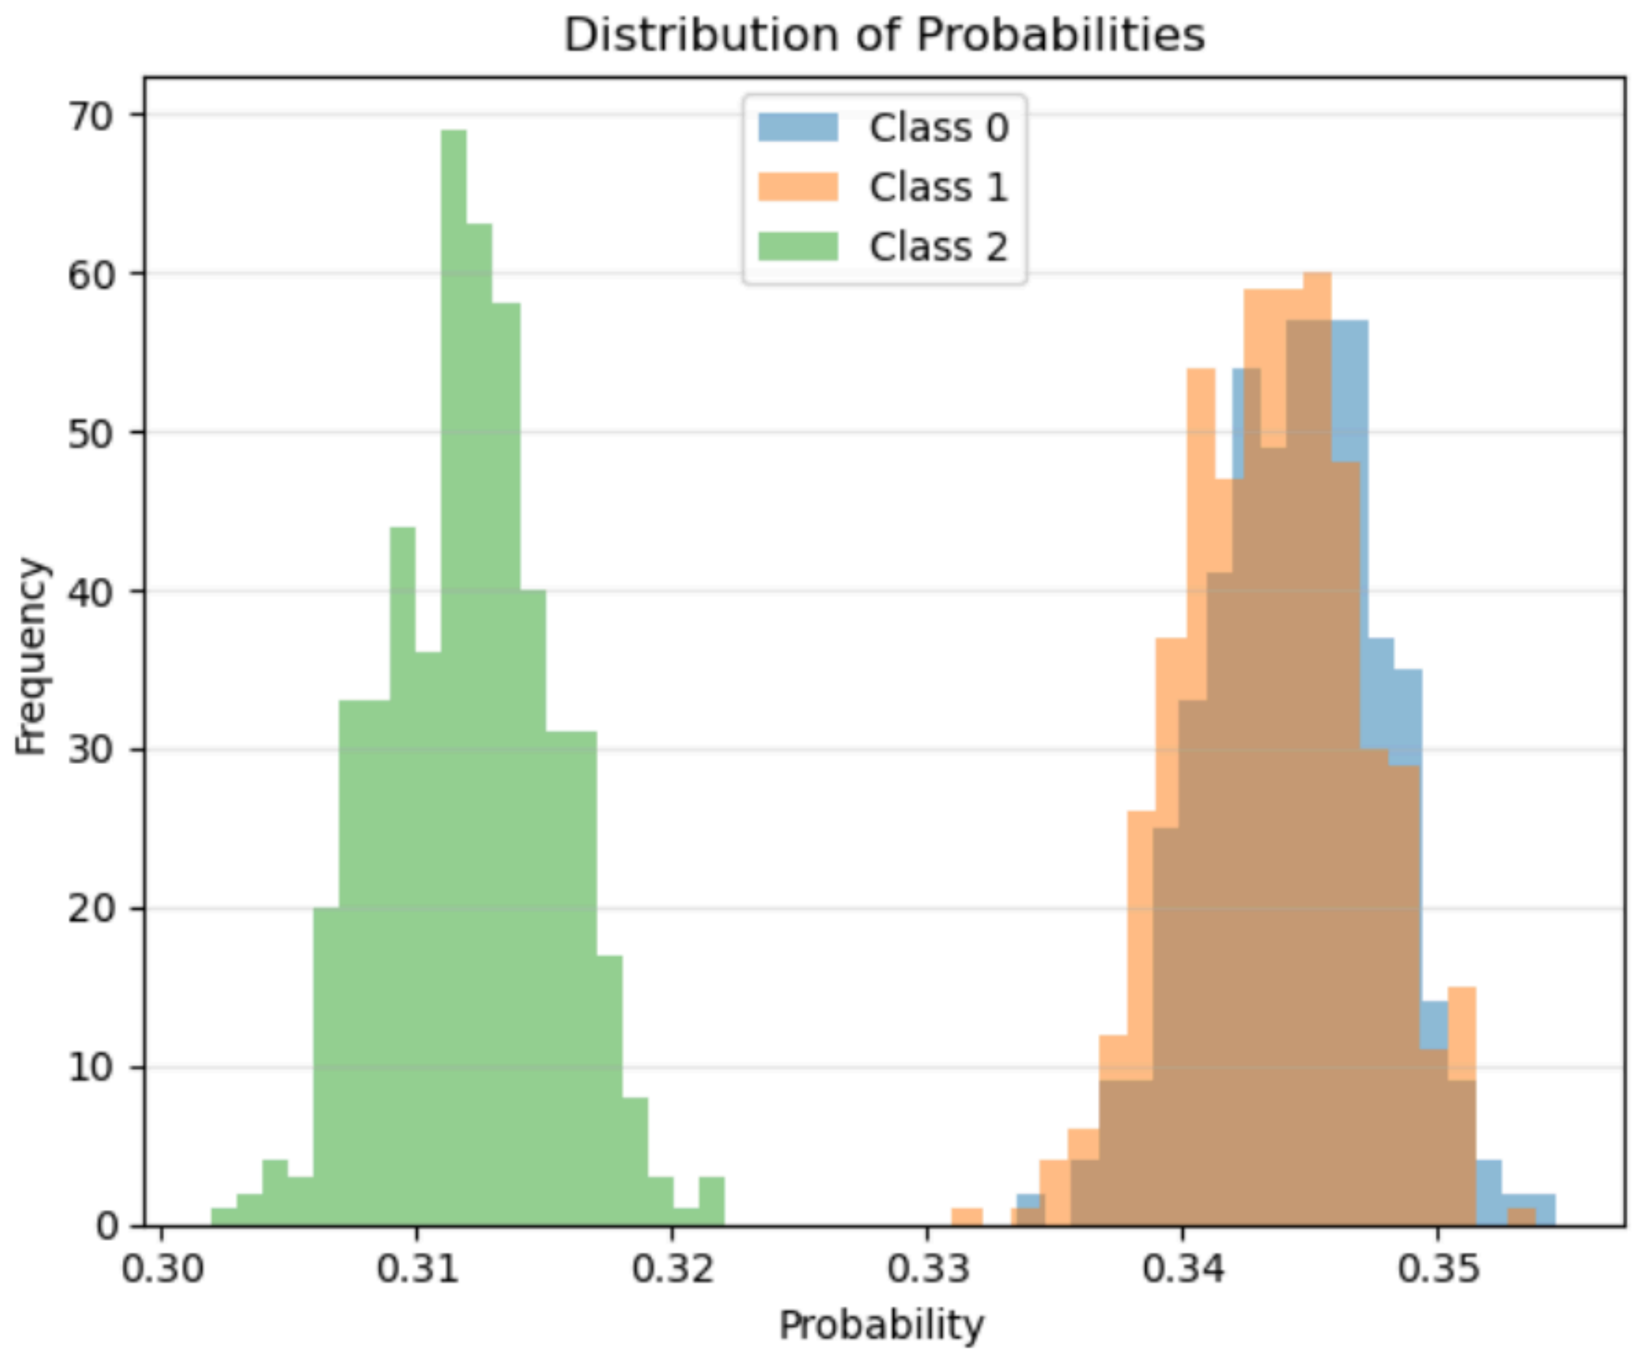

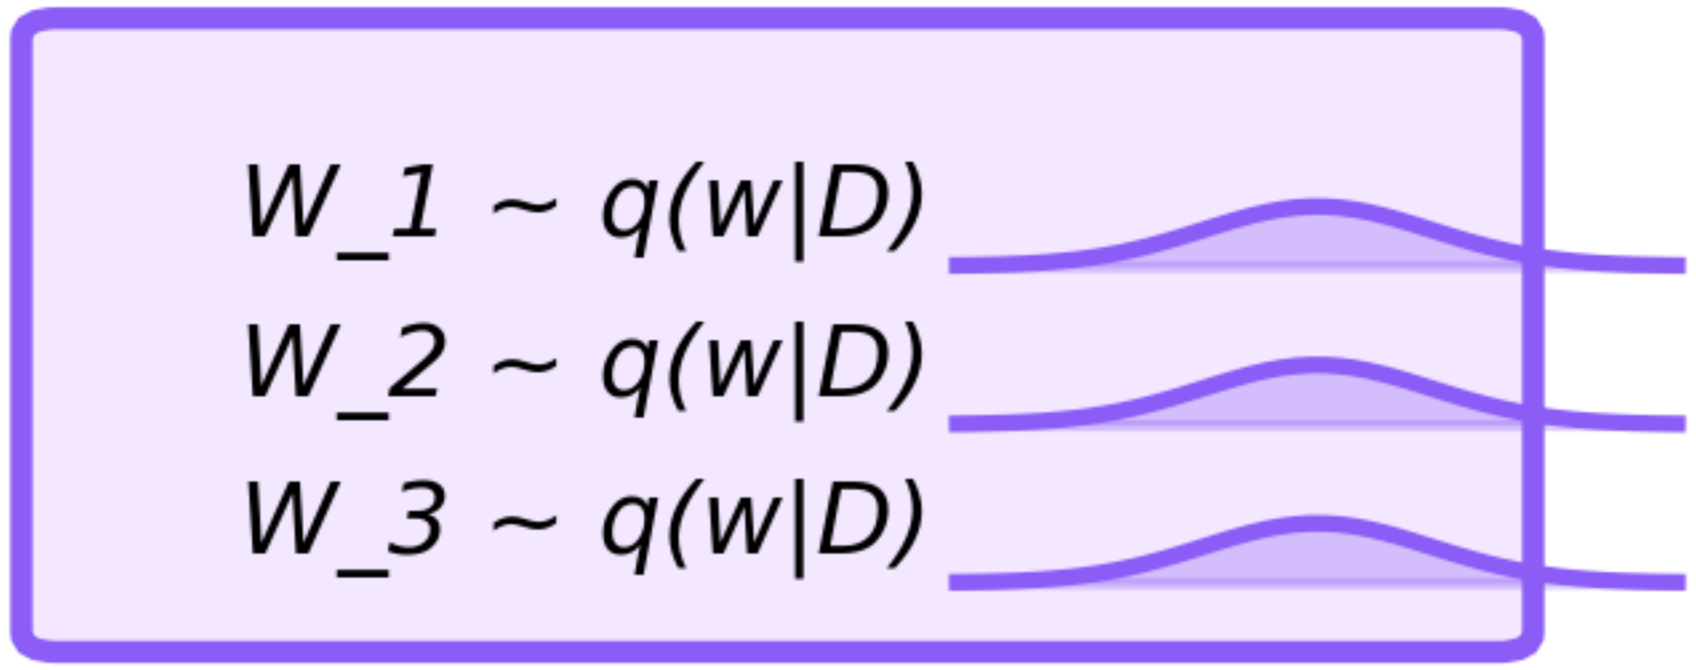

Non-informative Prior

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

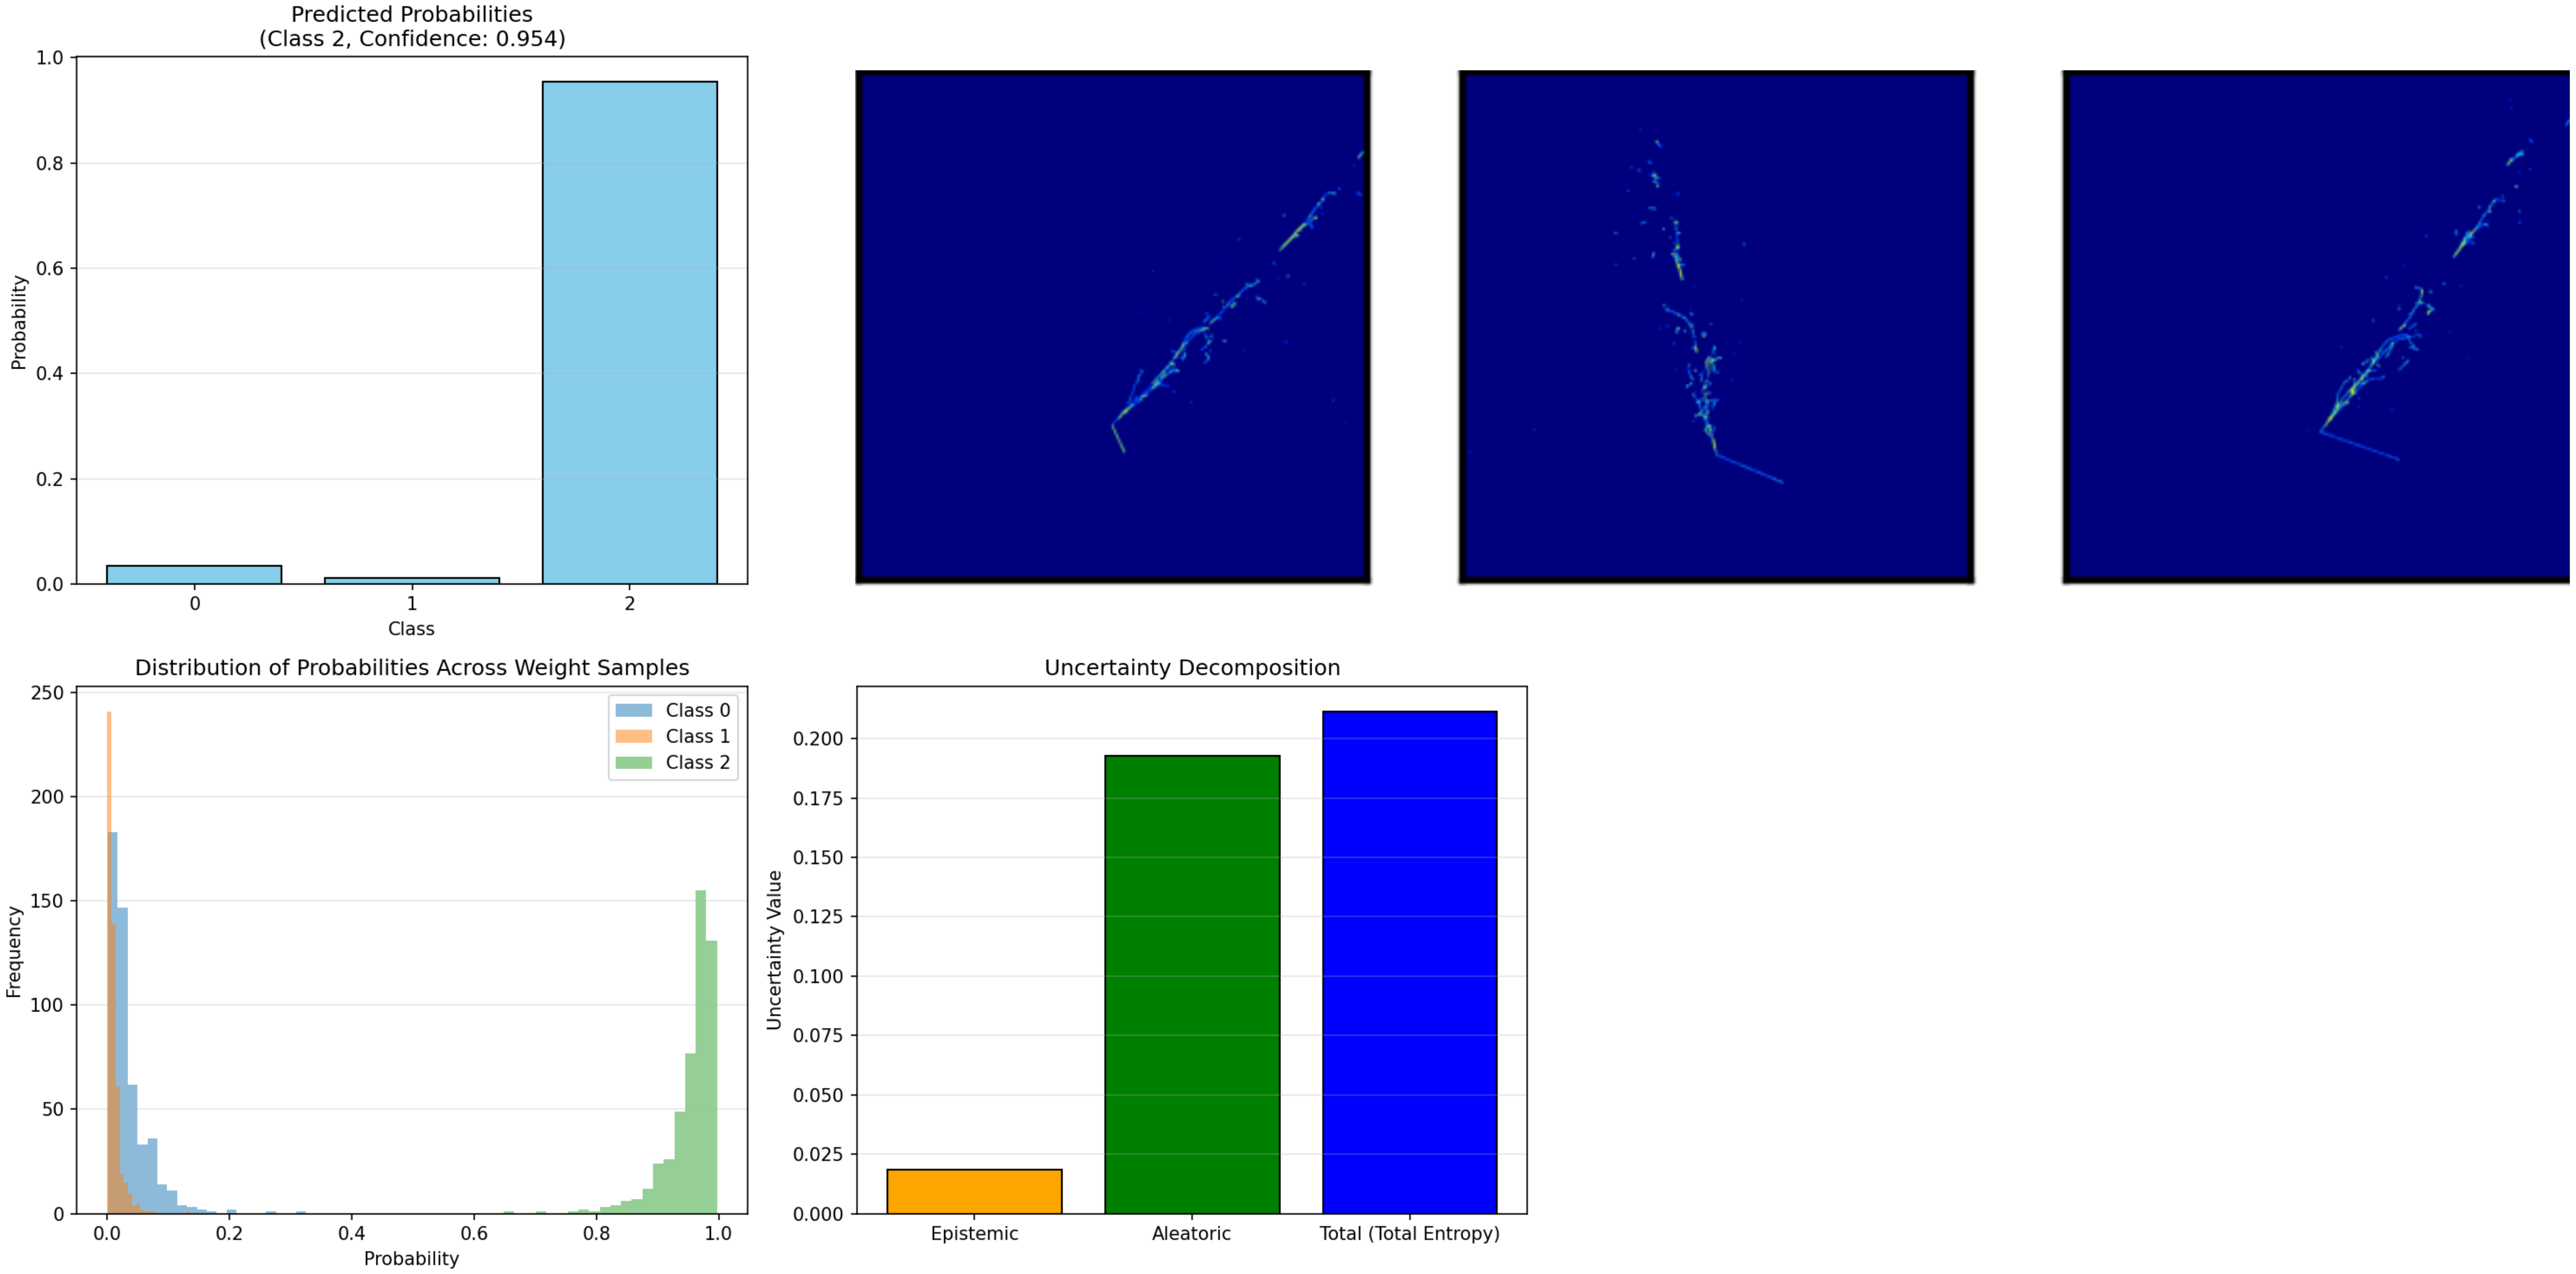

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

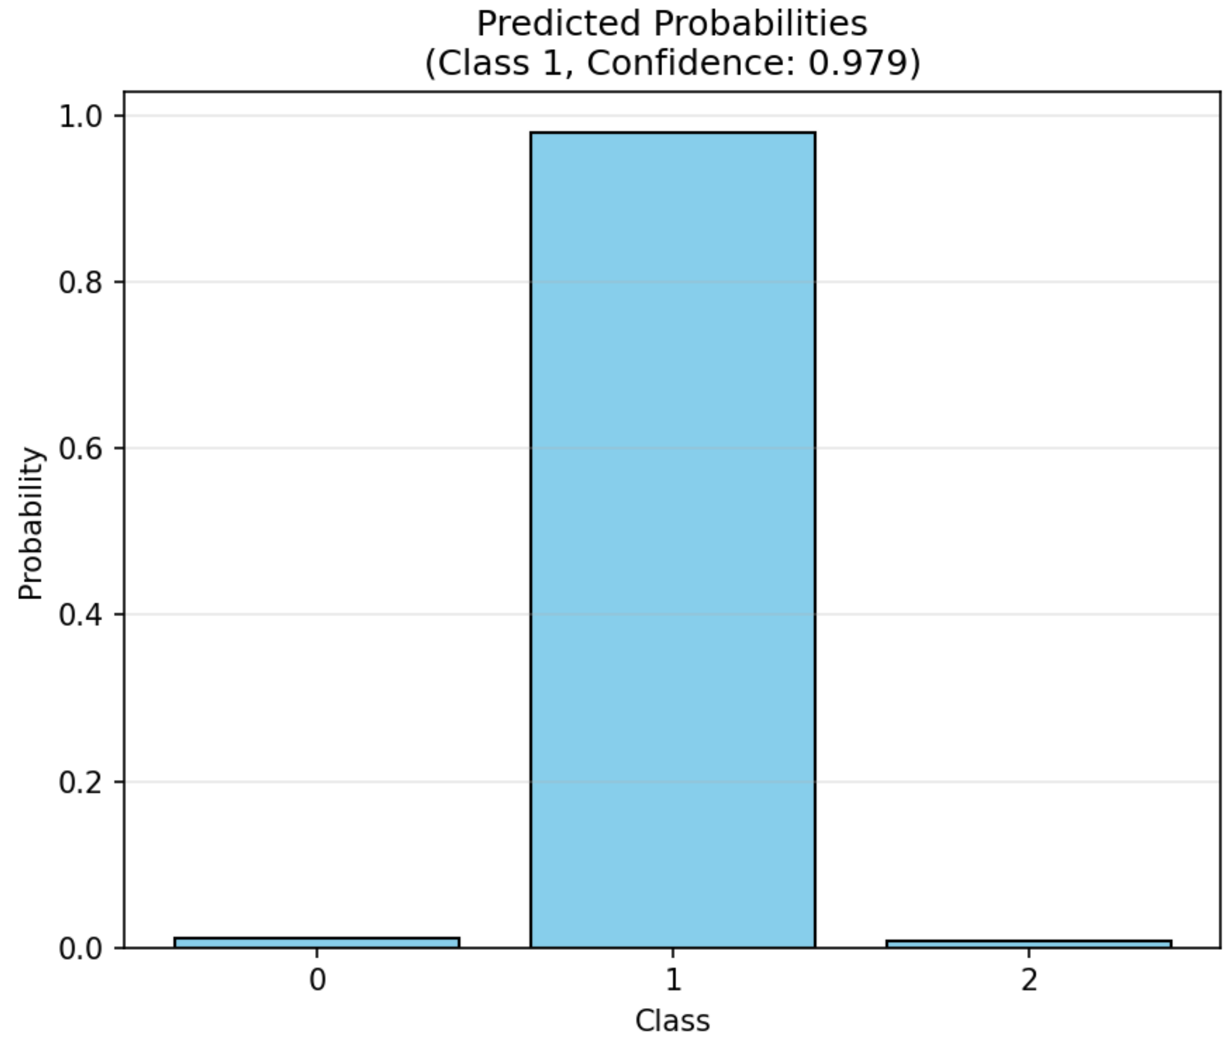

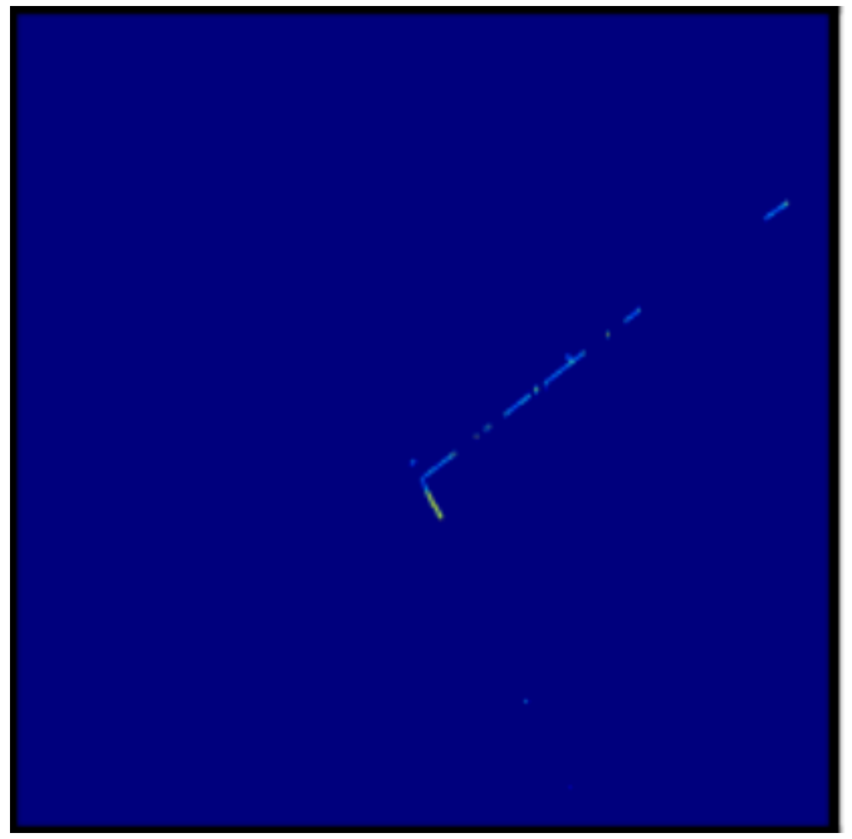

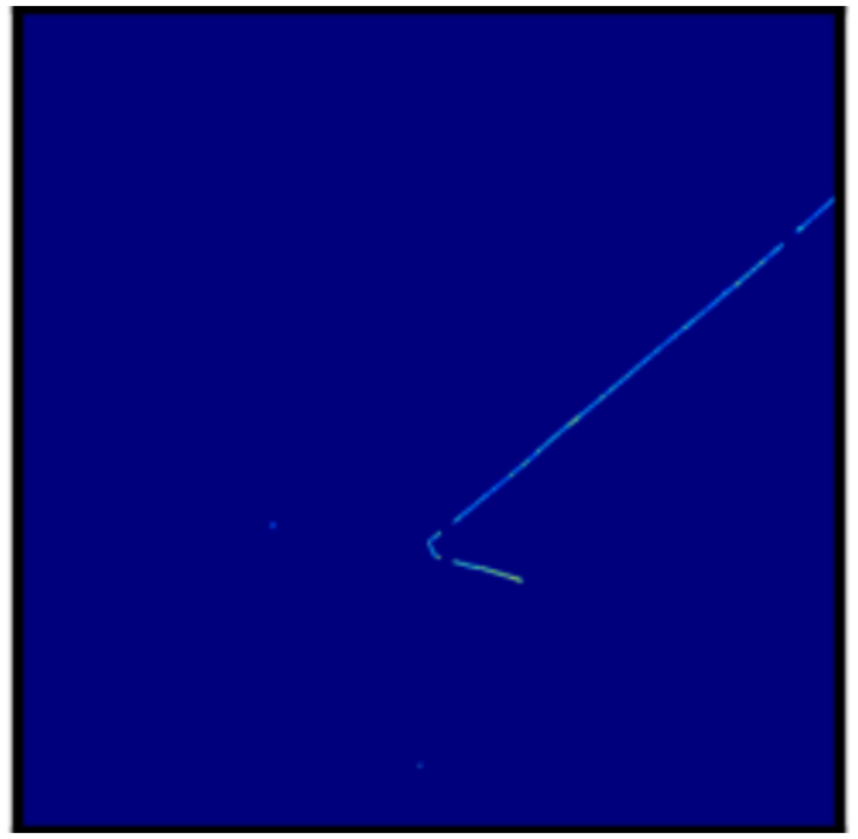

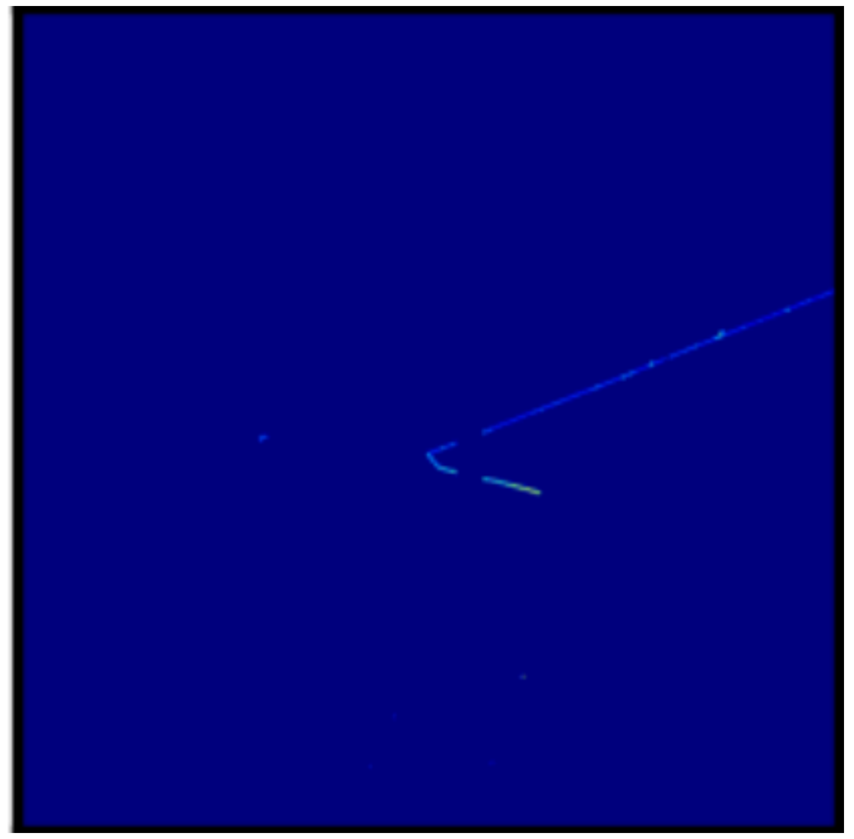

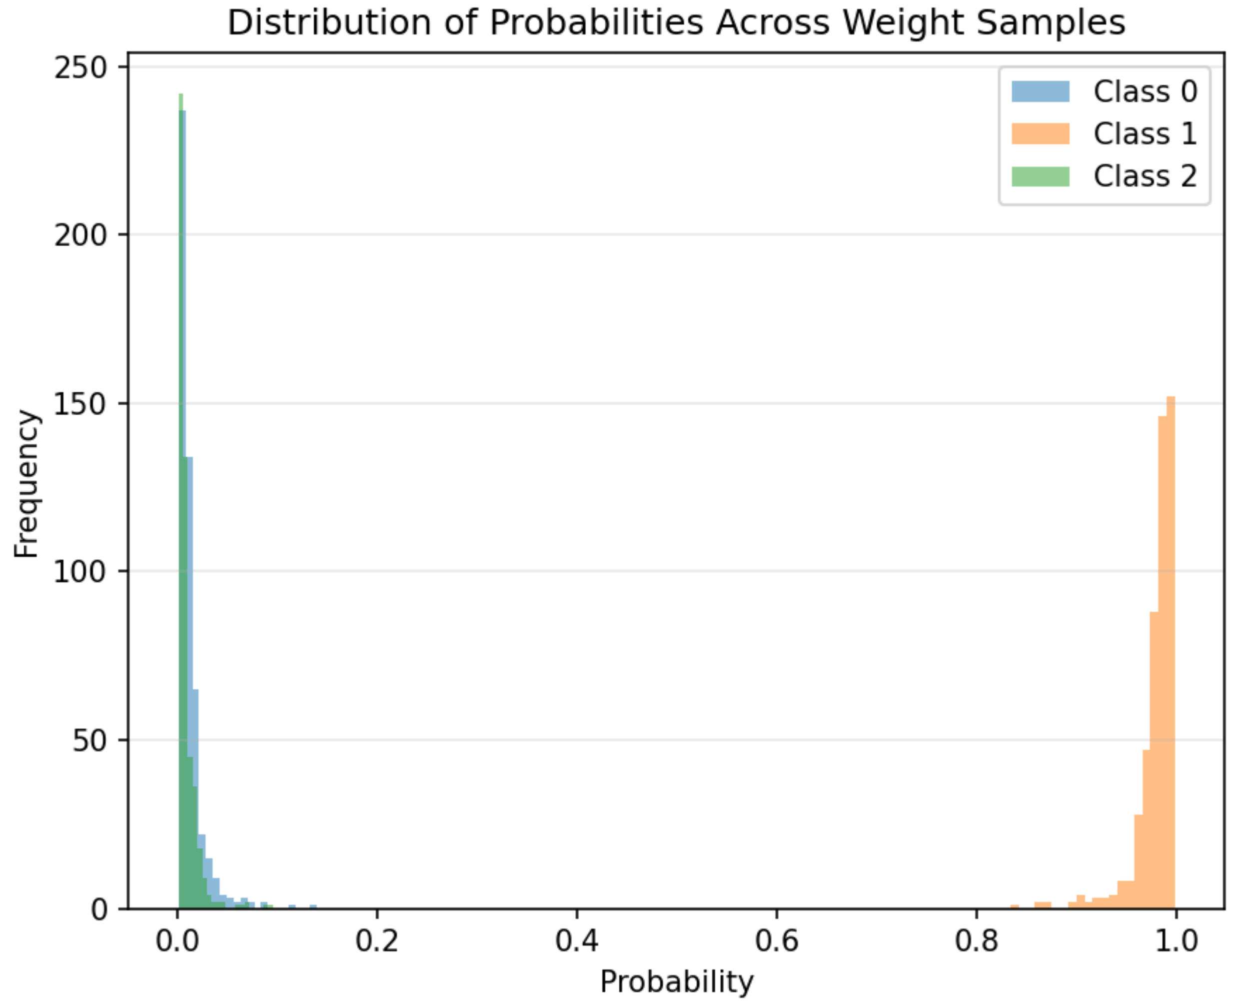

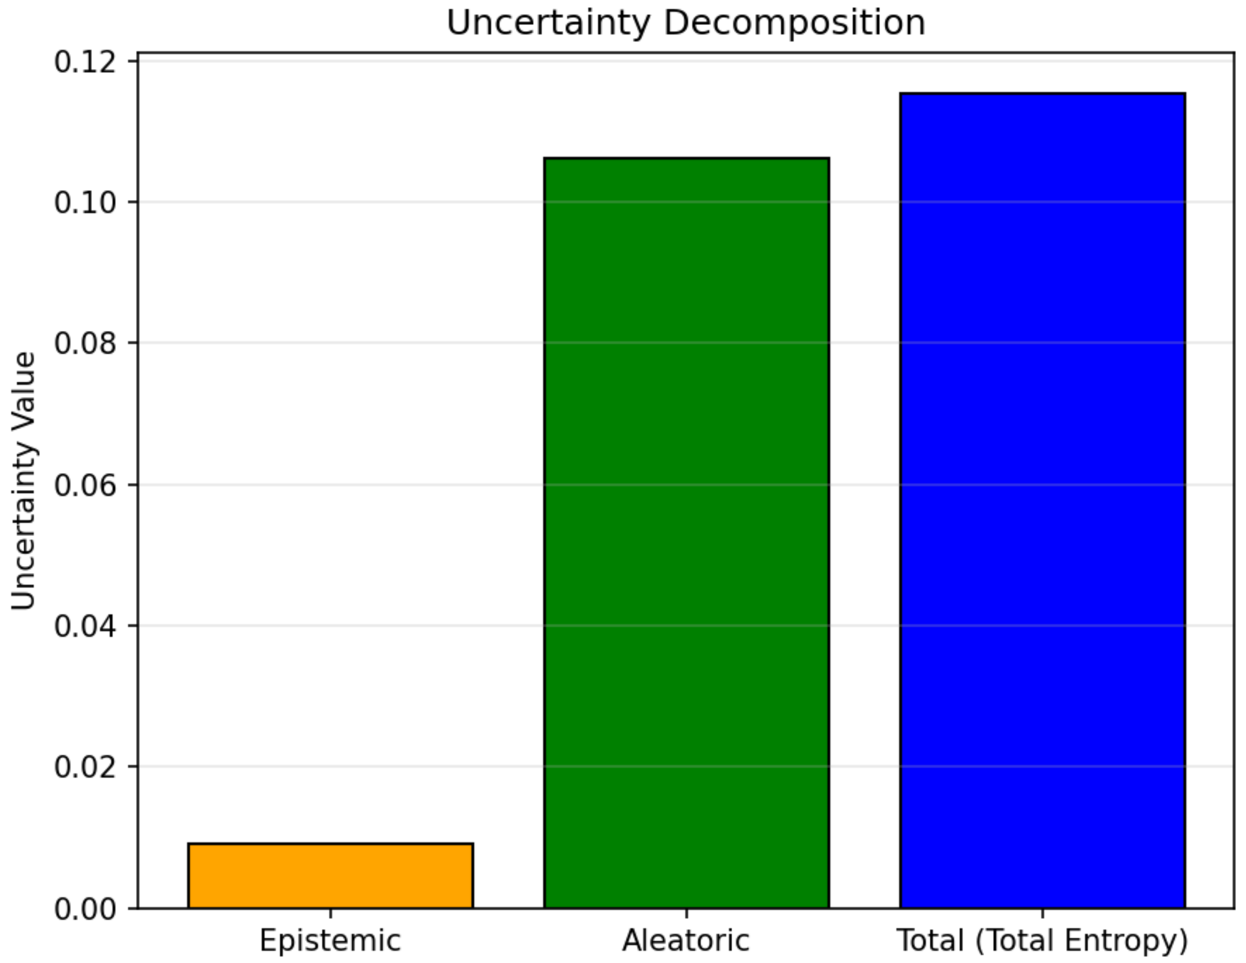

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

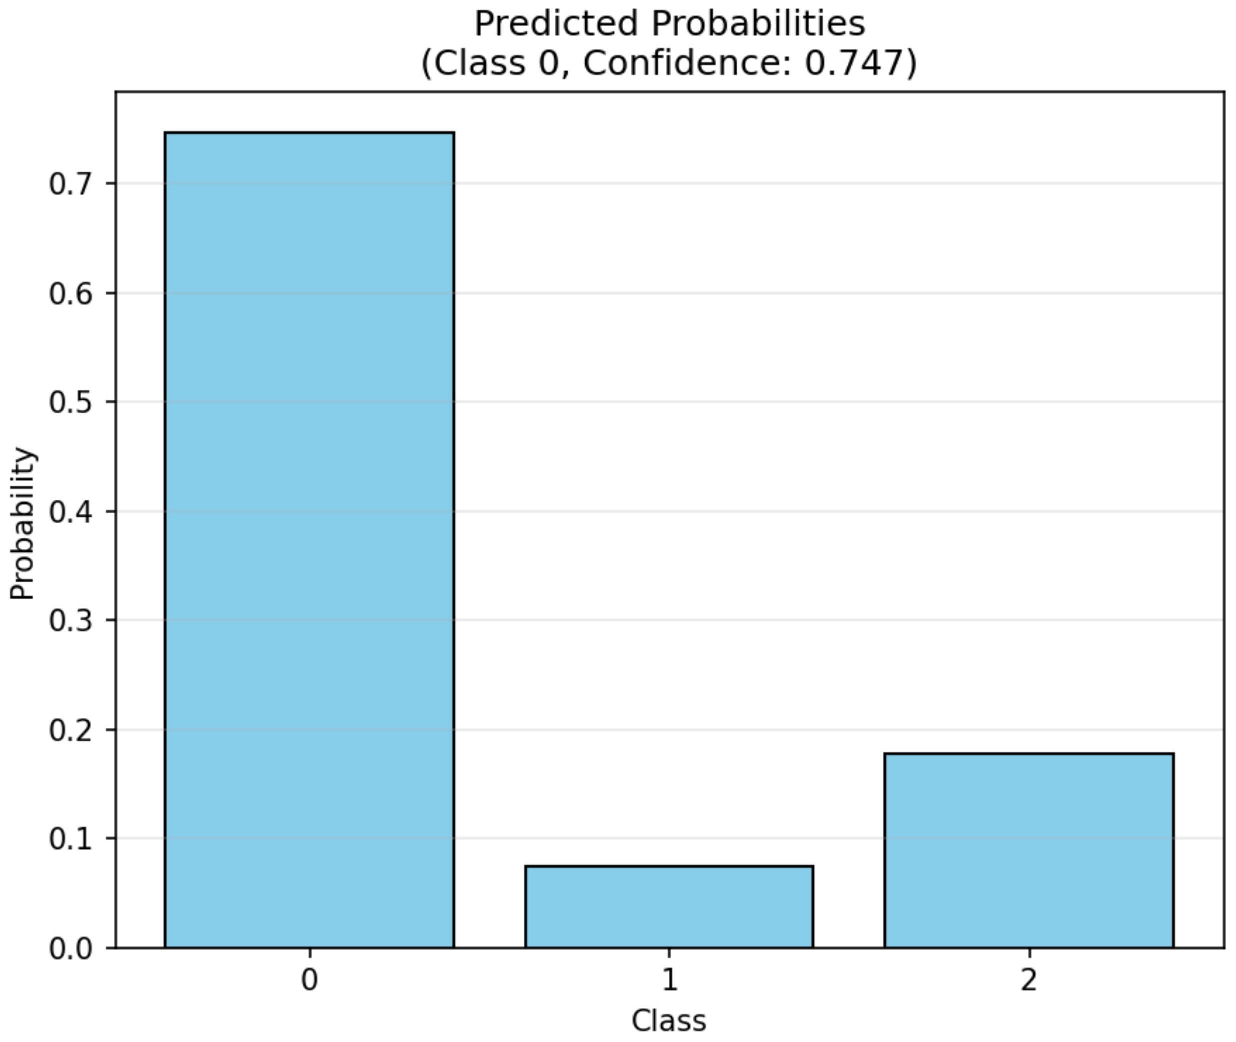

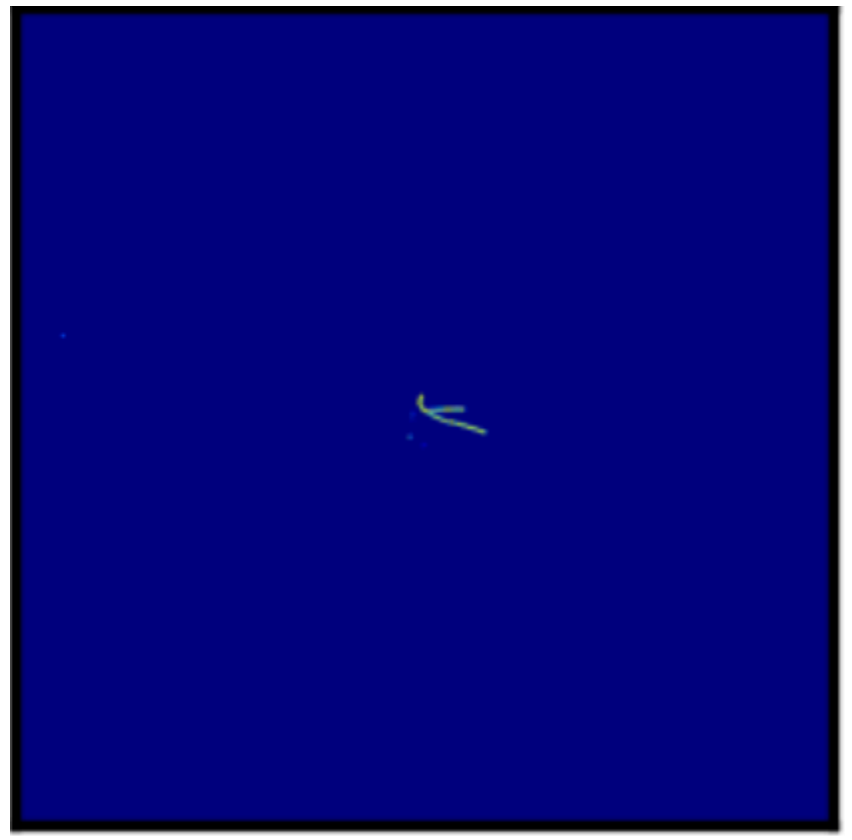

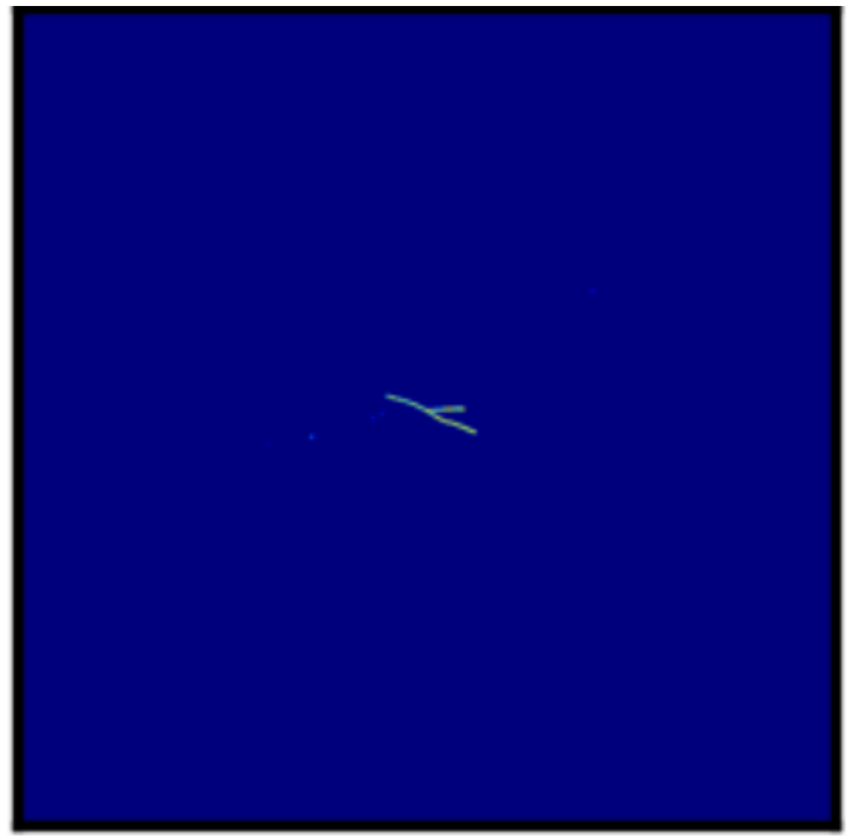

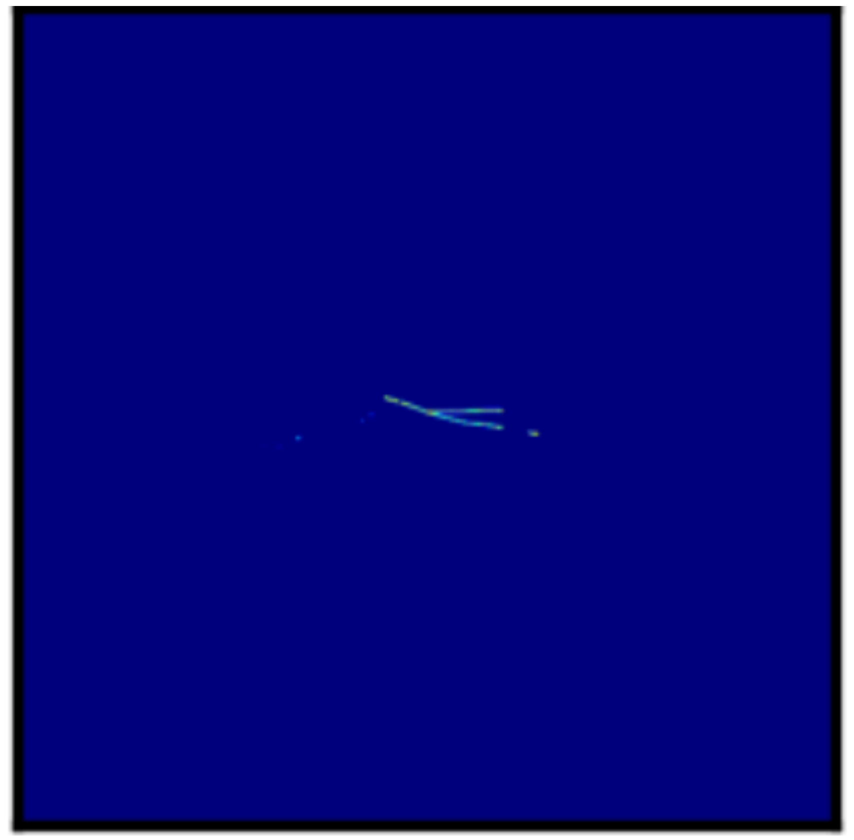

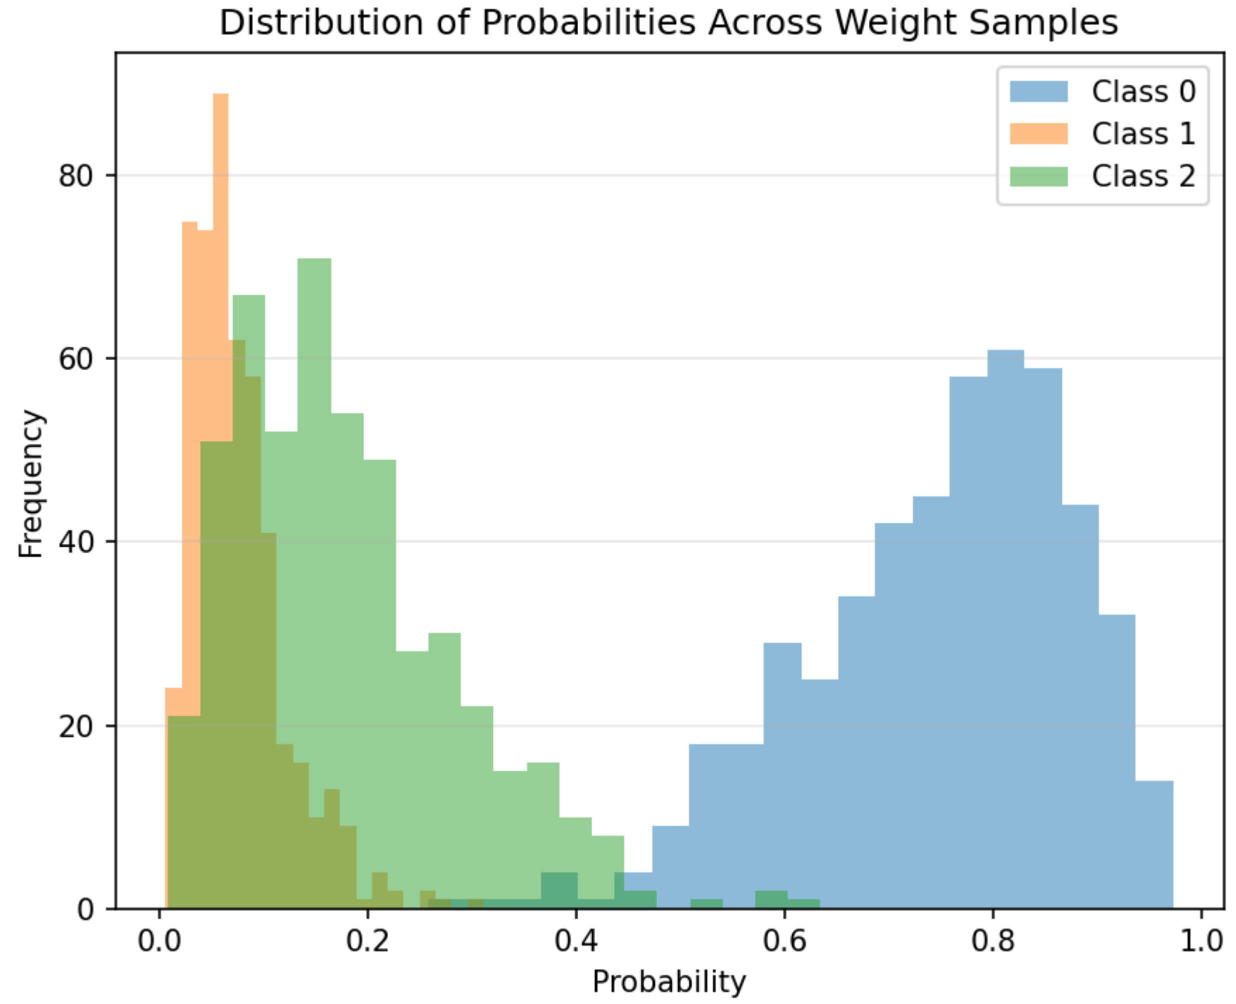

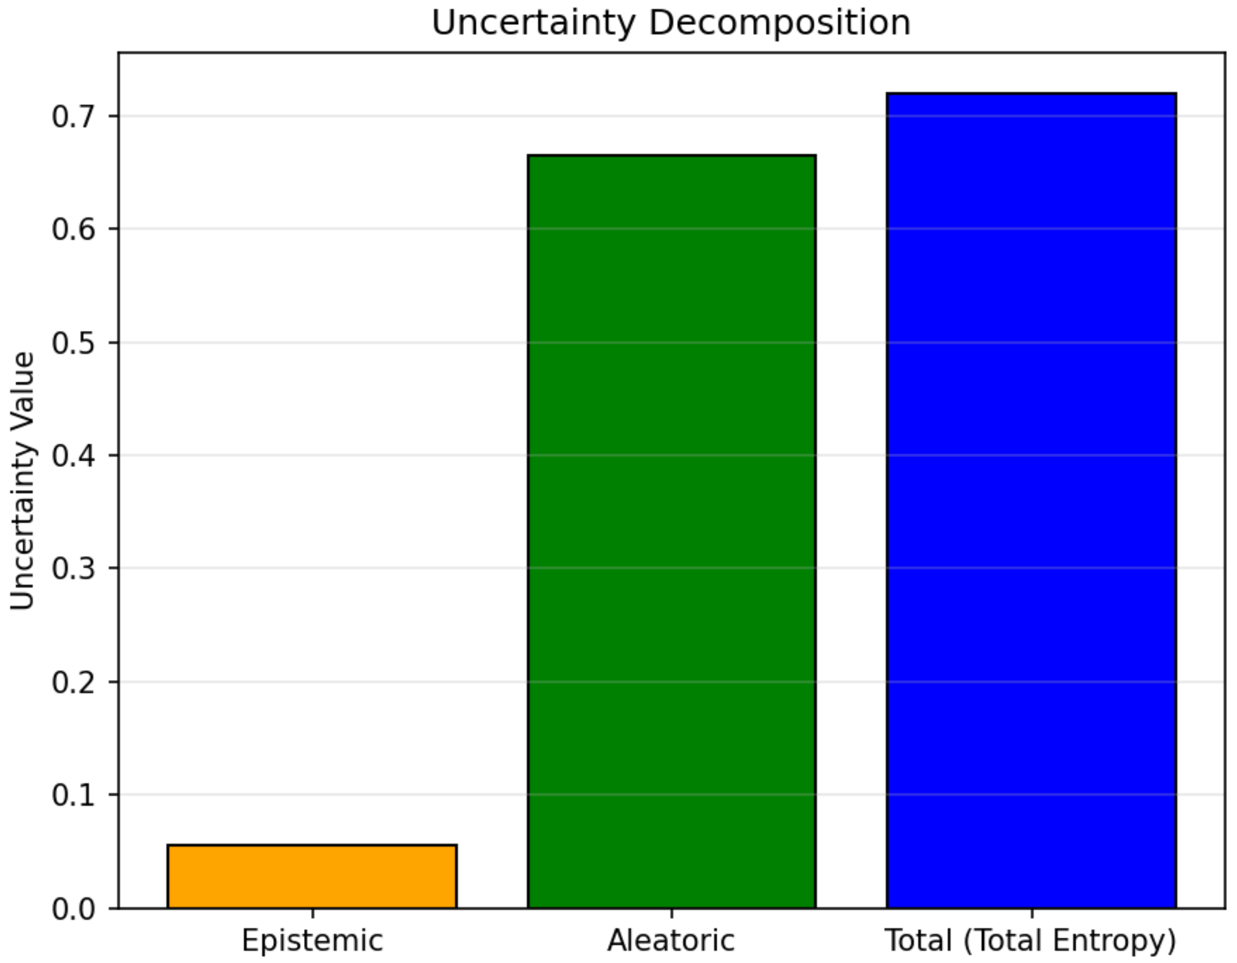

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

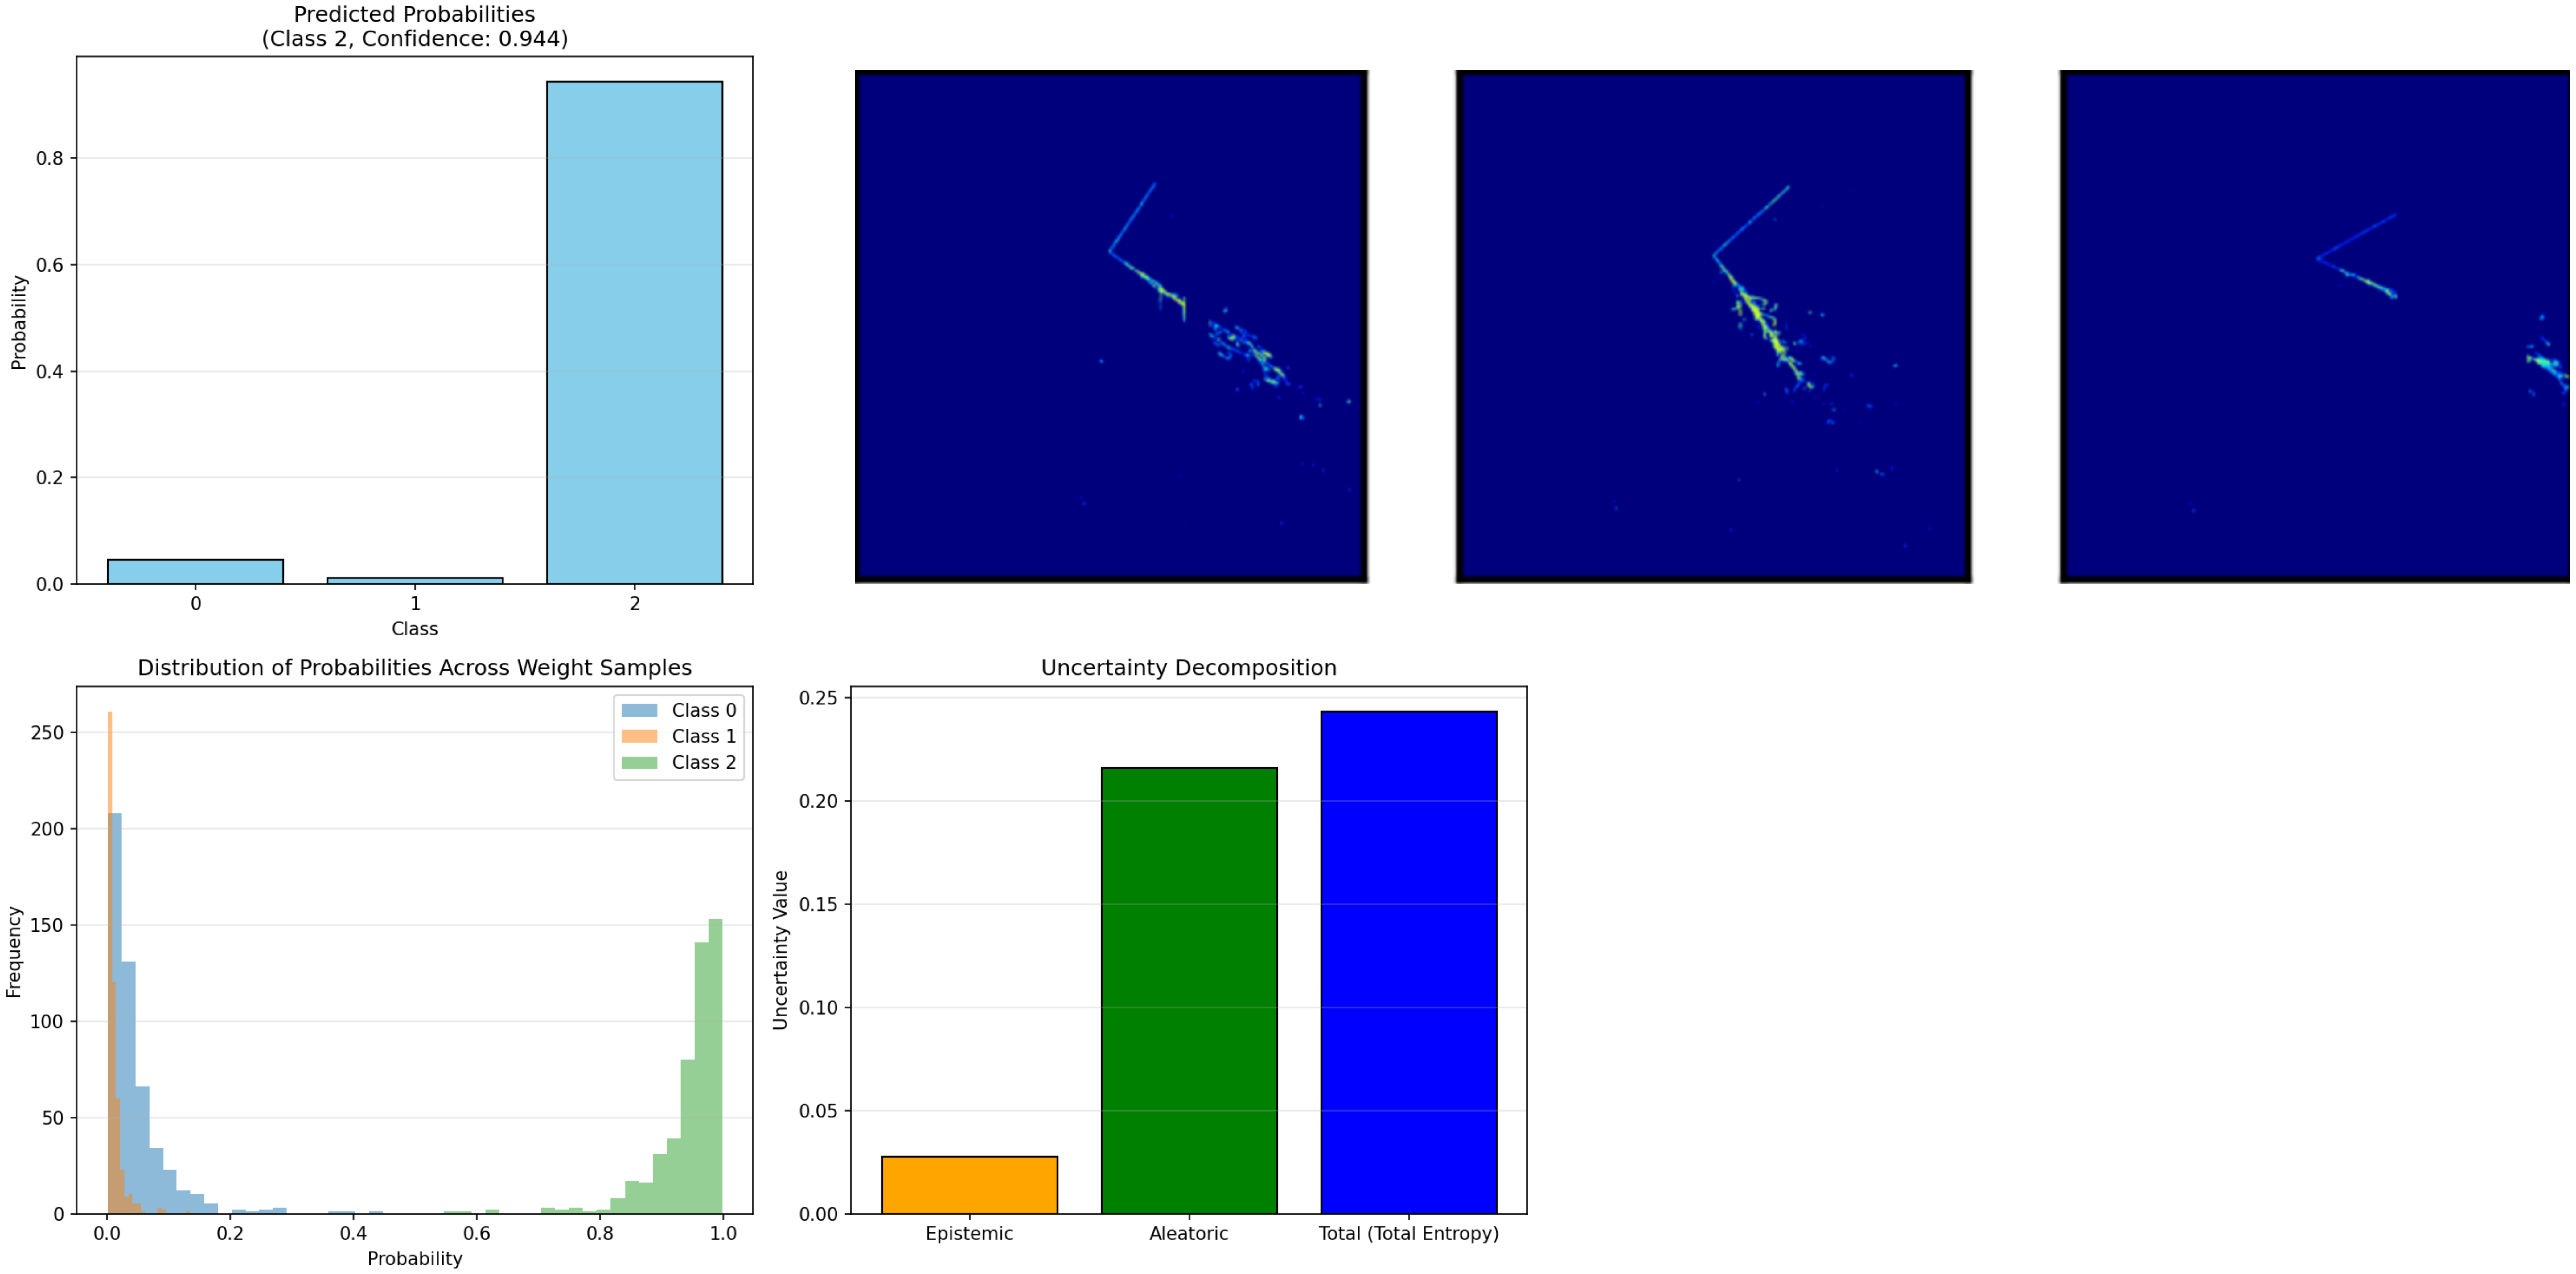

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

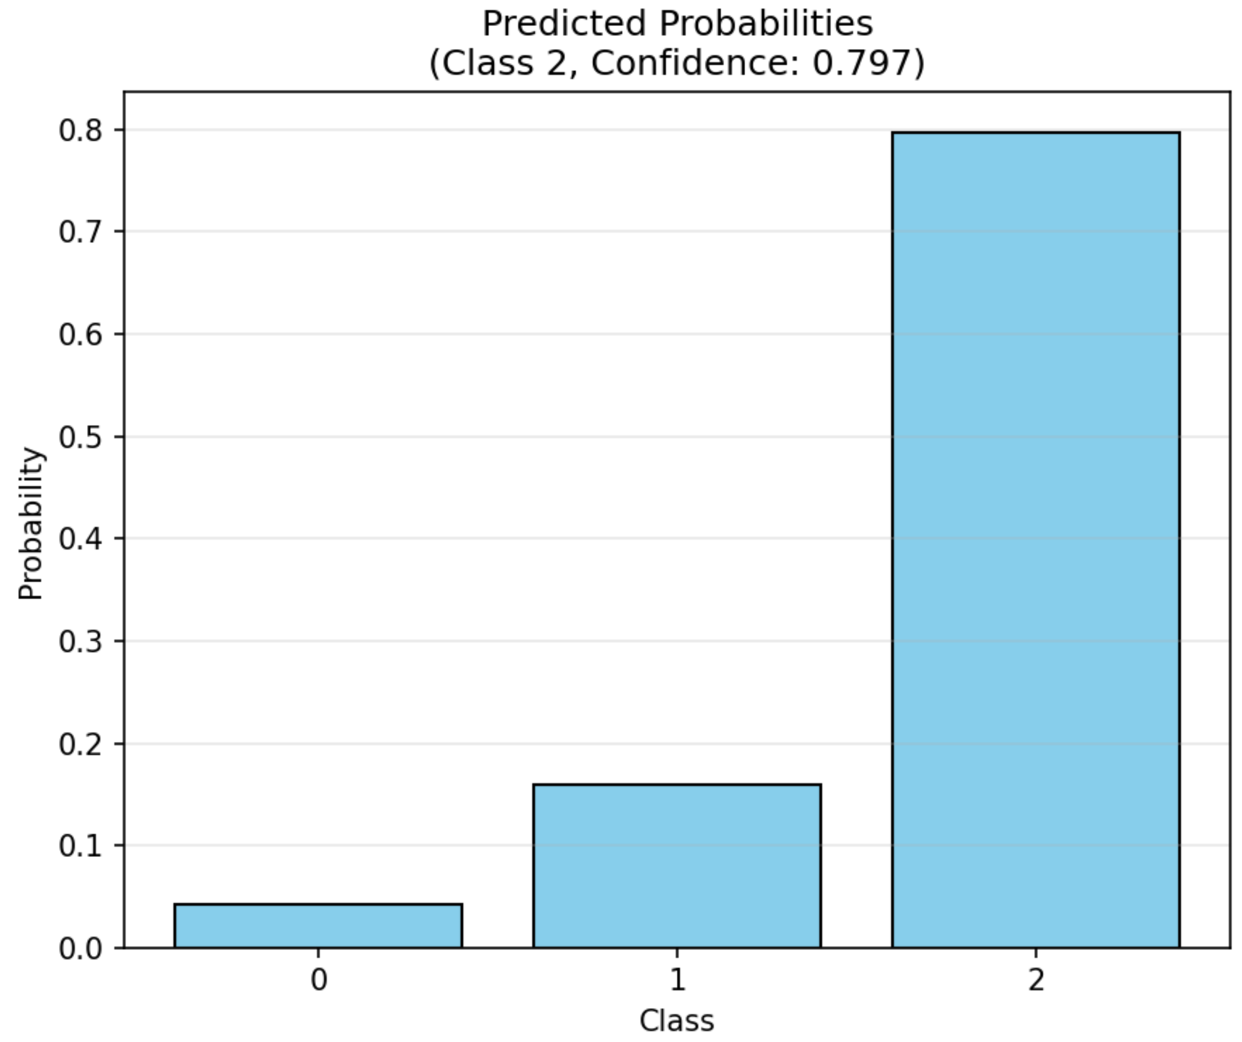

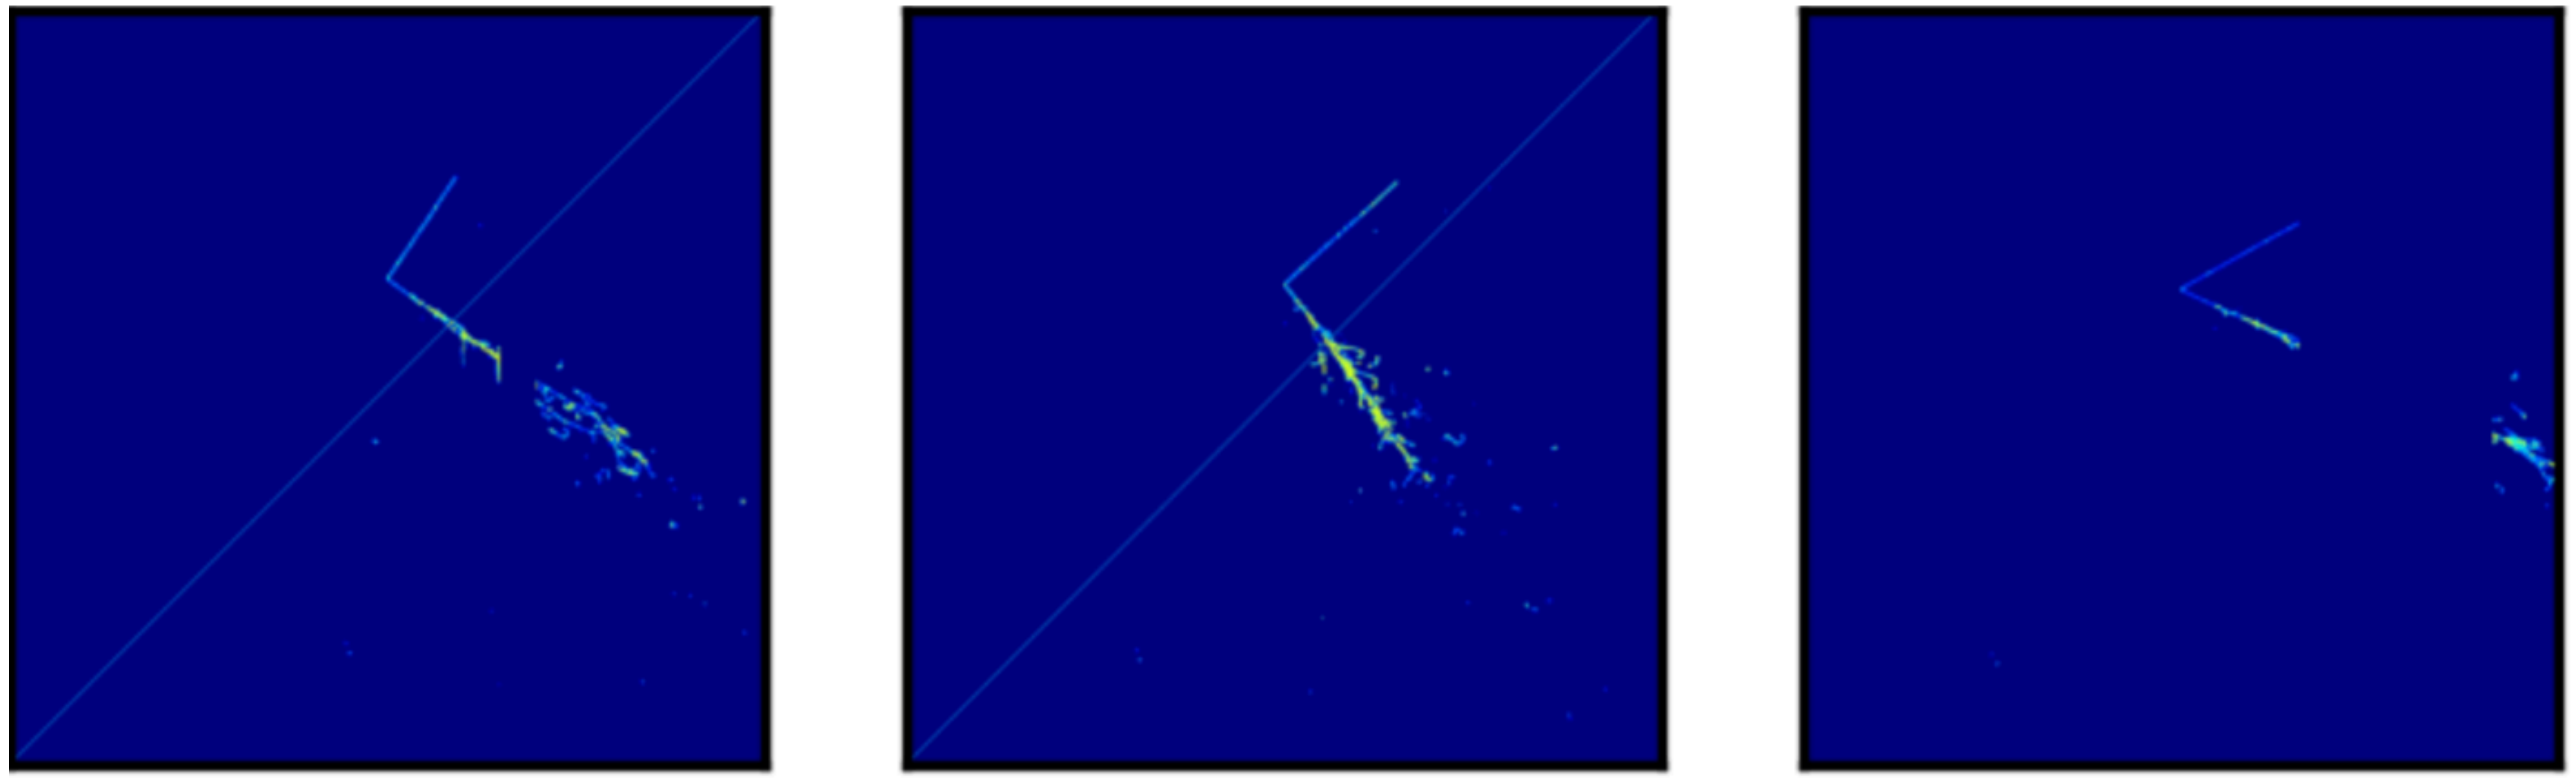

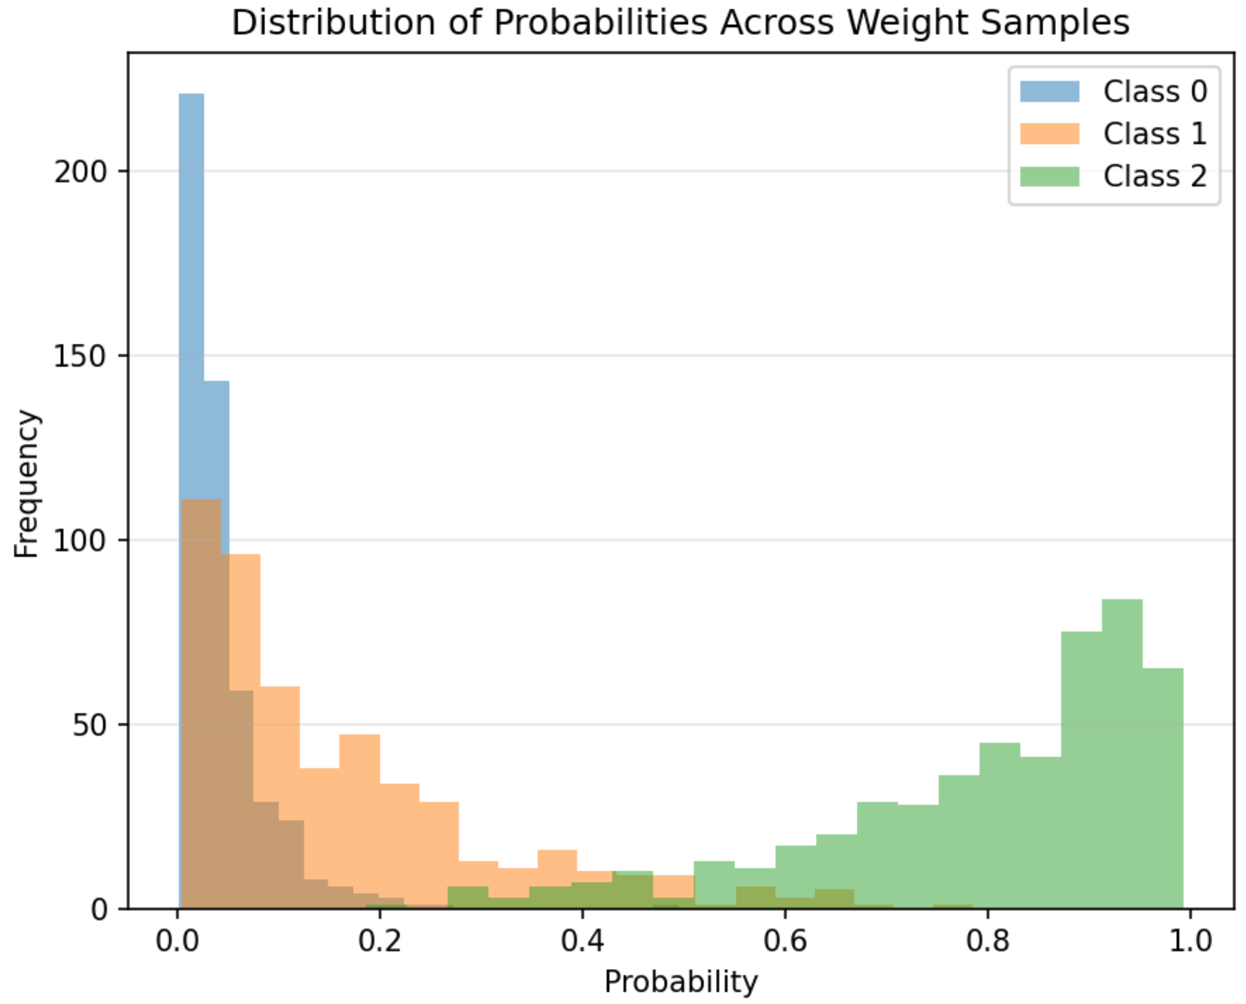

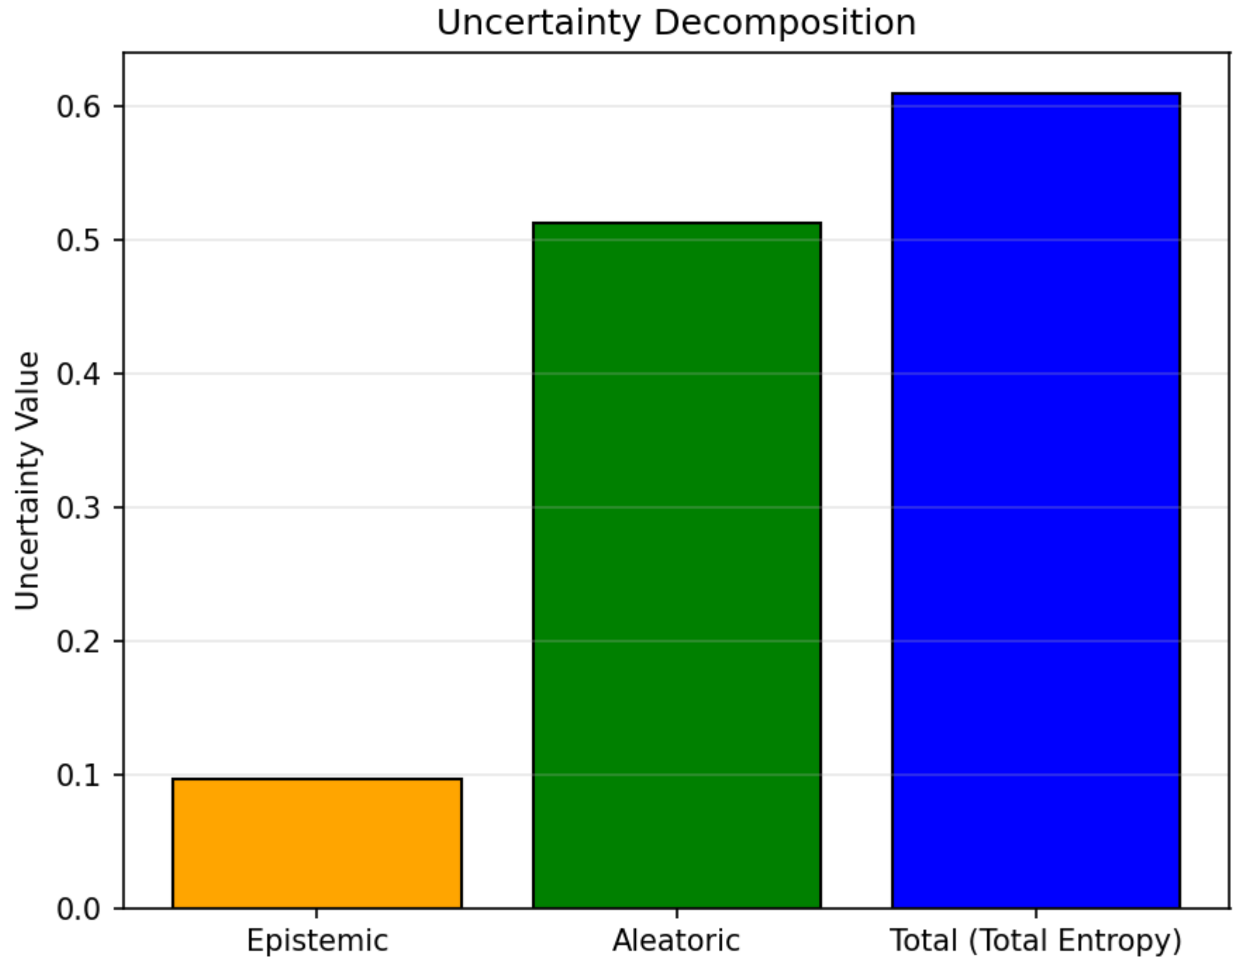

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

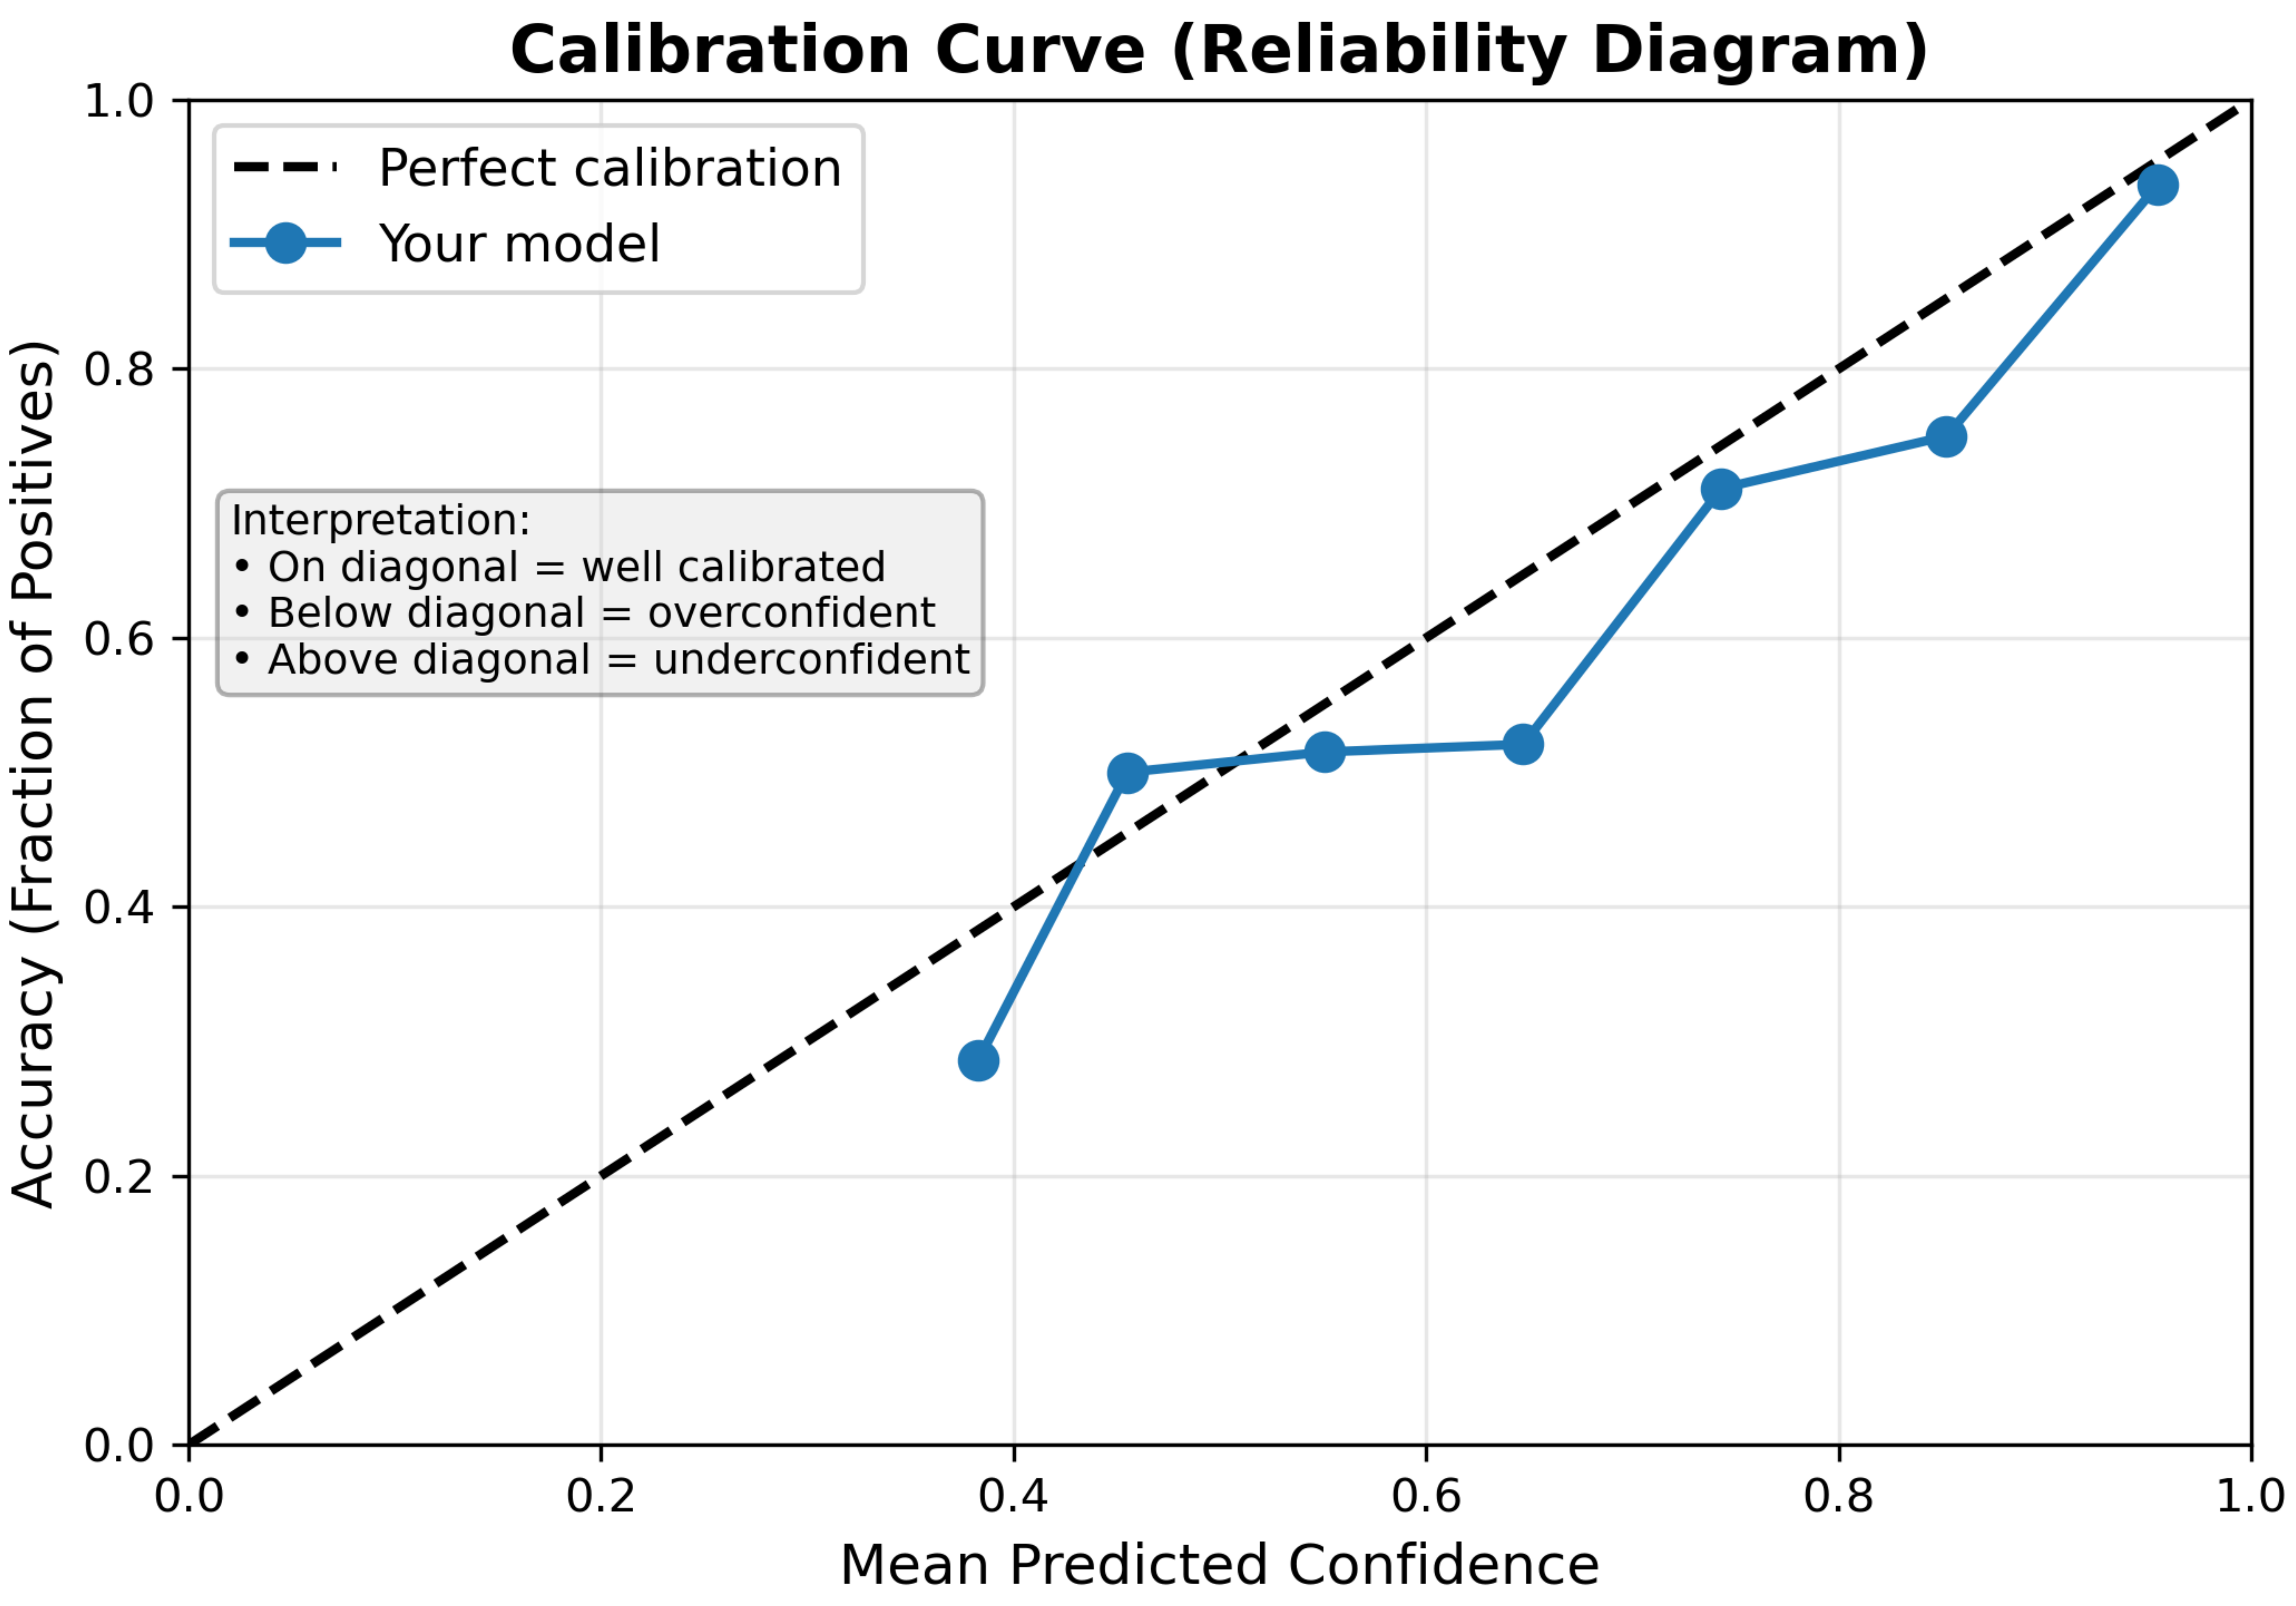

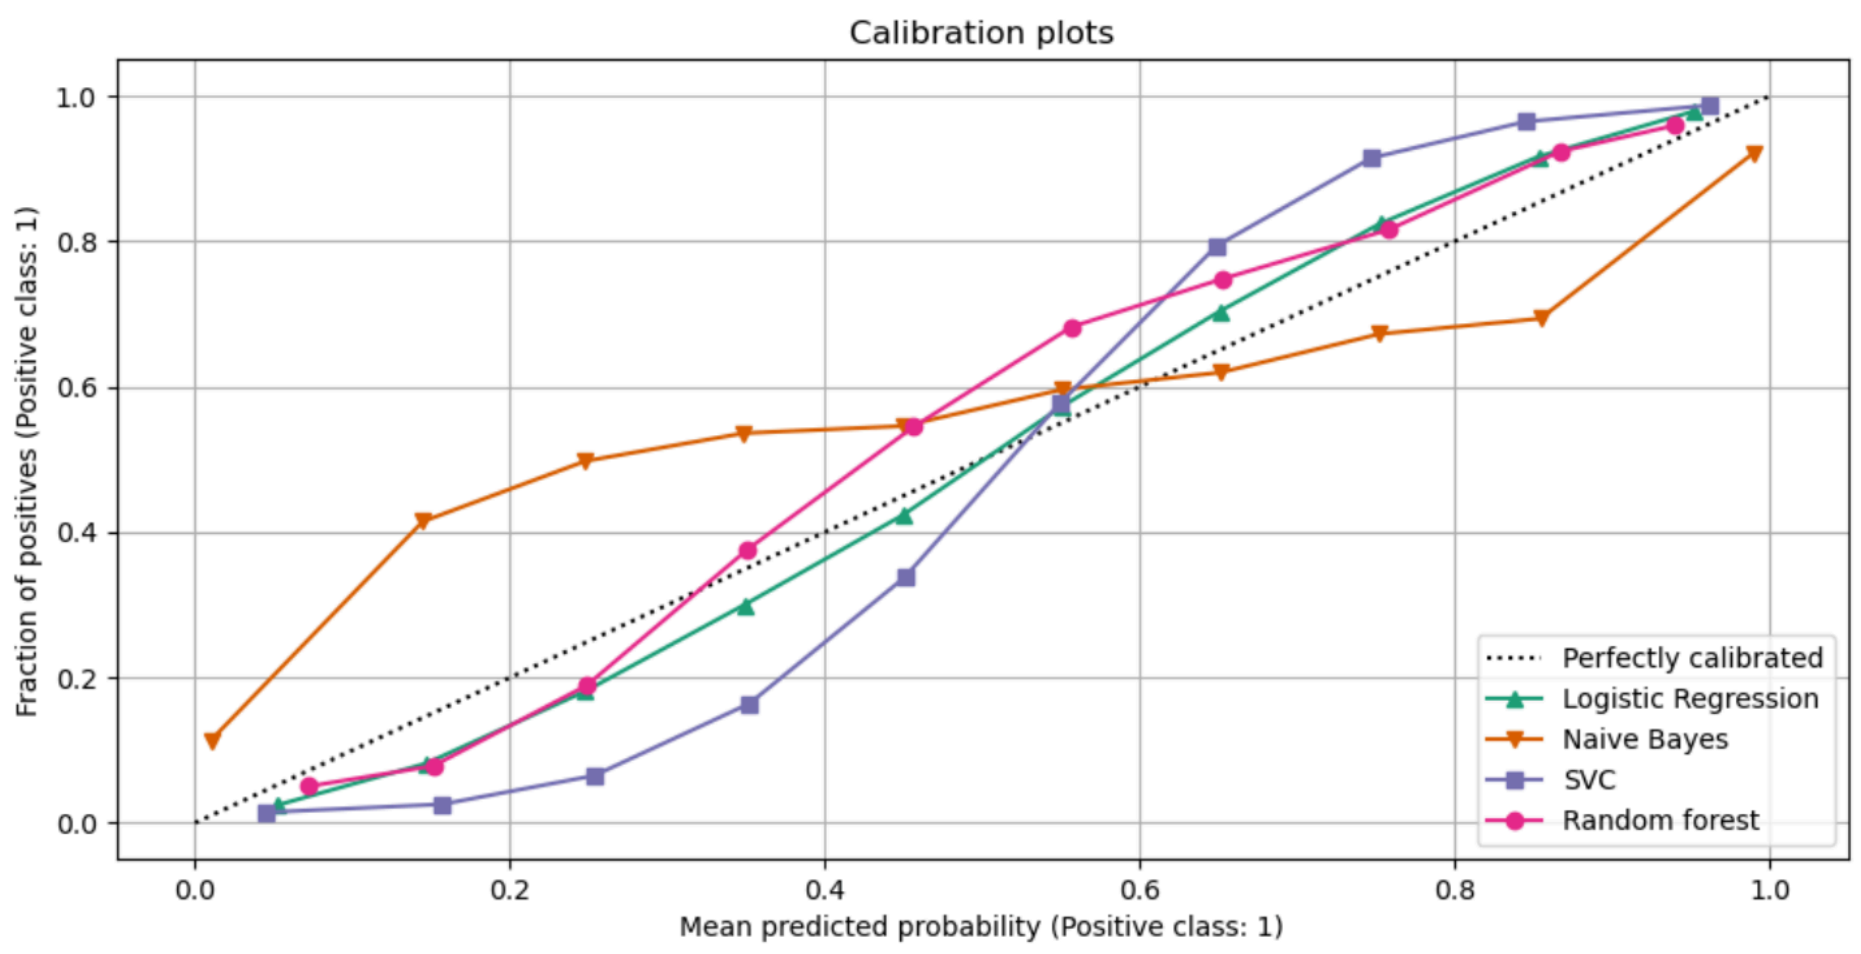

Uncertainty Quantification with BNN

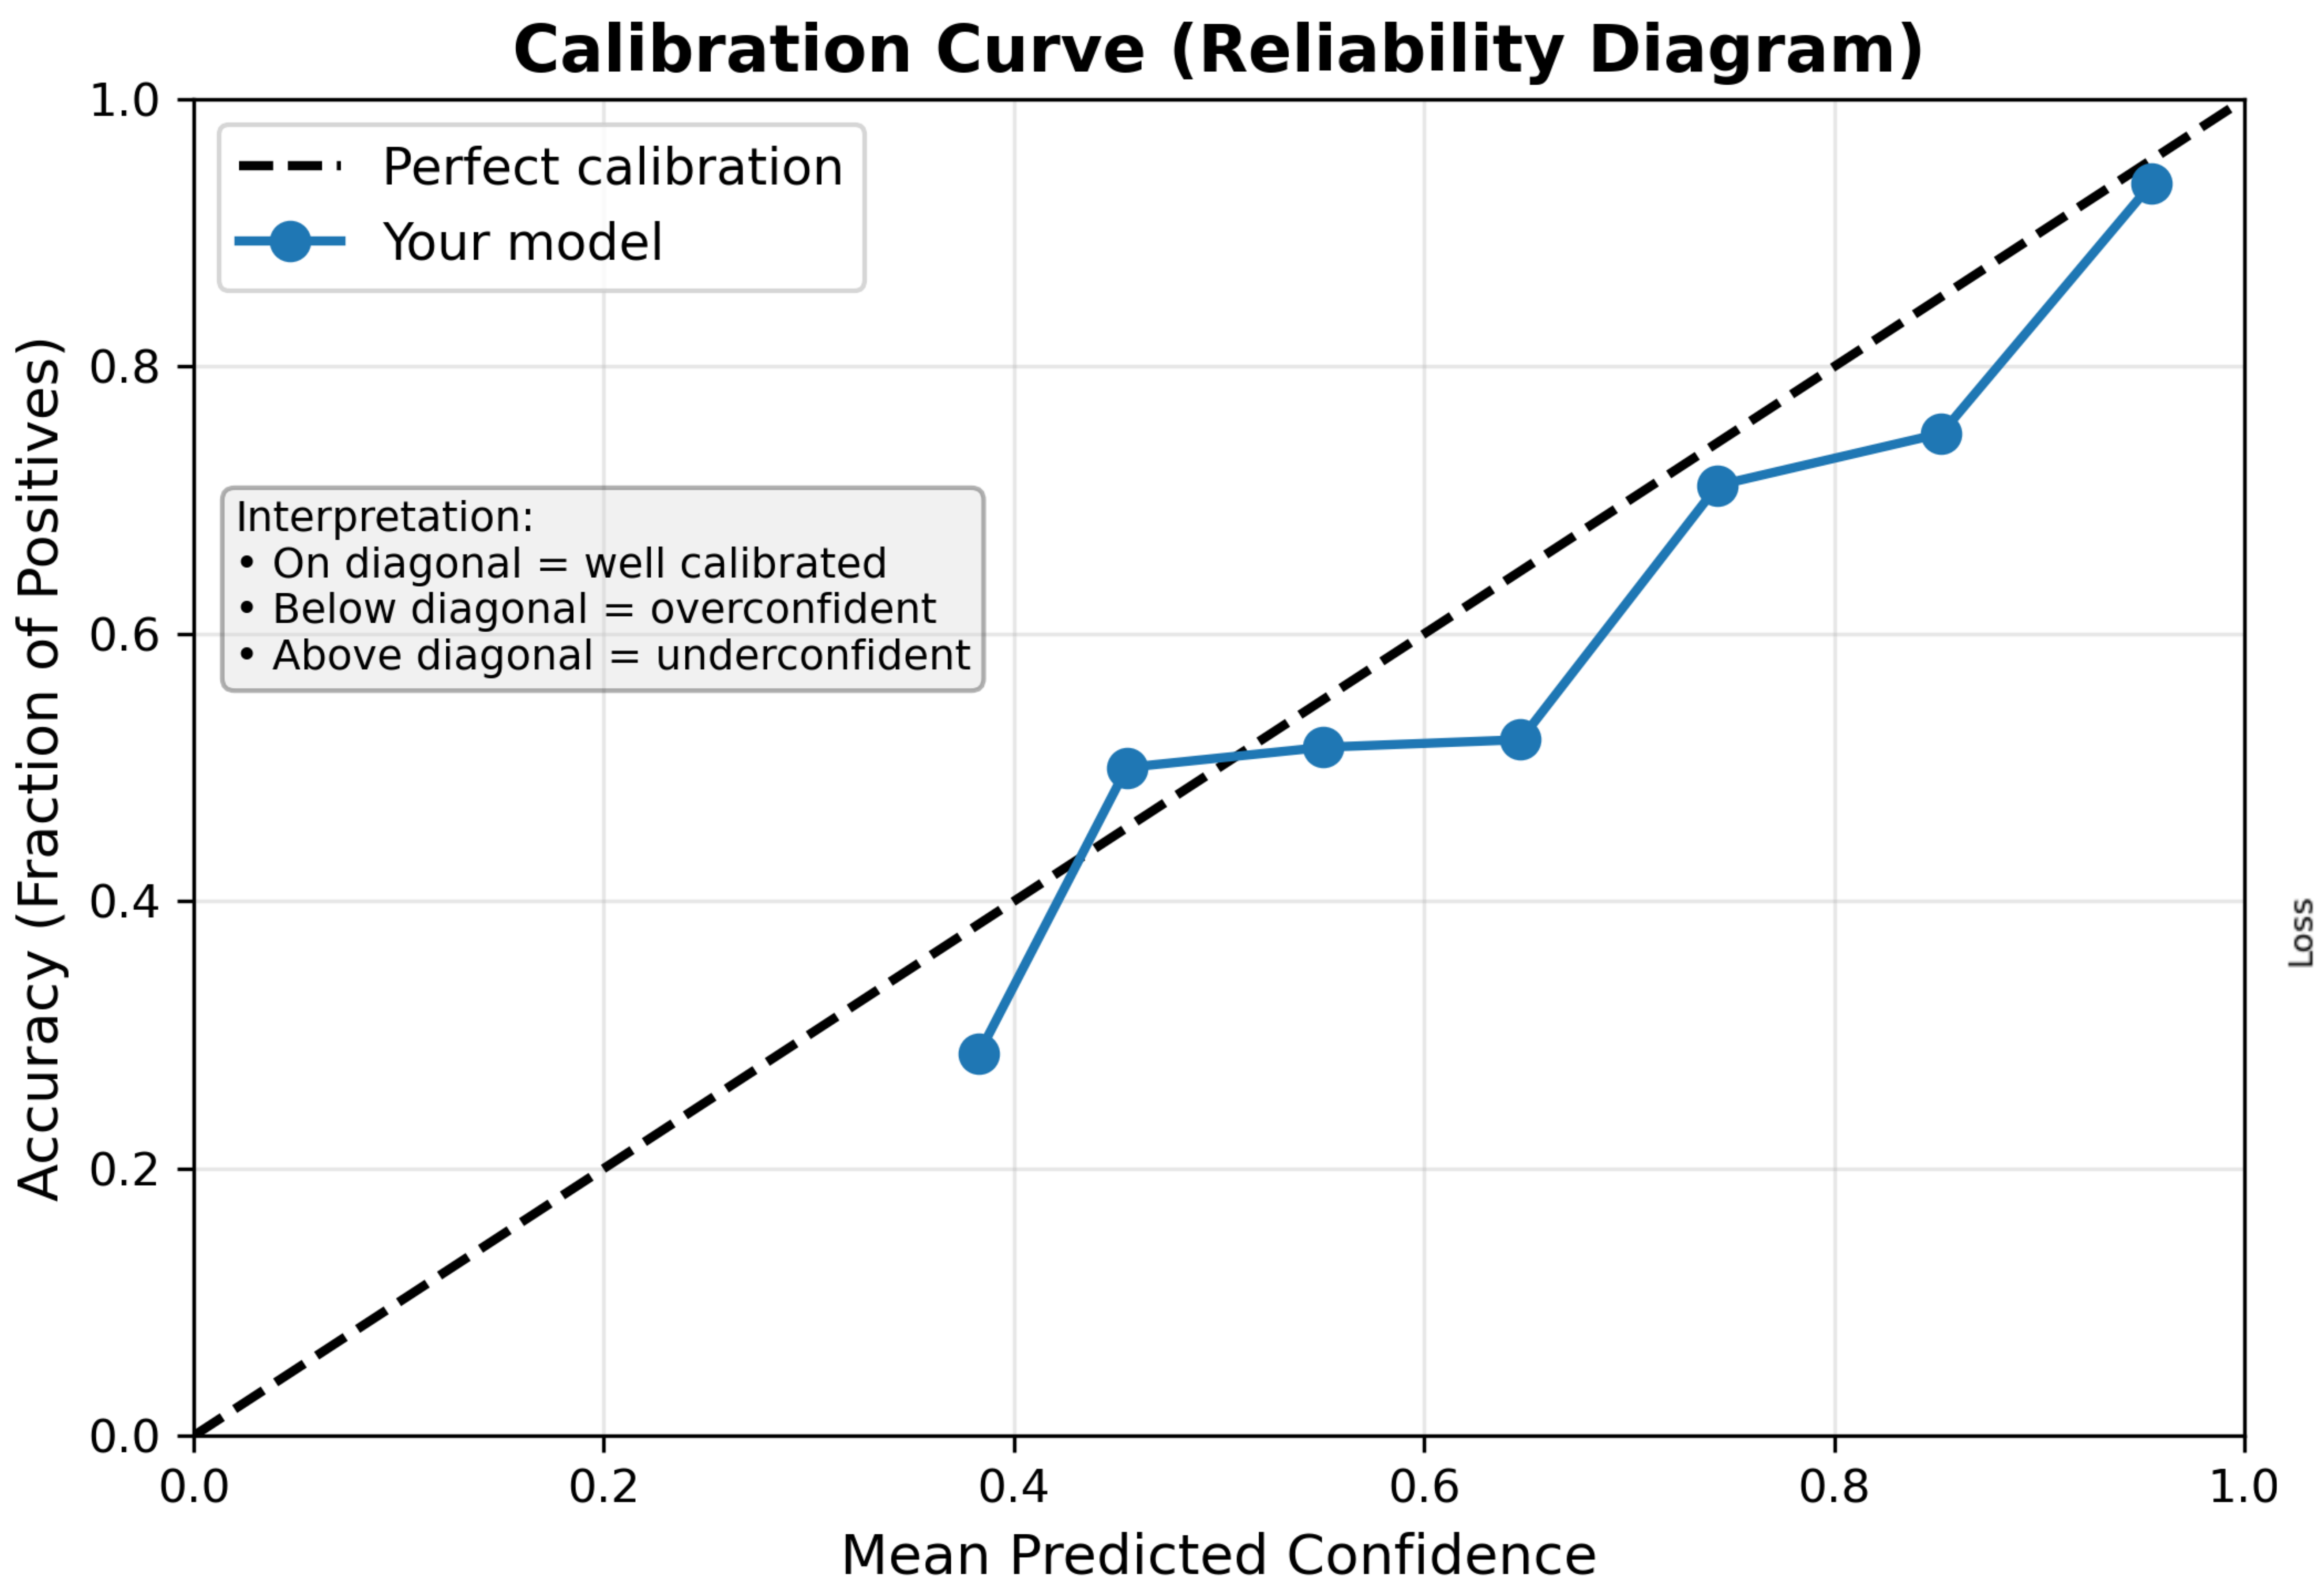

Uncertainty Quantification with BNN

Can I introduce physics knowledge, a.k.a Physics-informed prior?

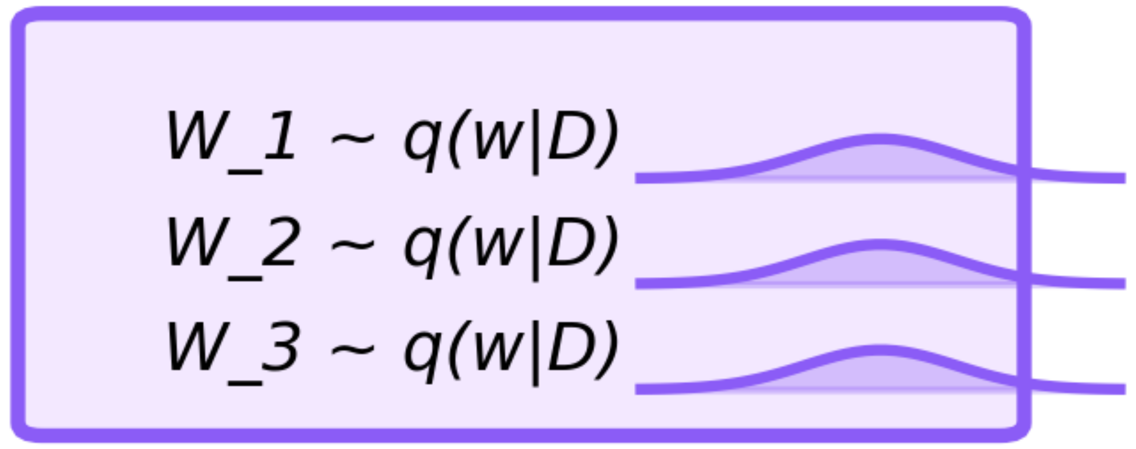

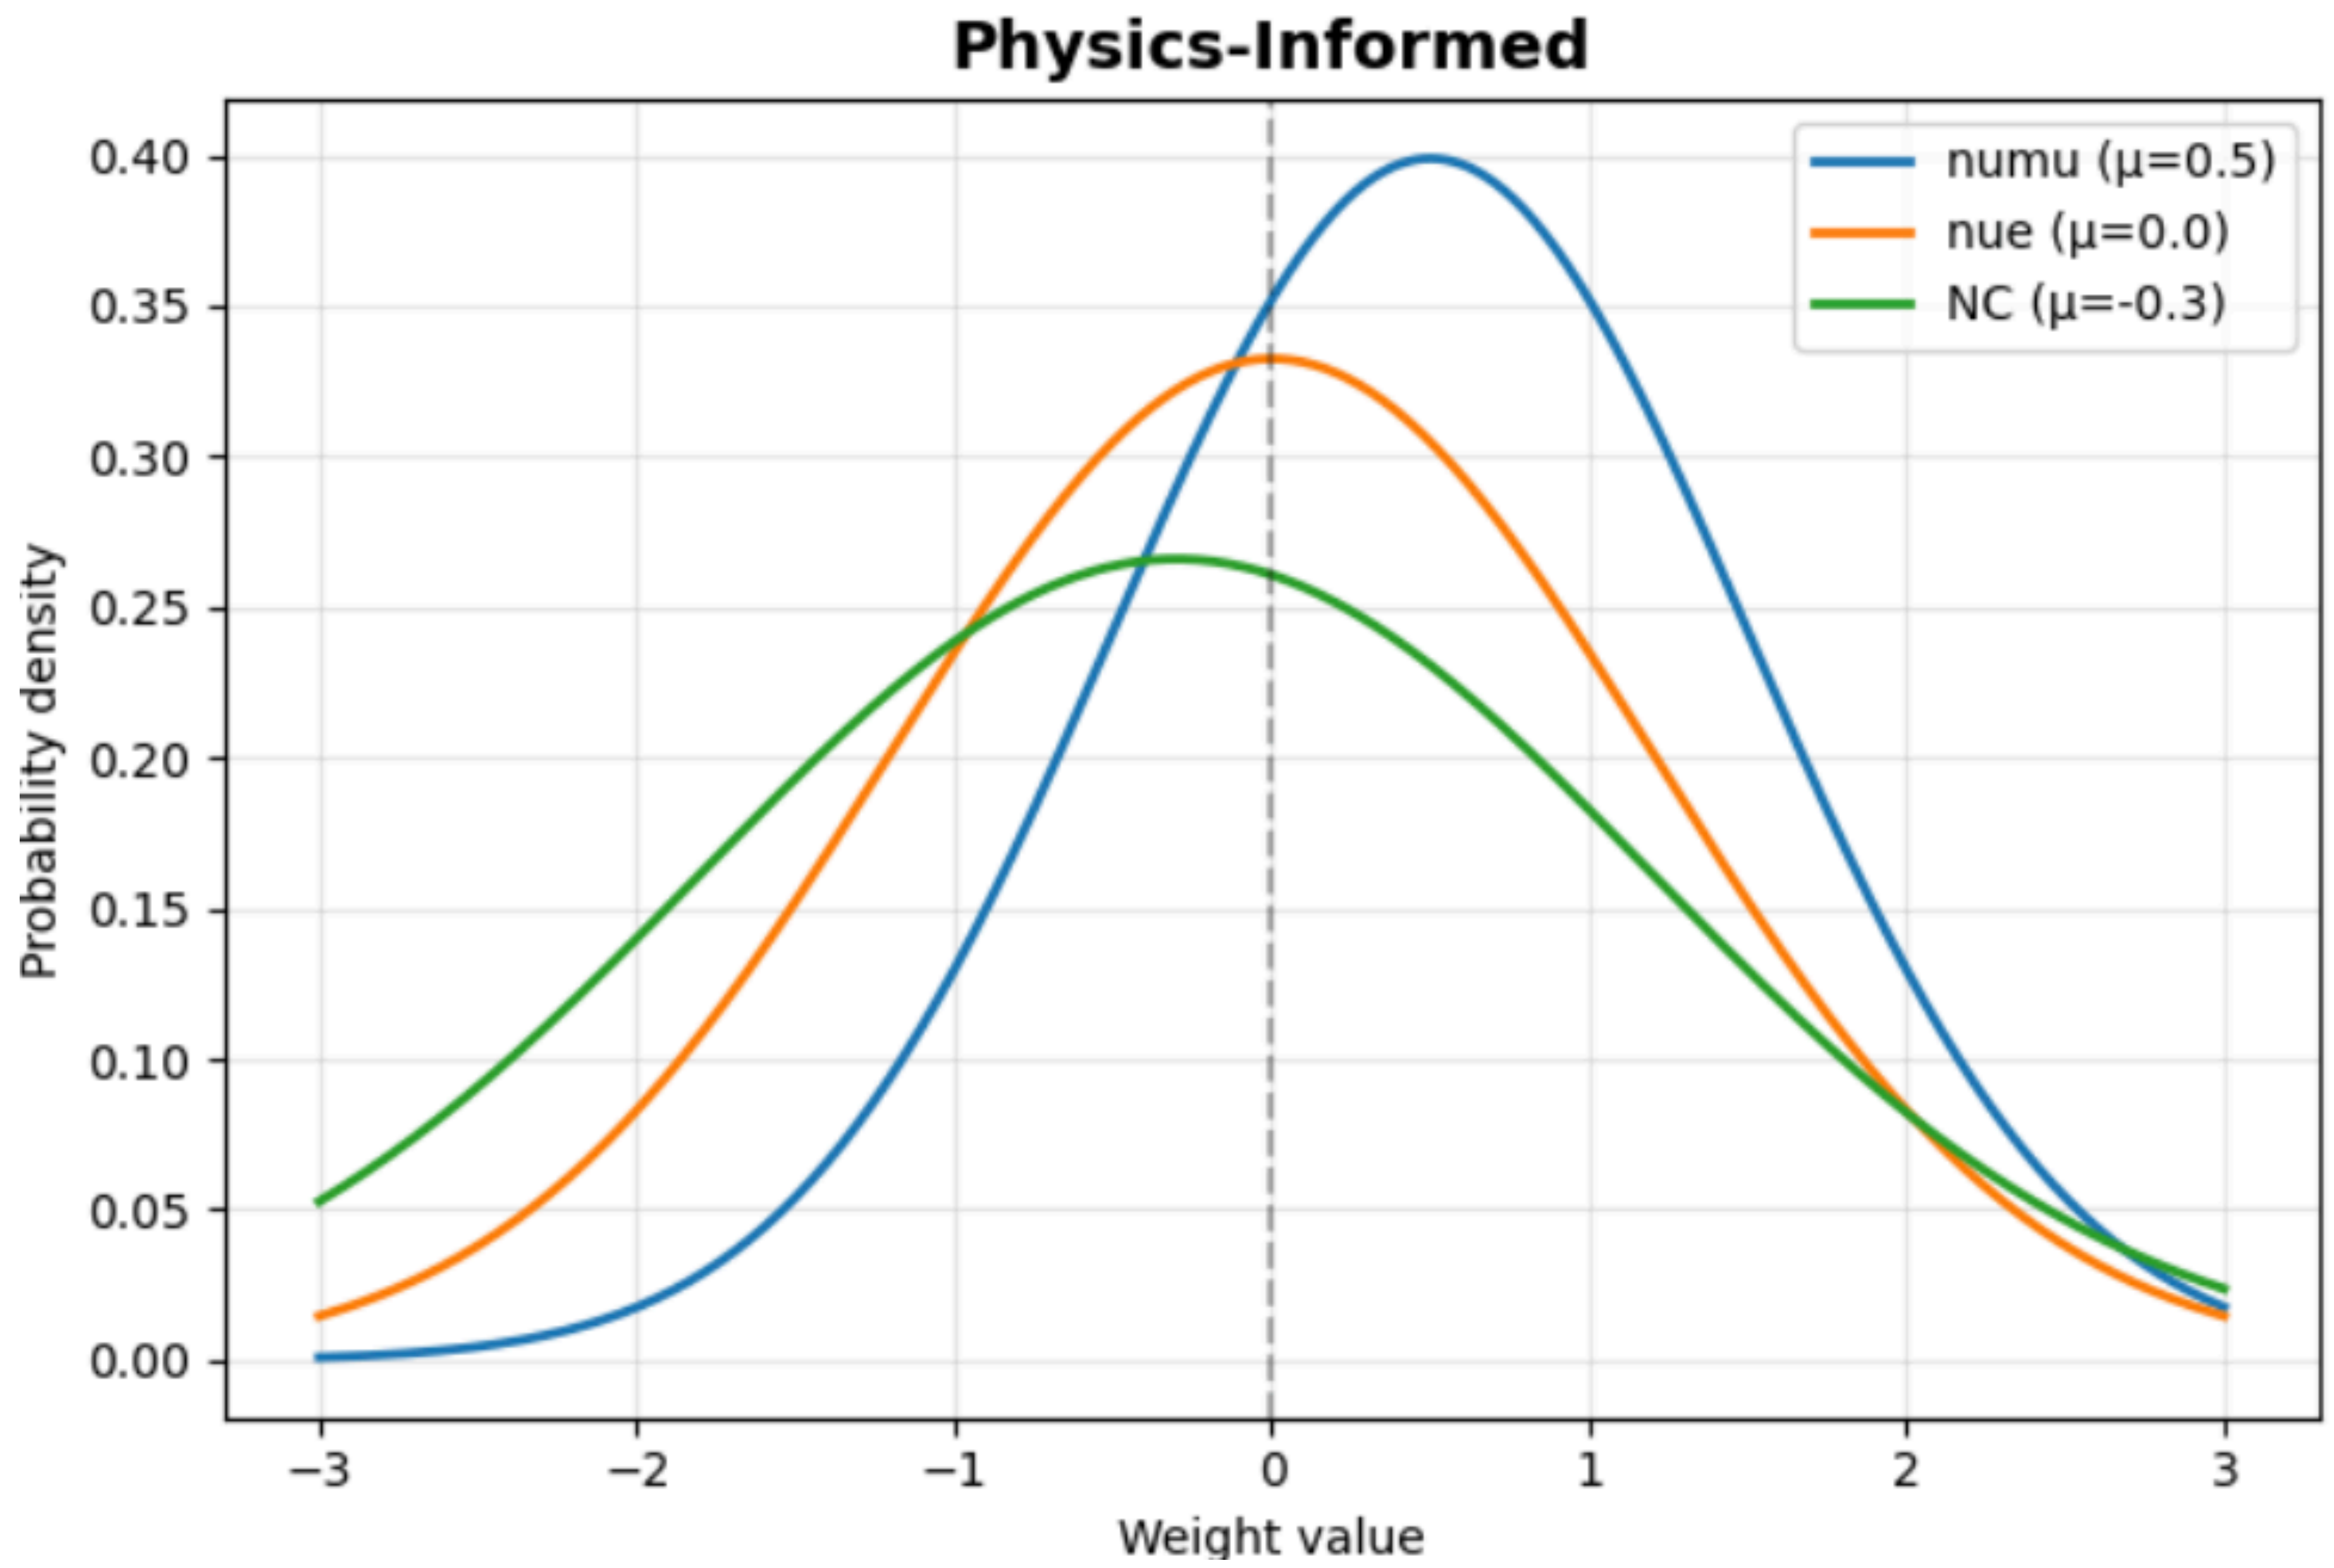

Class muon neutrino
- Produces long, straight tracks
- Clear directional signature

→ Prior mean: slightly positive (expect to find these features)

→ Prior std: small

Uncertainty Quantification with BNN

Can I introduce physics knowledge, a.k.a Physics-informed prior?

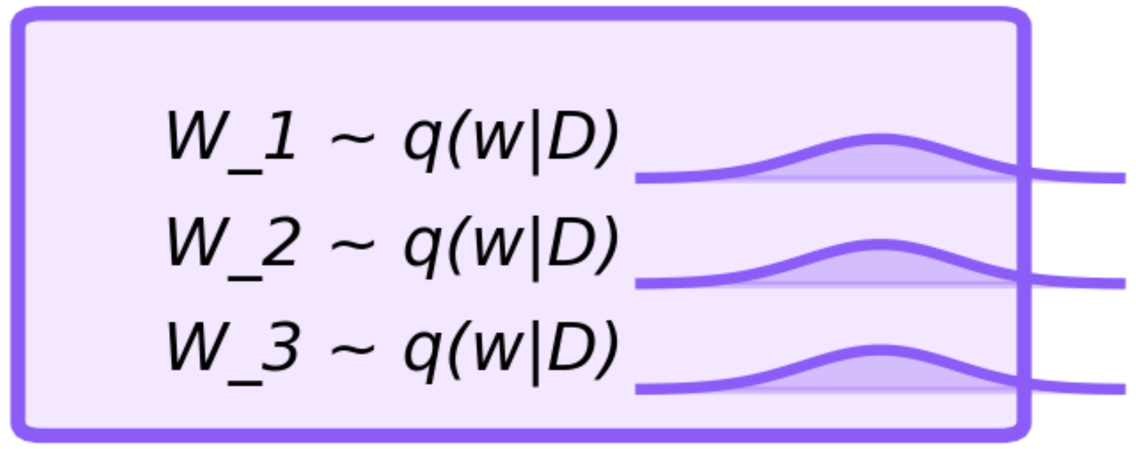

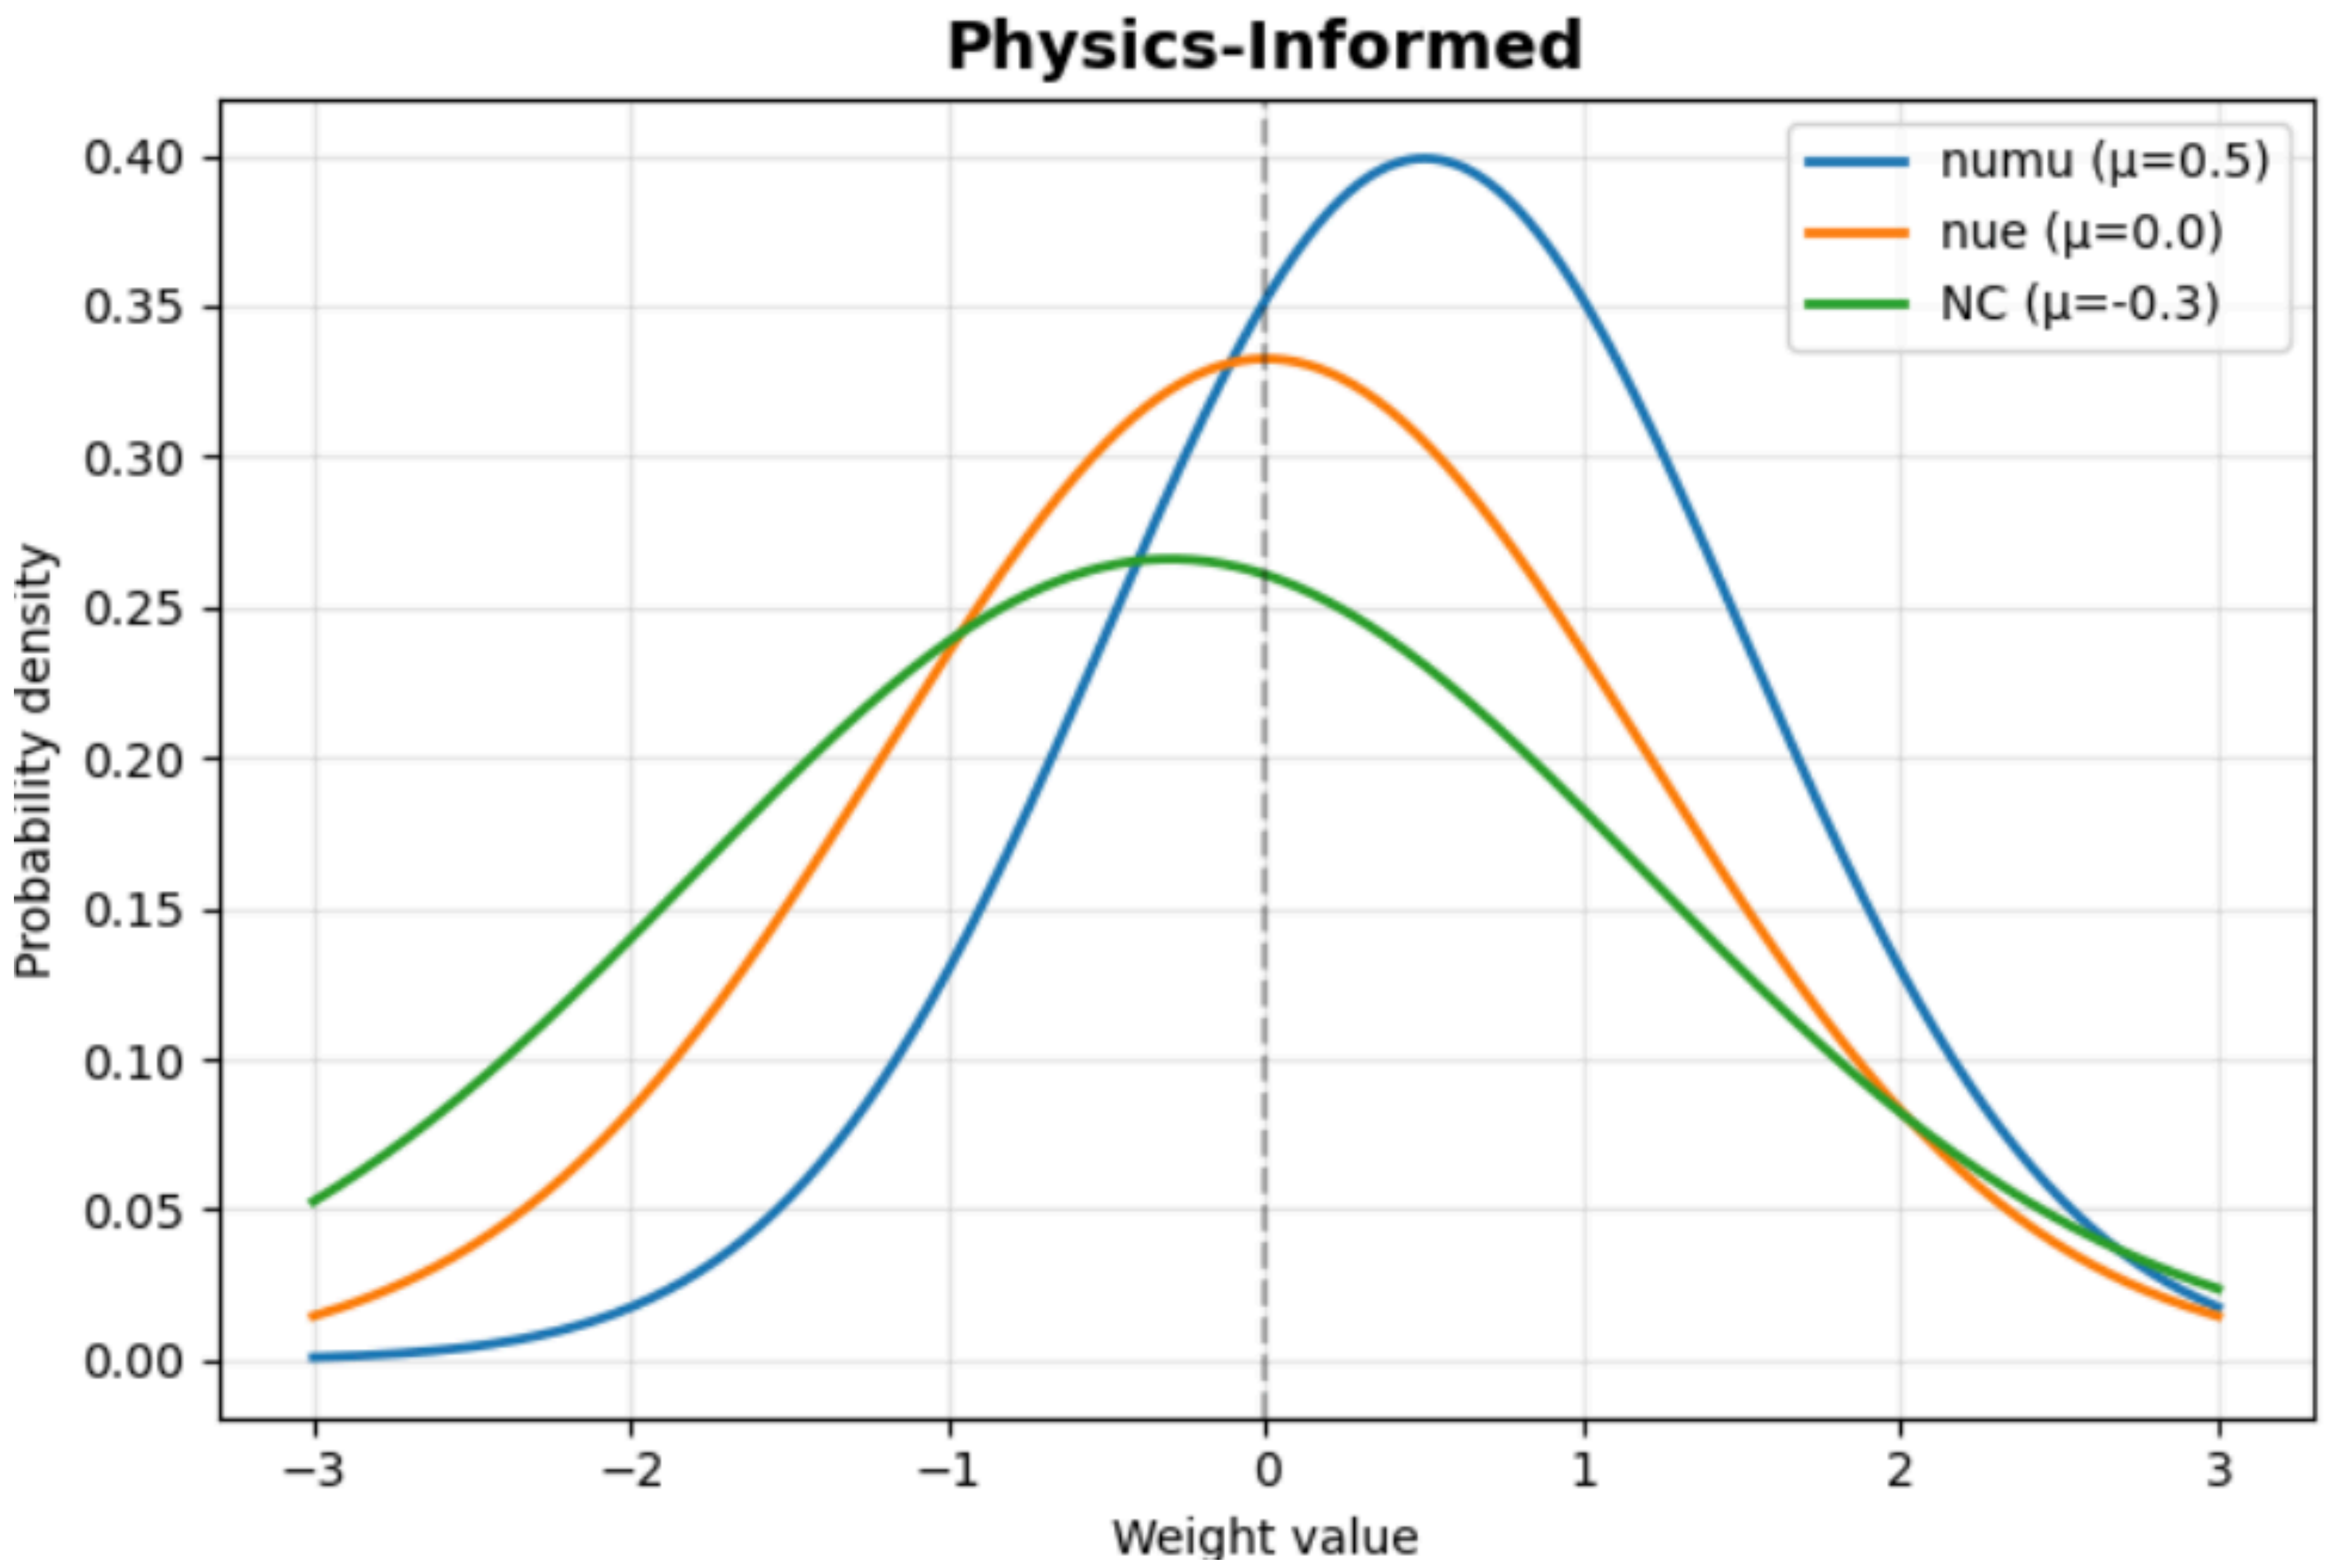

Class electron neutrino

- Produces electromagnetic showers
- Wider, more diffuse energy deposition

→ Prior mean: neutral (different topology than numu)

Uncertainty Quantification with BNN

Can I introduce physics knowledge, a.k.a Physics-informed prior?

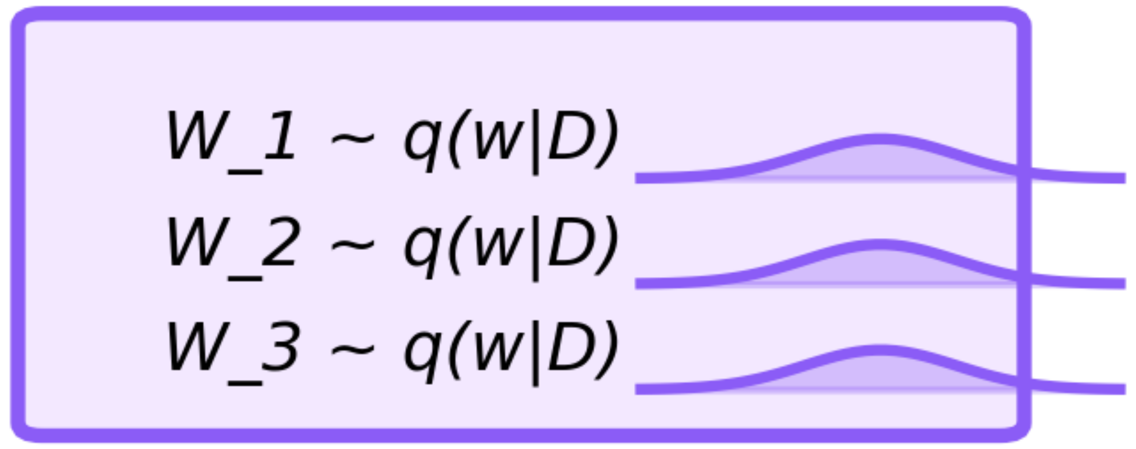

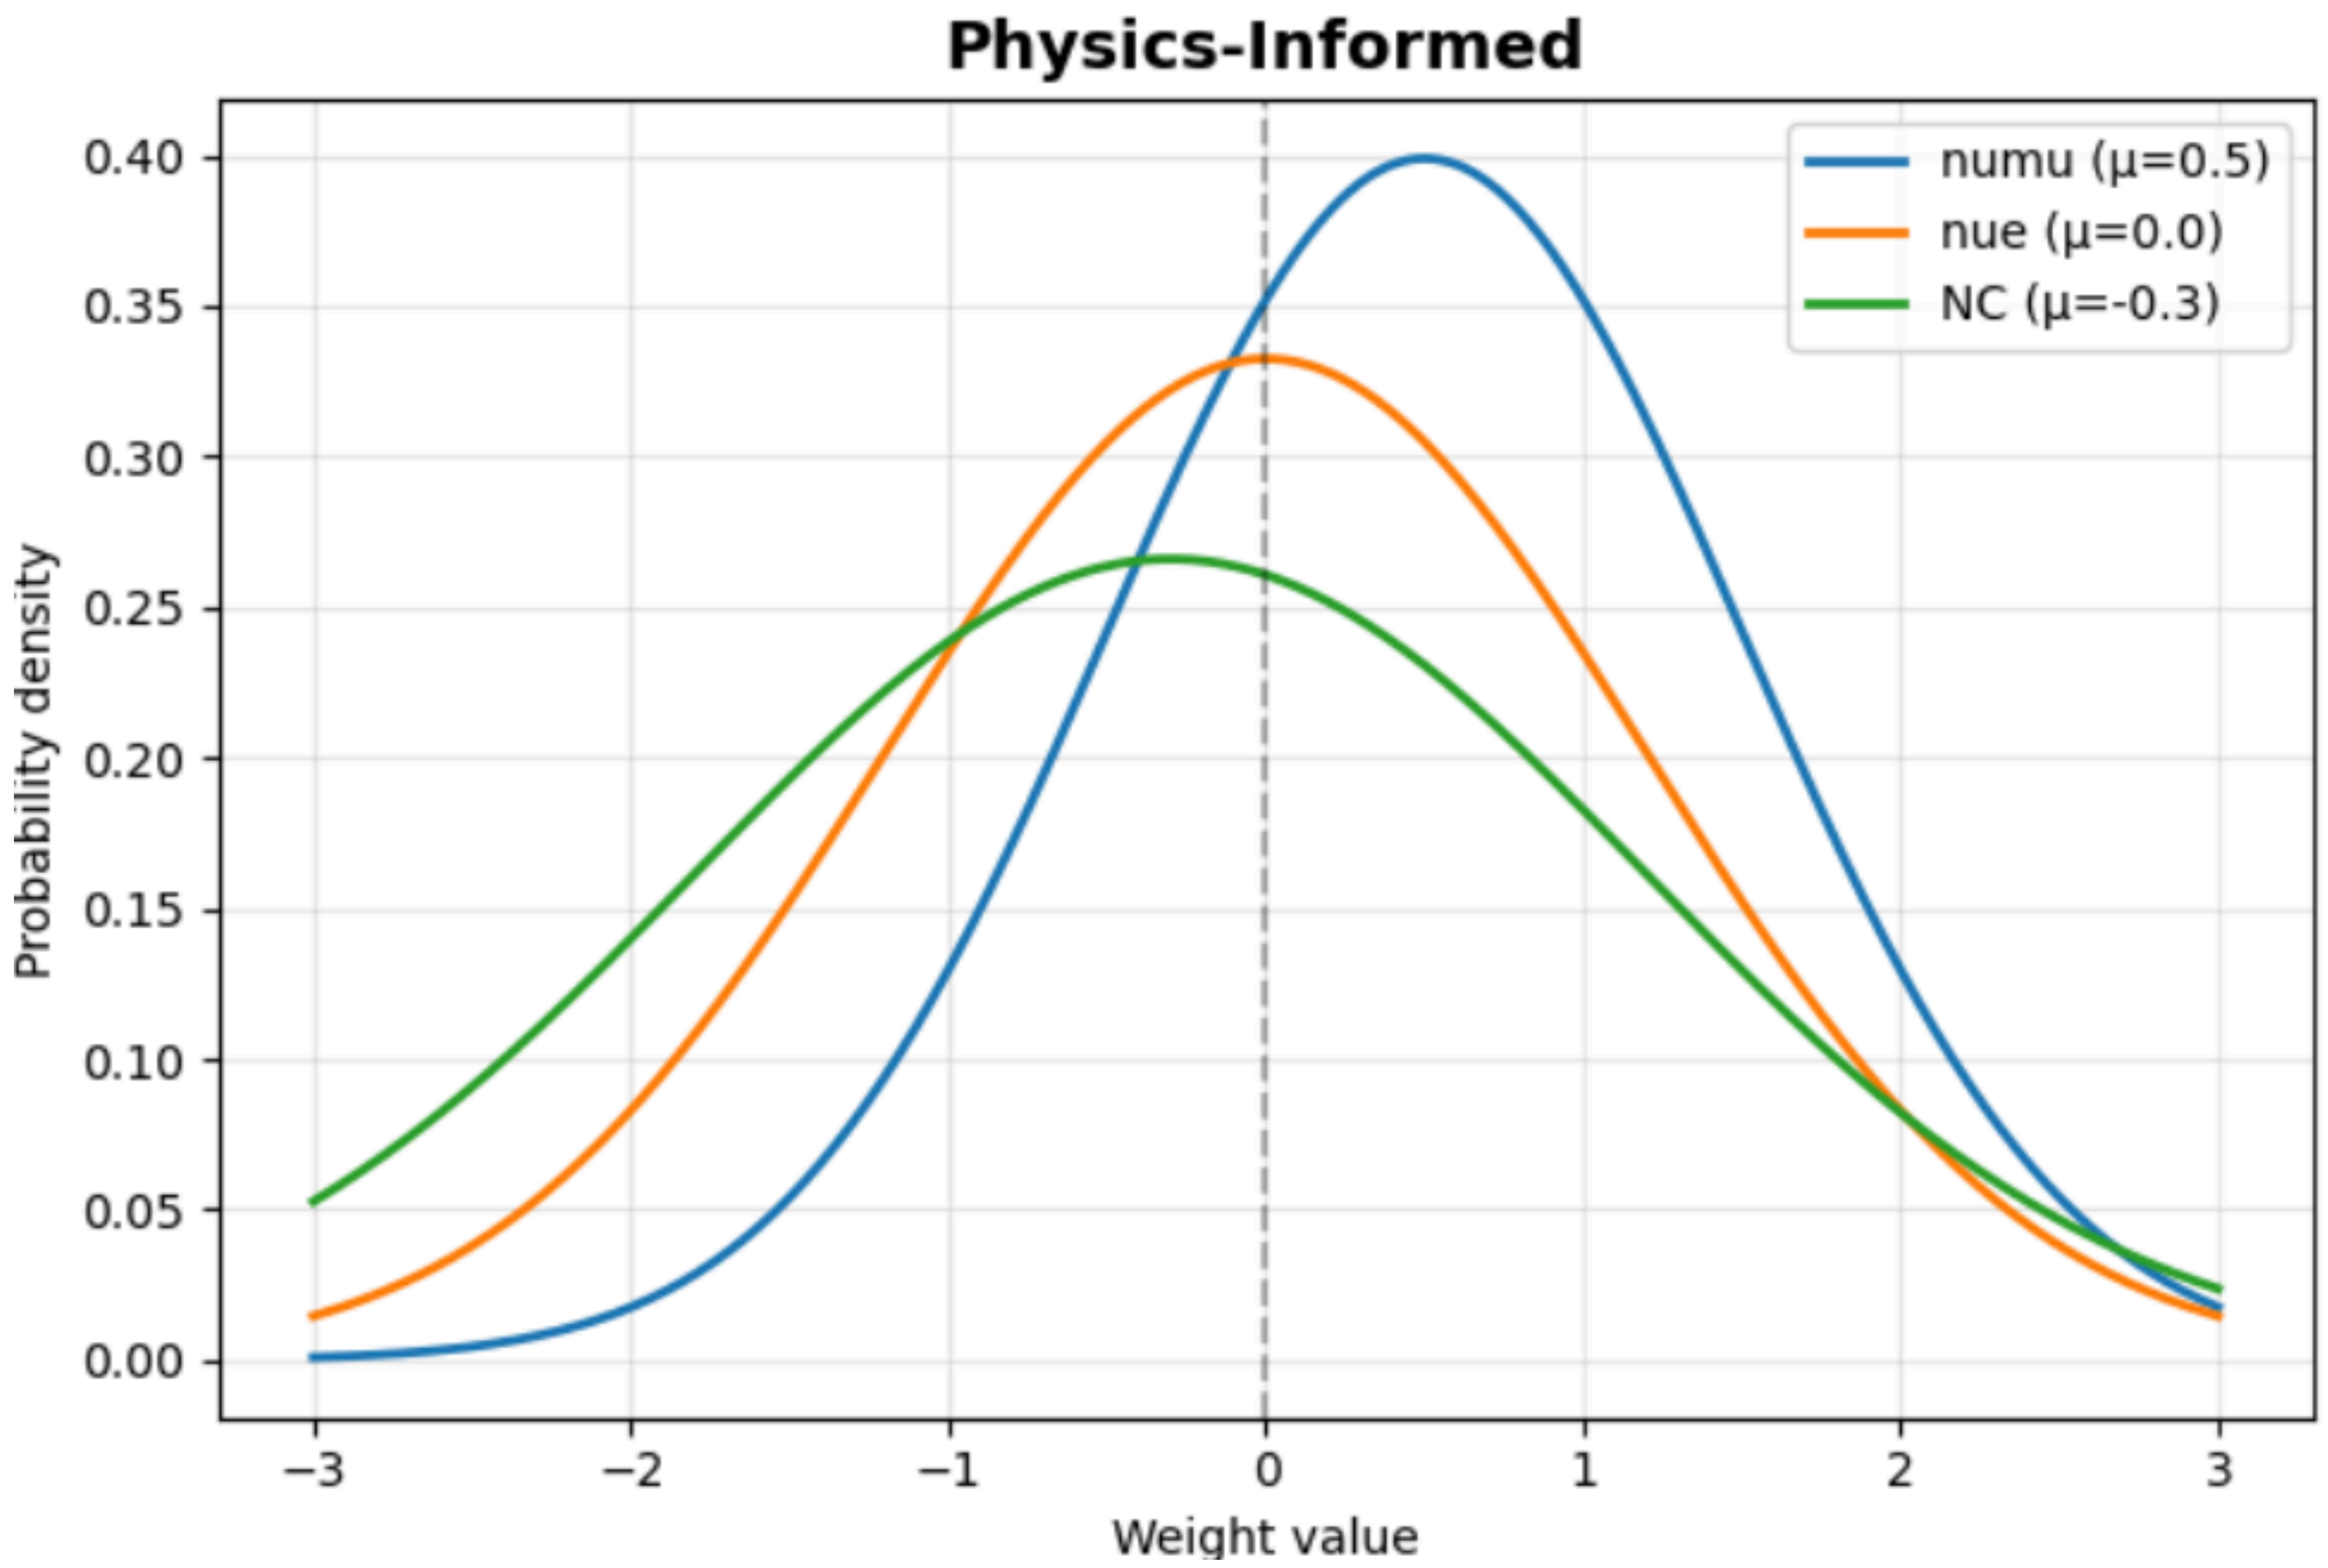

Class NC -

- No charged lepton (more ambiguous)
- Varied topologies
- Generally lower visible energy
- Often confused with others

A. Higuera, NPML 2025 University of Tokyo

Uncertainty Quantification with BNN

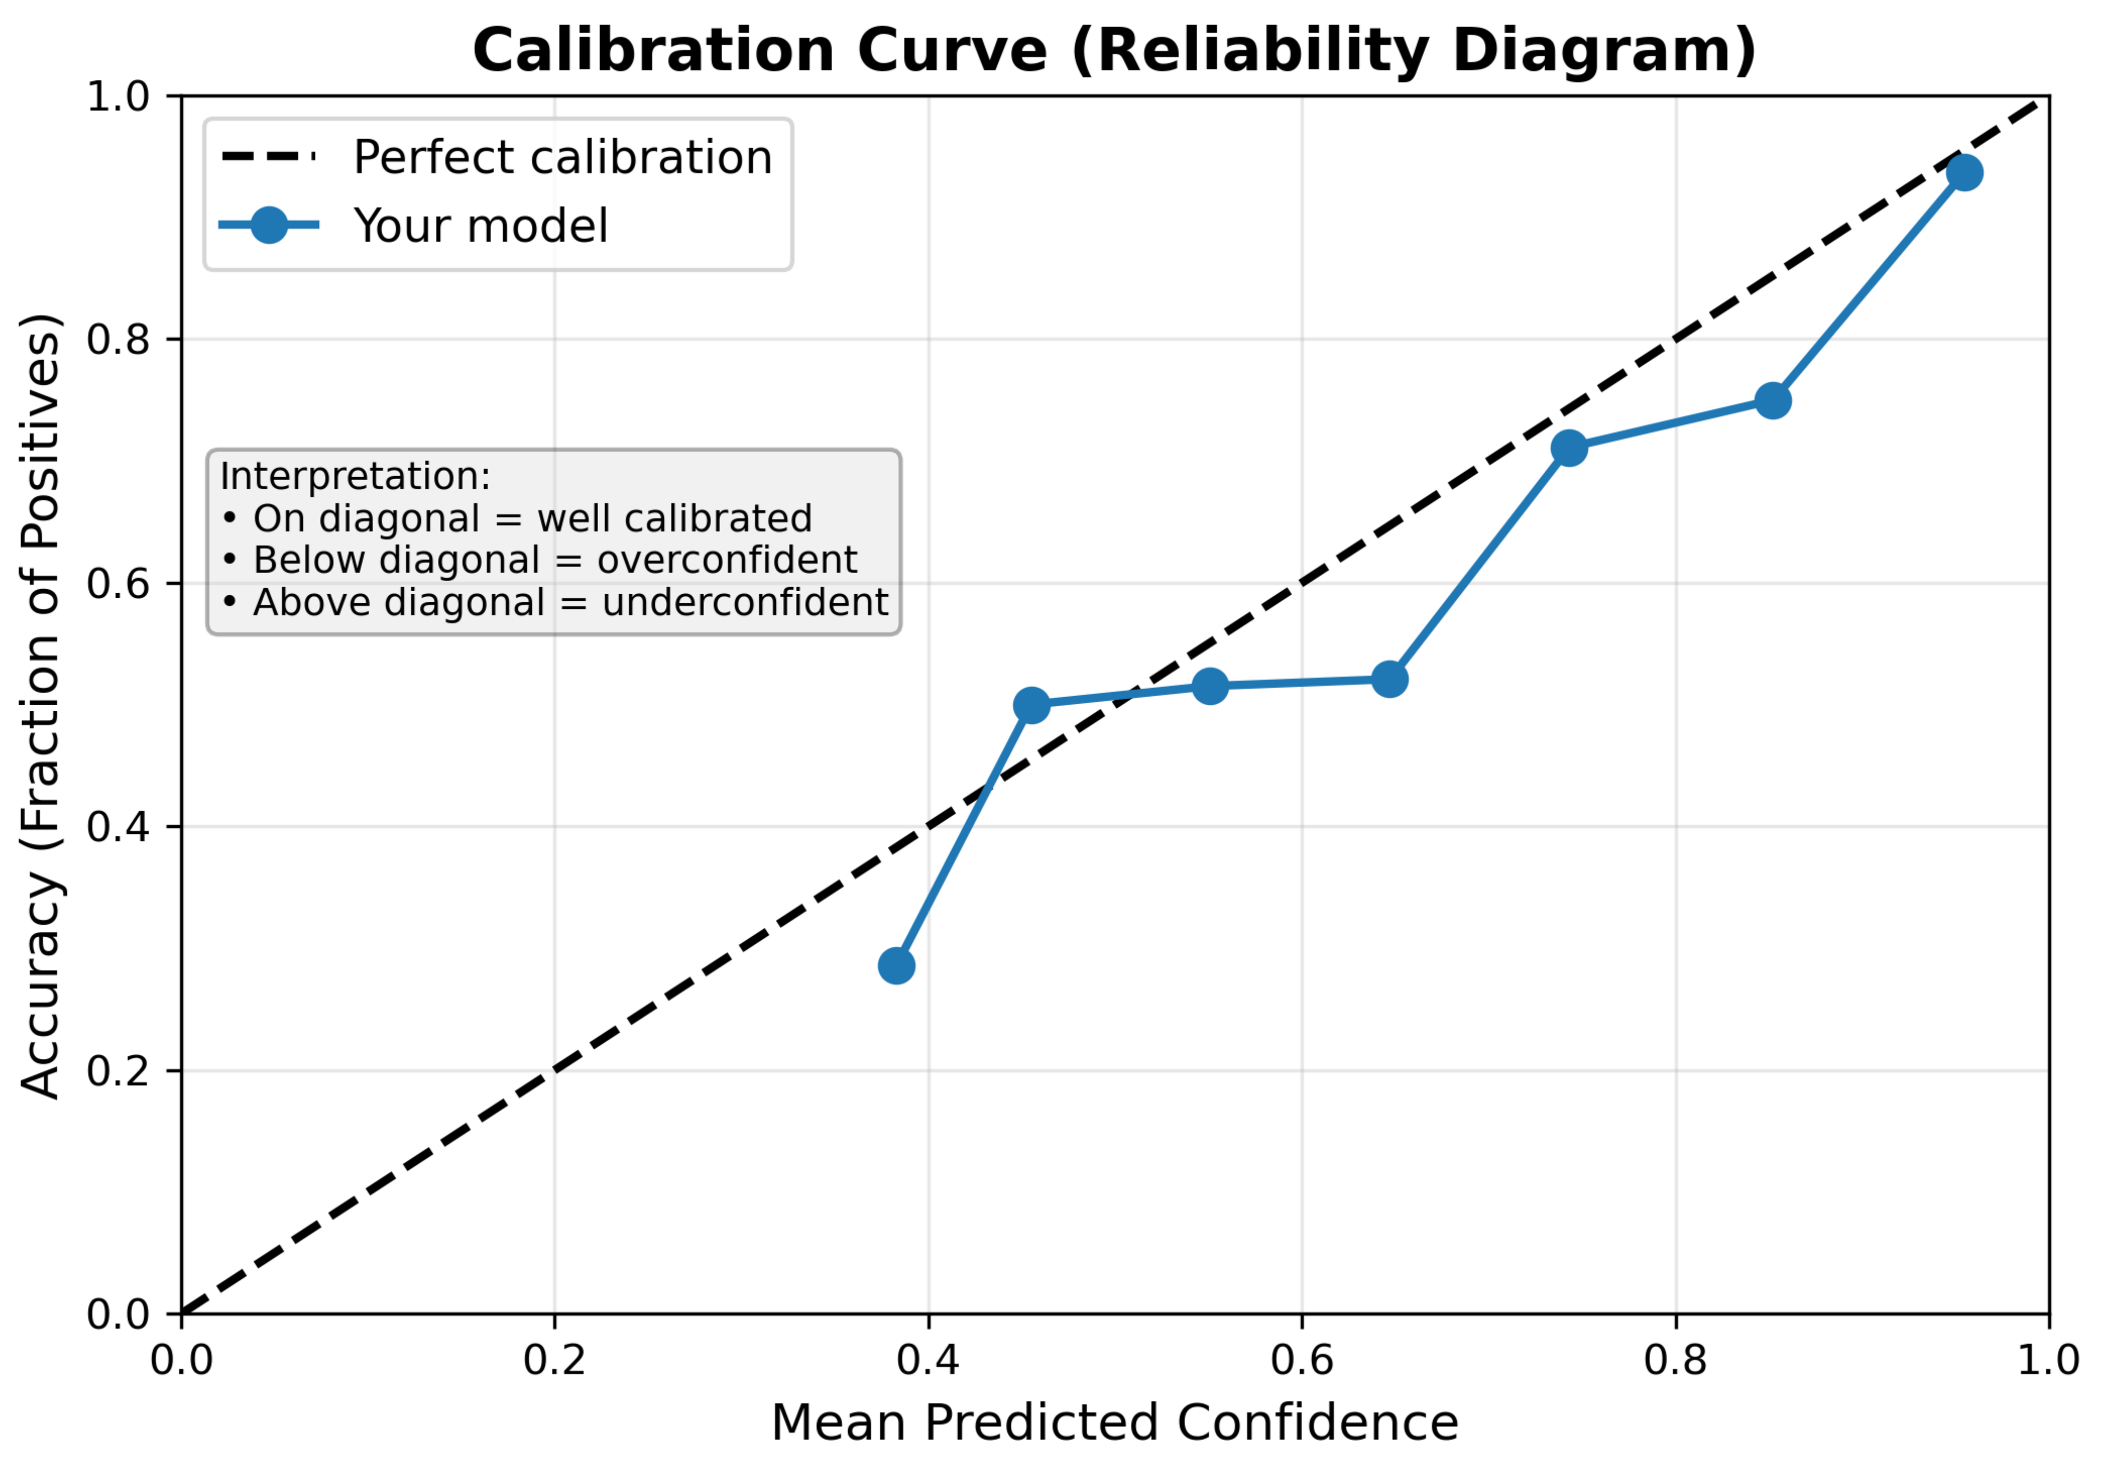

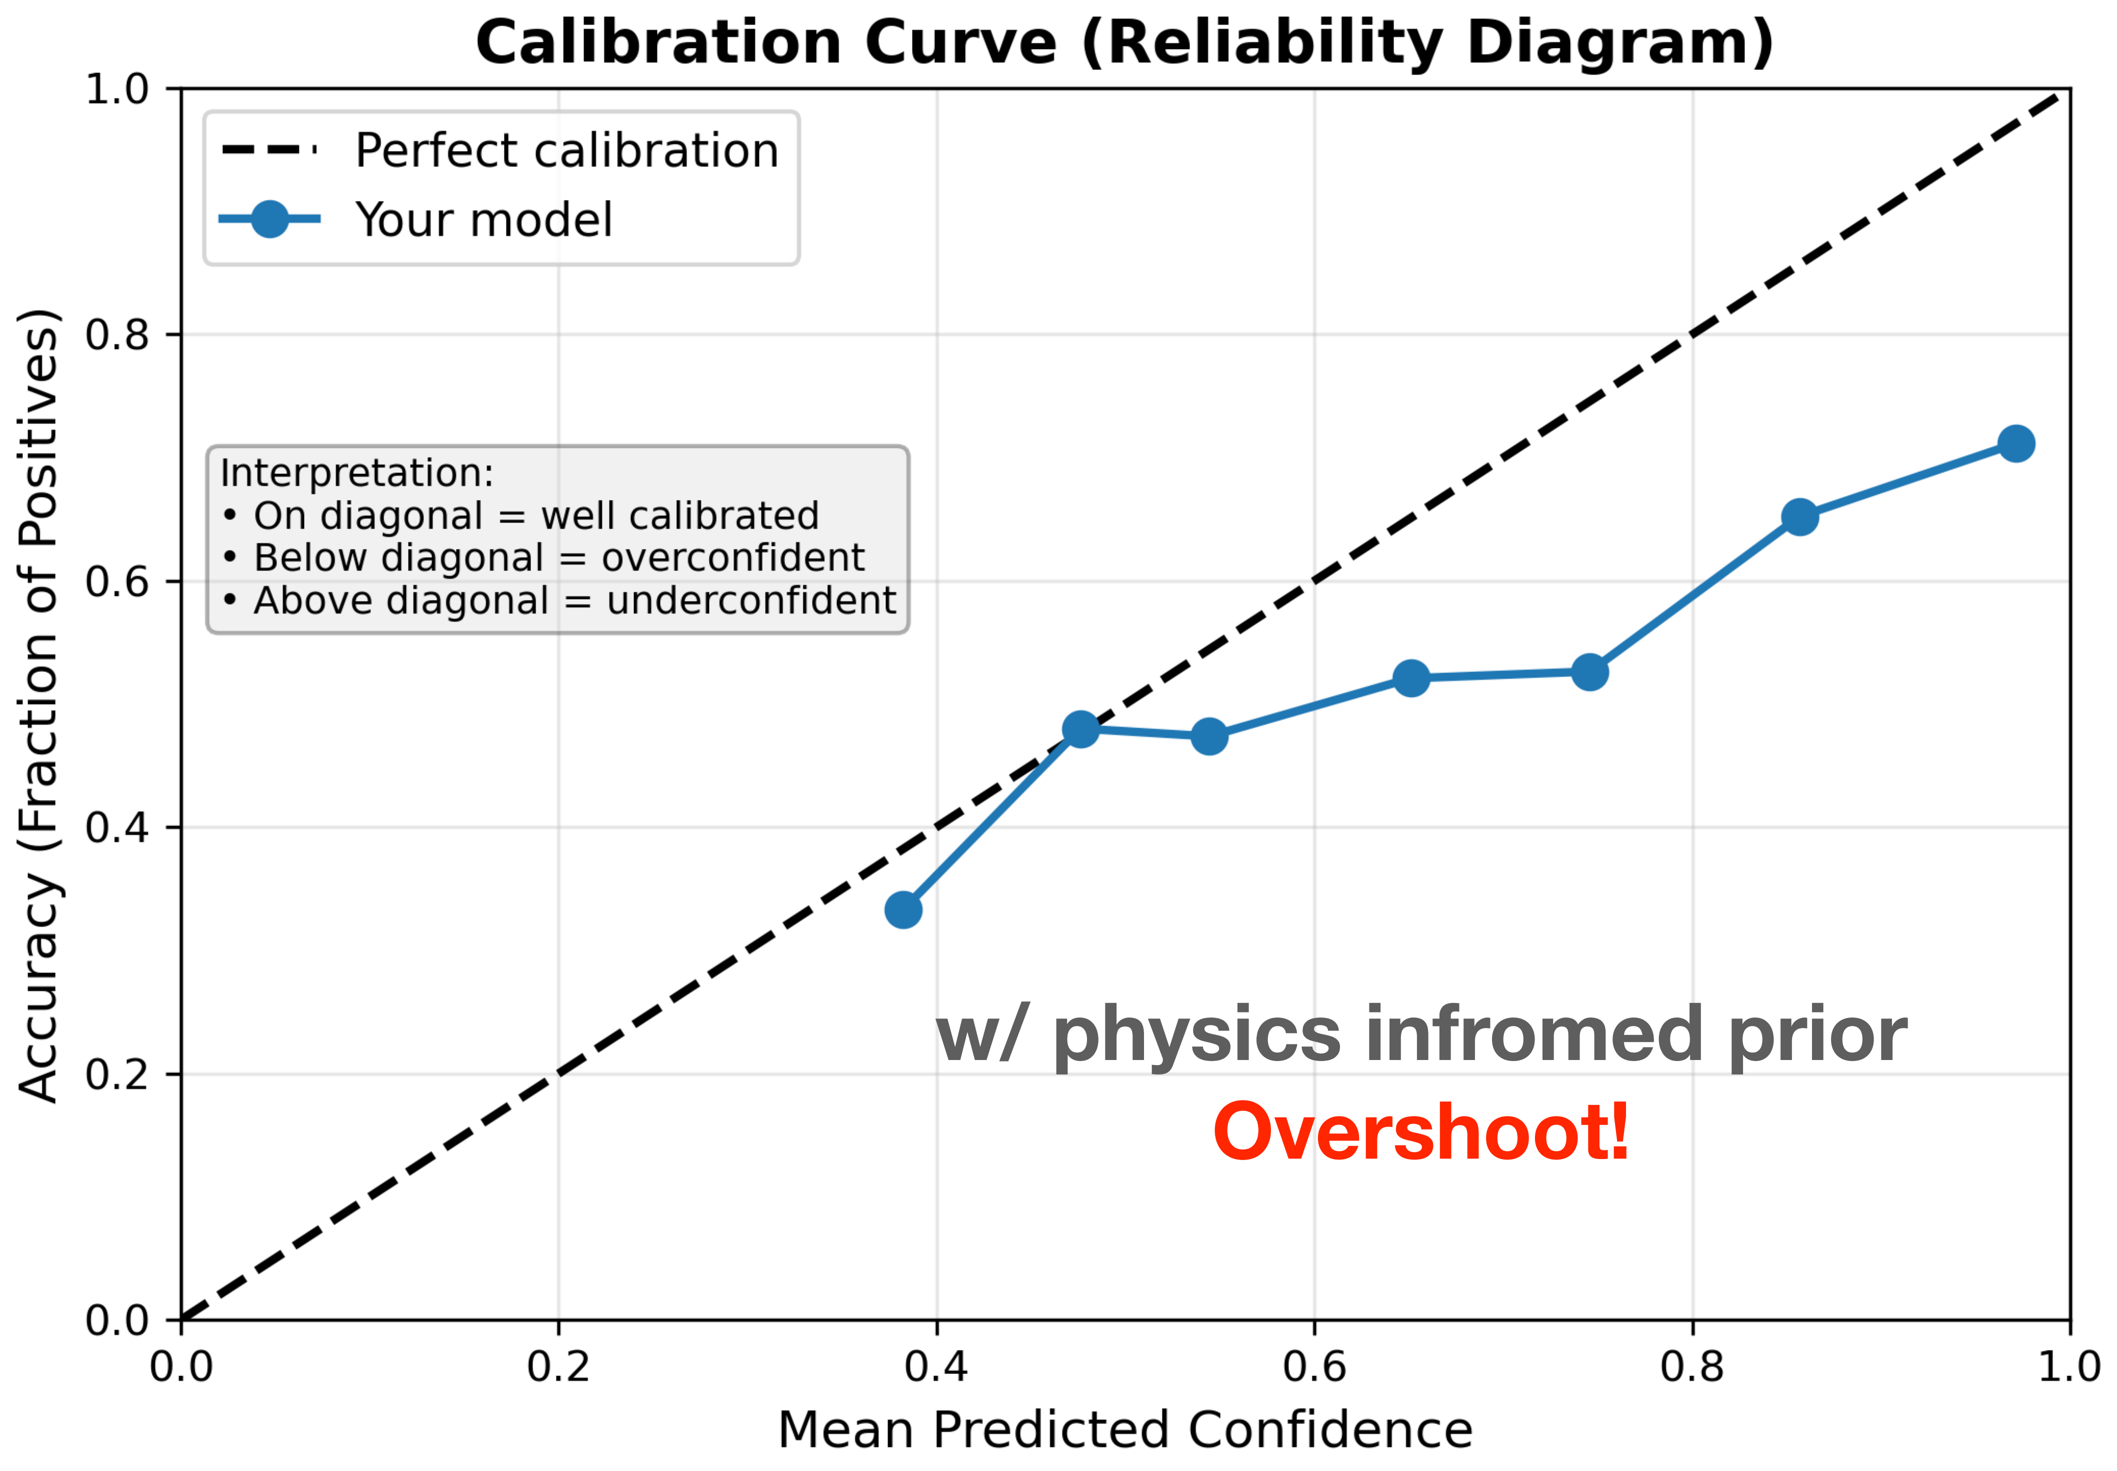

Uncertainty Quantification with BNN

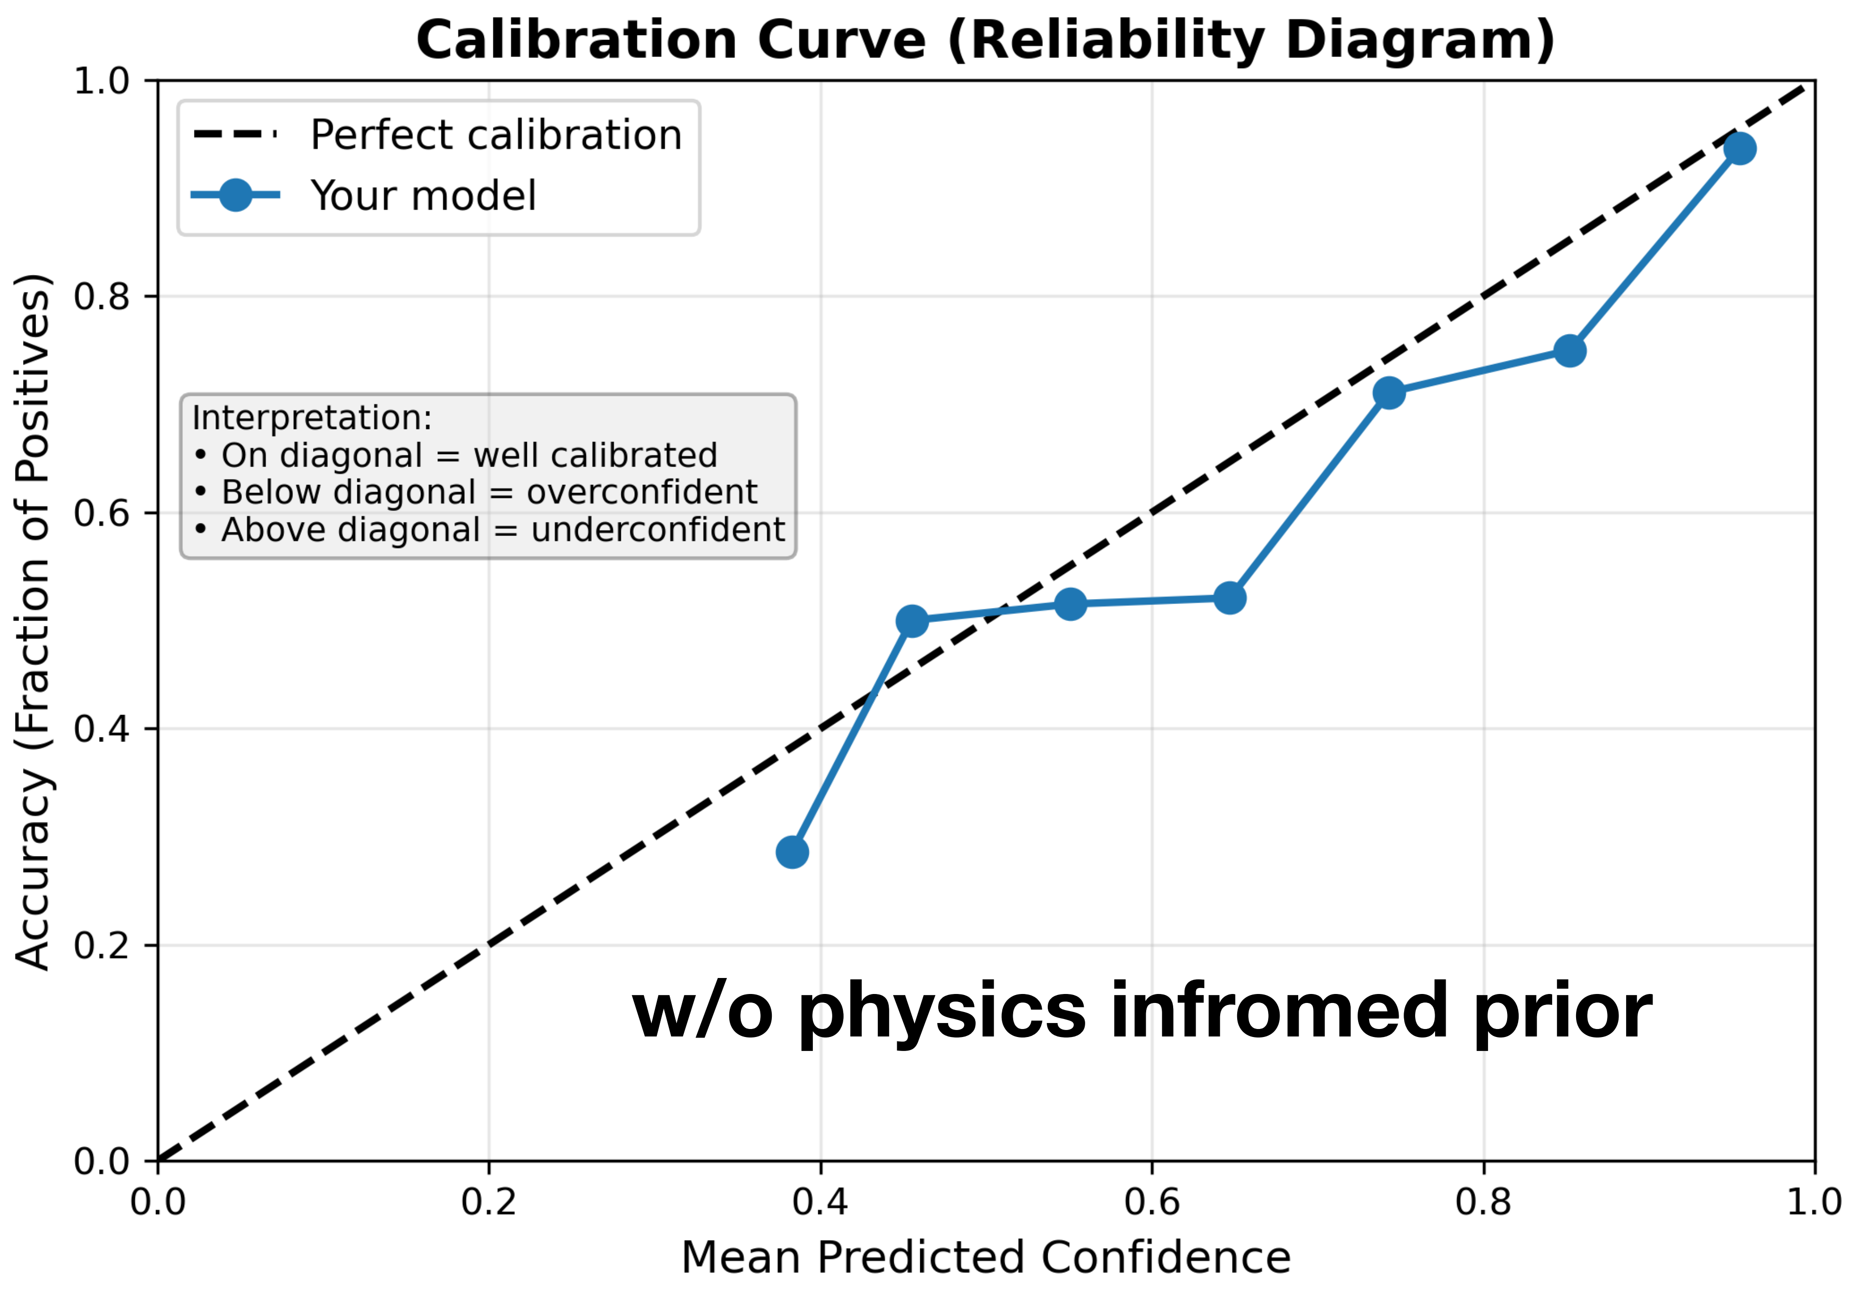

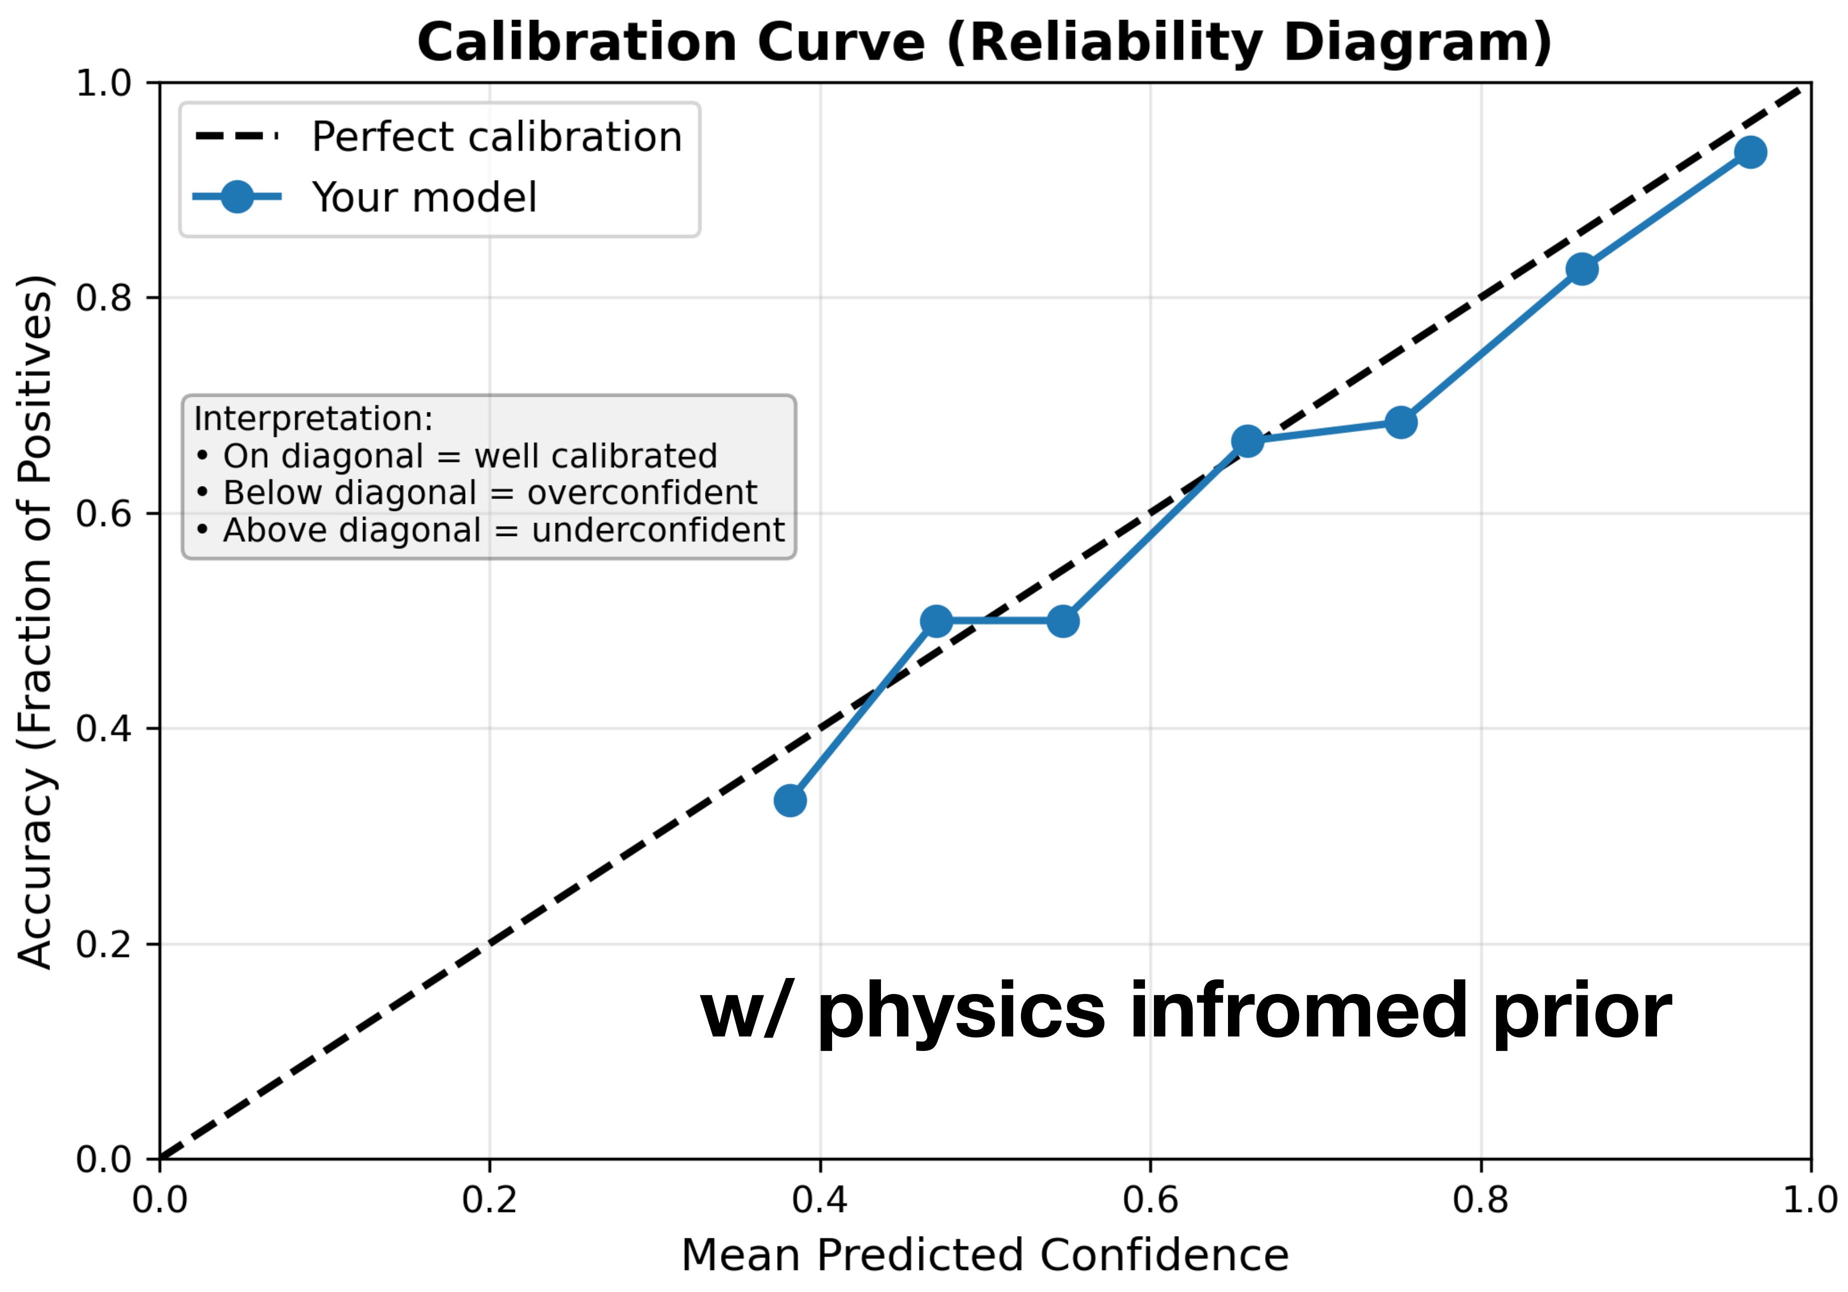

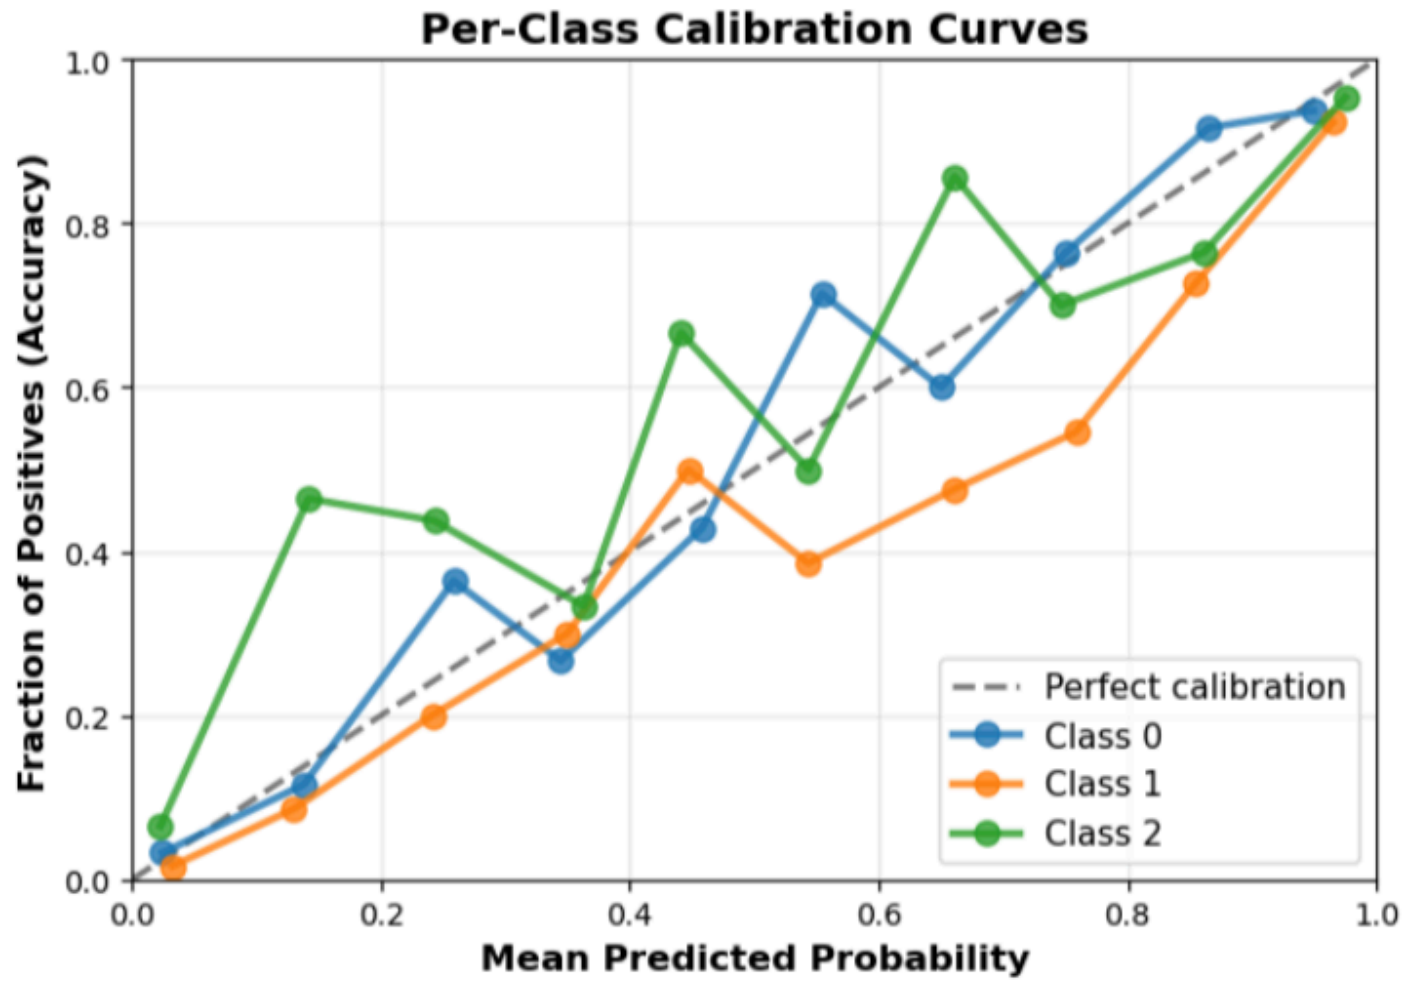

Comments

- Uncertainty quantification is essential for reliable ML predictions in high-stakes domains
- Formal frameworks exist to decompose uncertainty into epistemic and aleatoric components,
- `softmax` is not a probability or measurement of model confidence
- Bayesian Neural Networks provide an interpretable, formal approach to include UQ
- Mature toolkits available: **TensorFlow Probability** **Pyro**, **Biltz**, **nflows**
- Choice of method depends on computational budget, architecture, and required rigor
- If we can put error bars on neutrino flux, we can put them on our CNN/GCN/BDT/VLM!!!

Incorporating UQ in your ML Method

A. Higuera, NPML 2025 University of Tokyo

Extras

BNNs

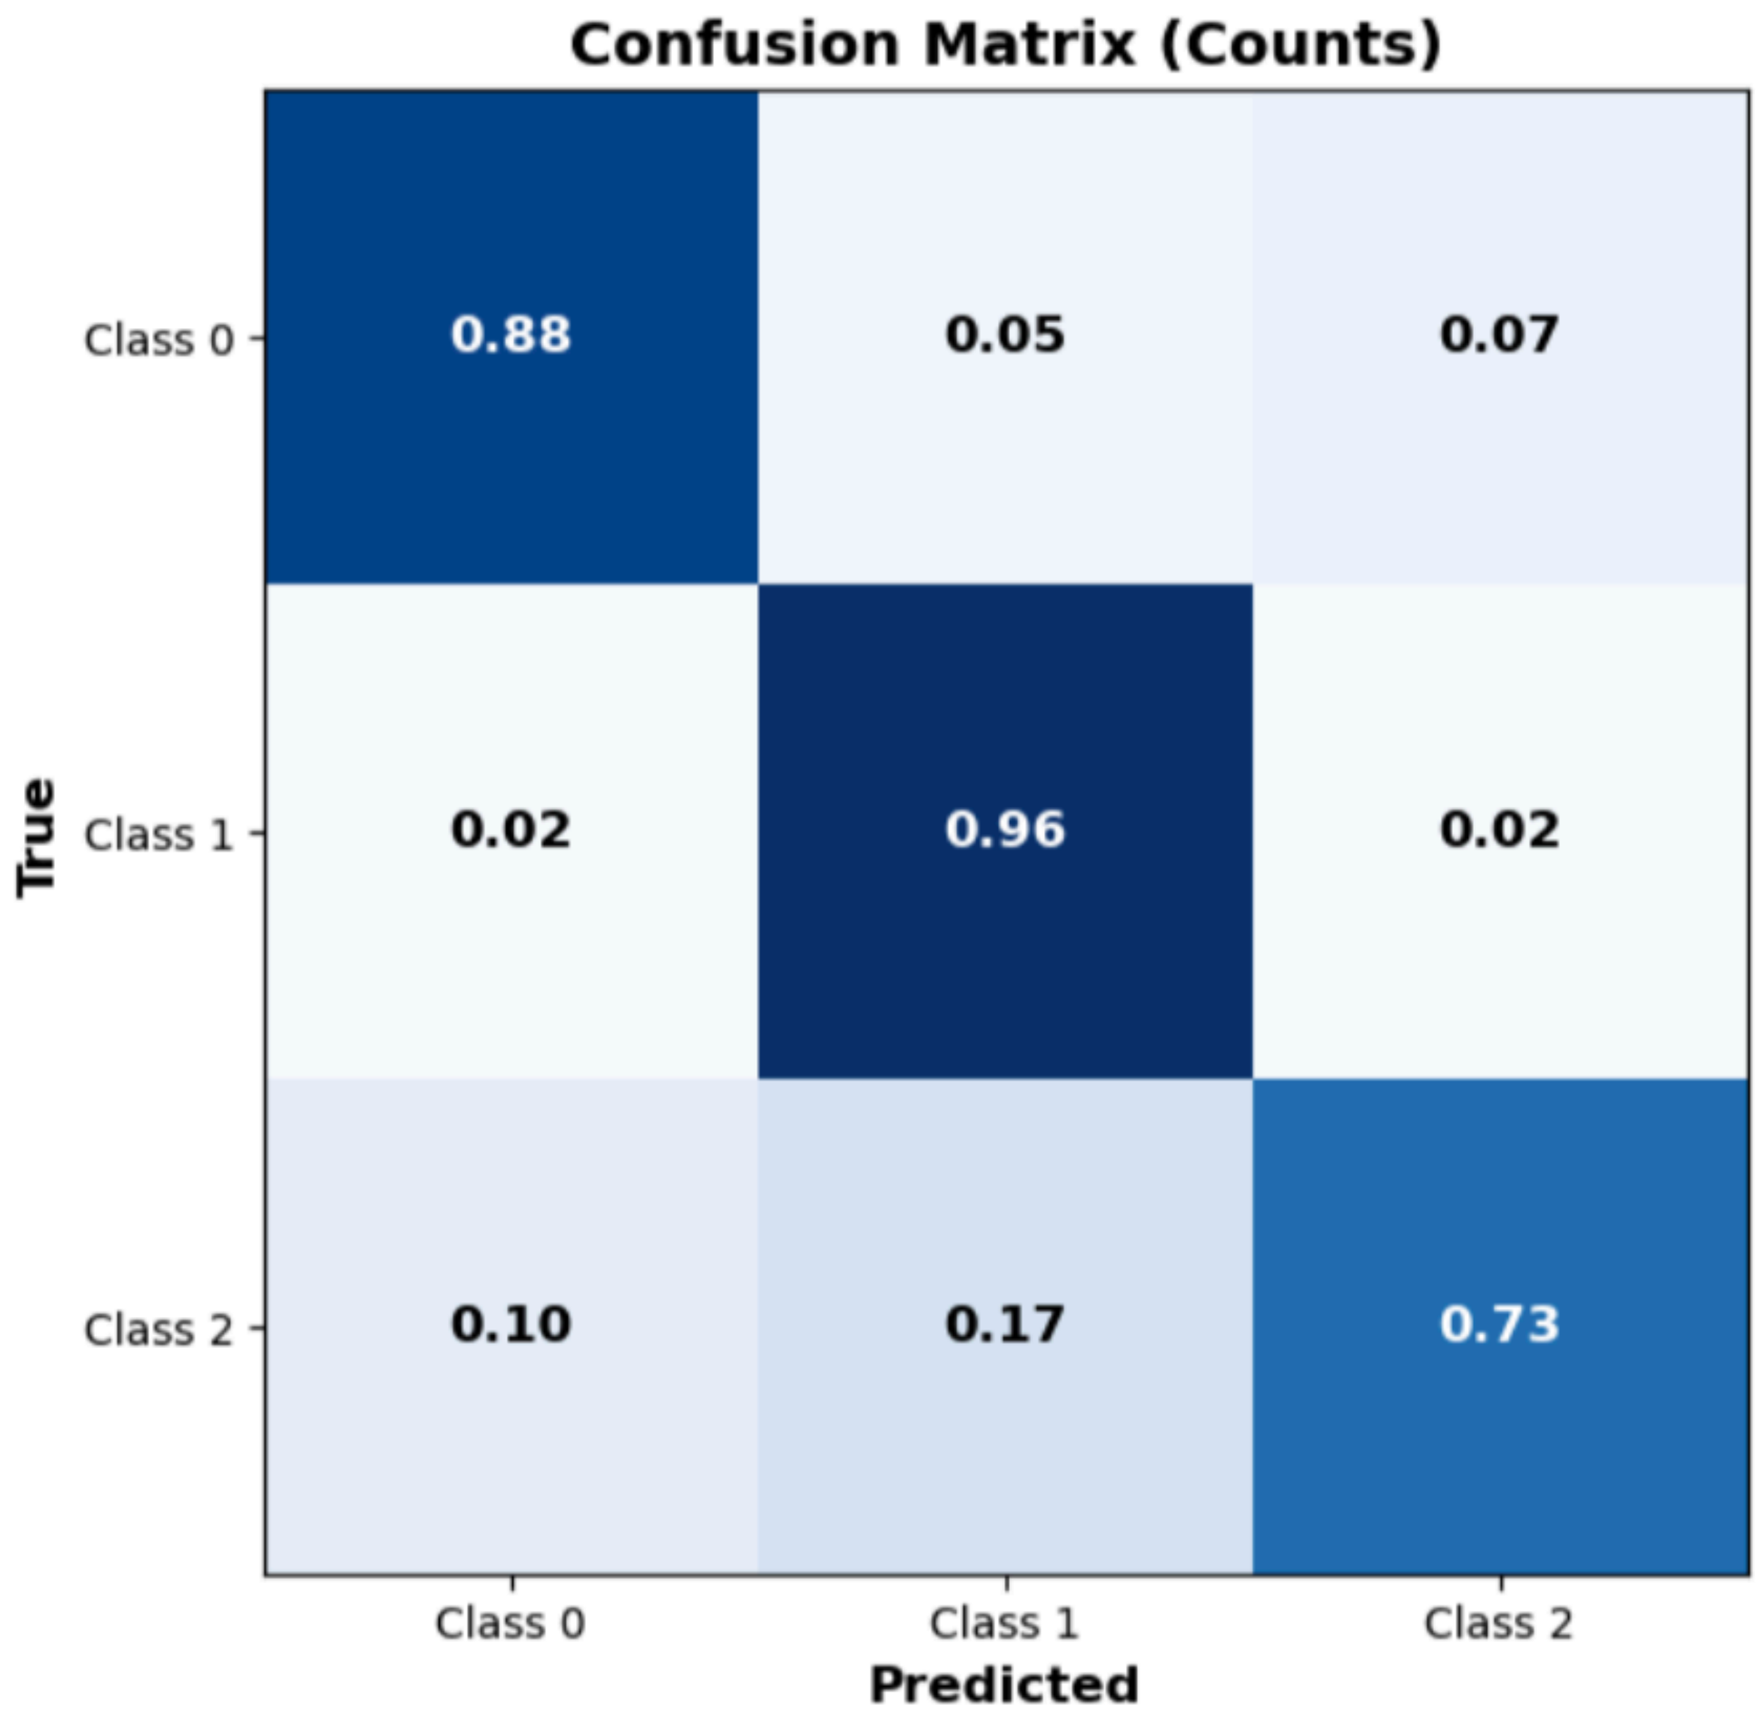

A. Higuera, NPML 2025 University of Tokyo