Bayesian Probabilistic Methods to Enable Cross-Cutting AI Research in Nuclear Science ("BUQ Phase 2")

Peter Jacobs

Lawrence Berkeley National Laboratory

University of California, Berkeley

2025 AI/ML PI Exchange Meeting
Nov. 19-20, 2025

BUQ Phase 2 project

This proposal requests funding to develop and implement general ML-based approaches to Bayesian probabilistic analysis methods, including uncertainty quantification, emulation, inference, and de–noising, for application to a broad range of Nuclear Physics (NP) research areas. These areas include the measurement of the mass and fundamental nature of the neutrino; study of the Quark-Gluon Plasma that filled the early universe; and mapping of natural and anthropogenic radiation environments. The proposal co–PIs are Nuclear Physicists with leading roles in each of these areas, and data scientists developing state–of–the–art ML-based methods for probabilistic Bayesian analysis.

PI: Peter Jacobs, LBNL

Co-PIs:

- LBNL: Brian Fujikawa, Alan Poon, Jayson Vavrek
- UC Berkeley/LBNL: Yury Kolomensky, Uros Seljak
- Duke: Simon Mak
- UC San Diego: Aobo Li
- Wayne State University: Chun Shen

https://sites.google.com/lbl.gov/bayesianuqproject/home

BUQ Phase 2: challenges

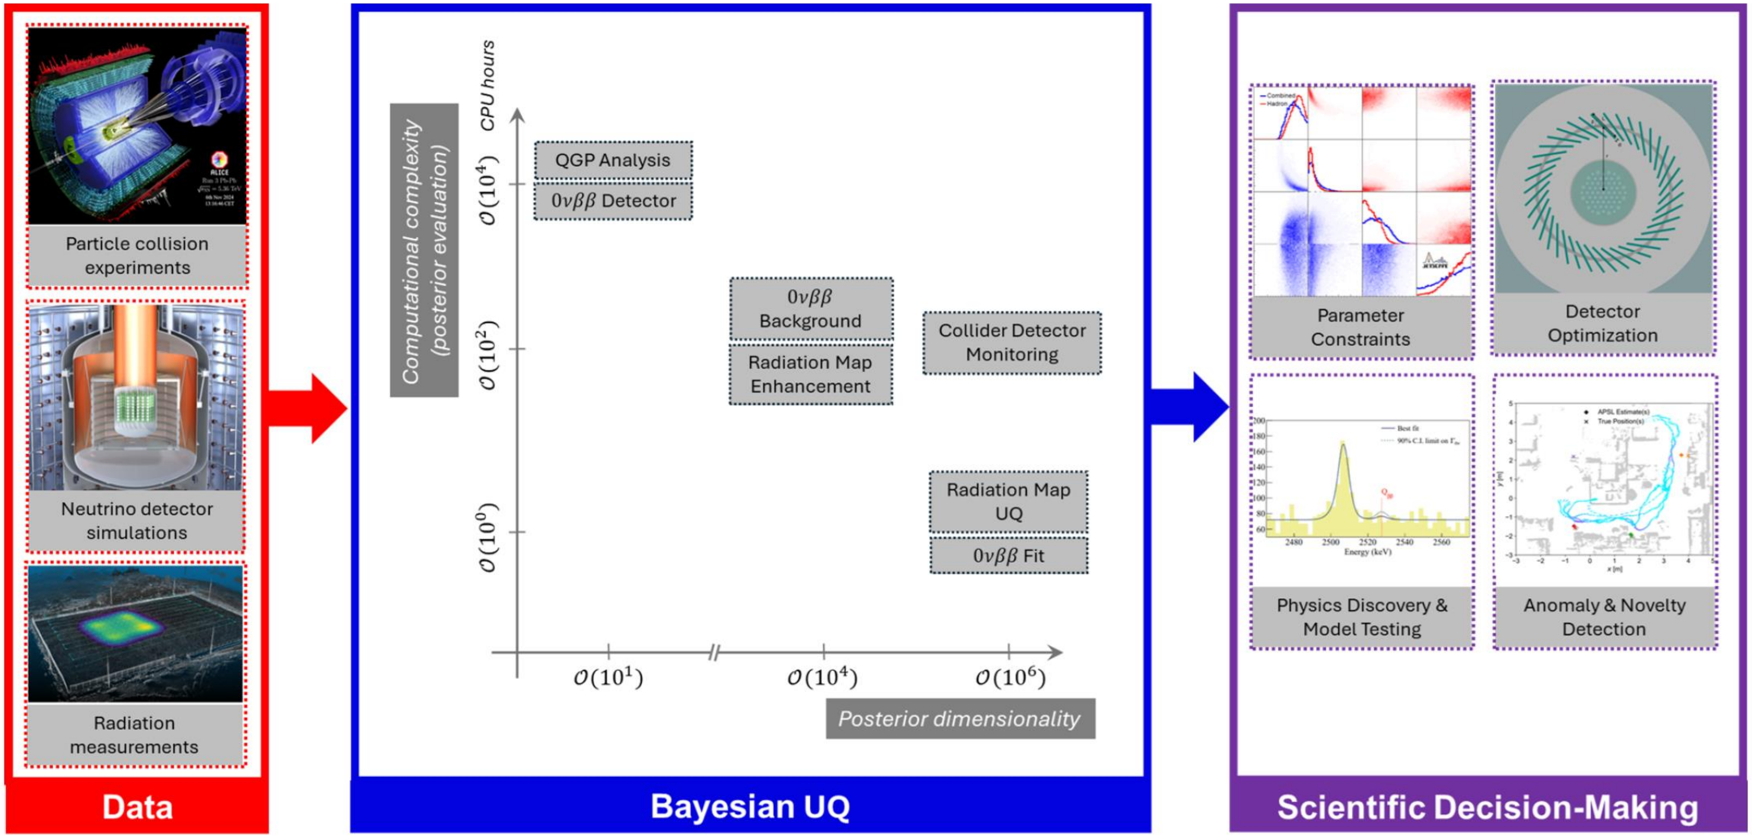

https://sites.google.com/lbl.gov/bayesianuqproject/home

BUQ Phase 2: ML-based solutions

	$0\nu\beta\beta$ Detector Optimization	$0\nu\beta\beta$ Background & Fit	QGP Collider Monitoring	QGP Analysis	Radiation Map Enhancement	Radiation Map UQ
Bayesian Transfer Learning				✓	✓	
Bayesian Multi-Fidelity Learning	✓ ○			✓	✓	
Langevin Monte Carlo		✓ ○		✓		✓ ○
Bayesian Manifold Learning	○			○		
Bayesian Optimization	○					○
Boundary-Informed Surrogates	○			○		
Bayesian Image Change Detection			○			

✓: BUQ Phase 1 ○: BUQ Phase 2

https://sites.google.com/lbl.gov/bayesianuqproject/home

Selected results: neutrino physics

LEGEND: ton-scale detector to search for neutrino-less double beta decay

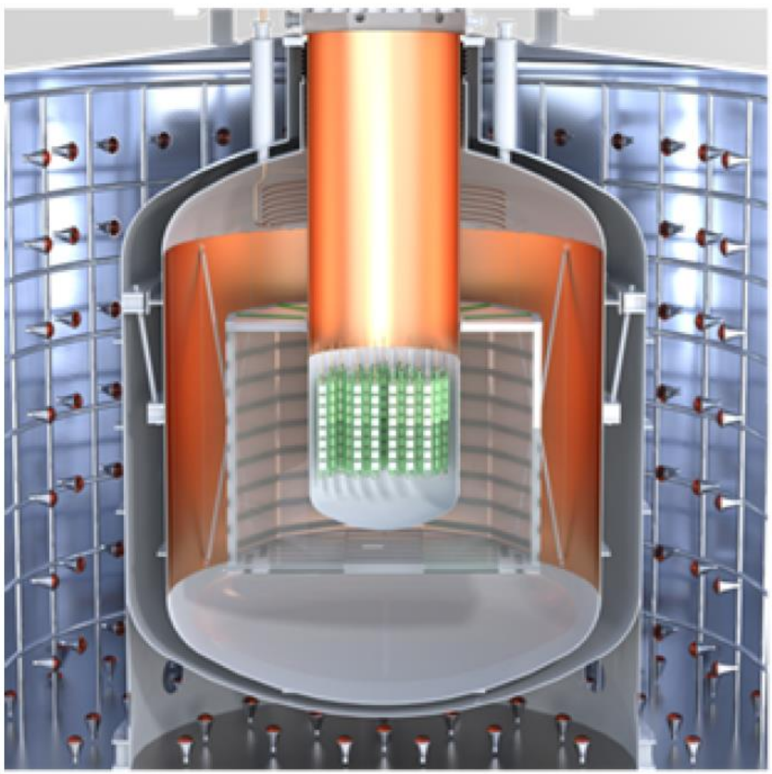

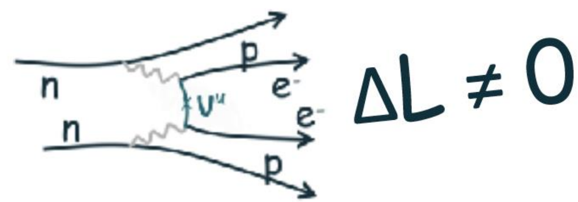

ReSUM

Optimize design of neutron moderator to minimize cosmogenic background

arXiv:2410.03873

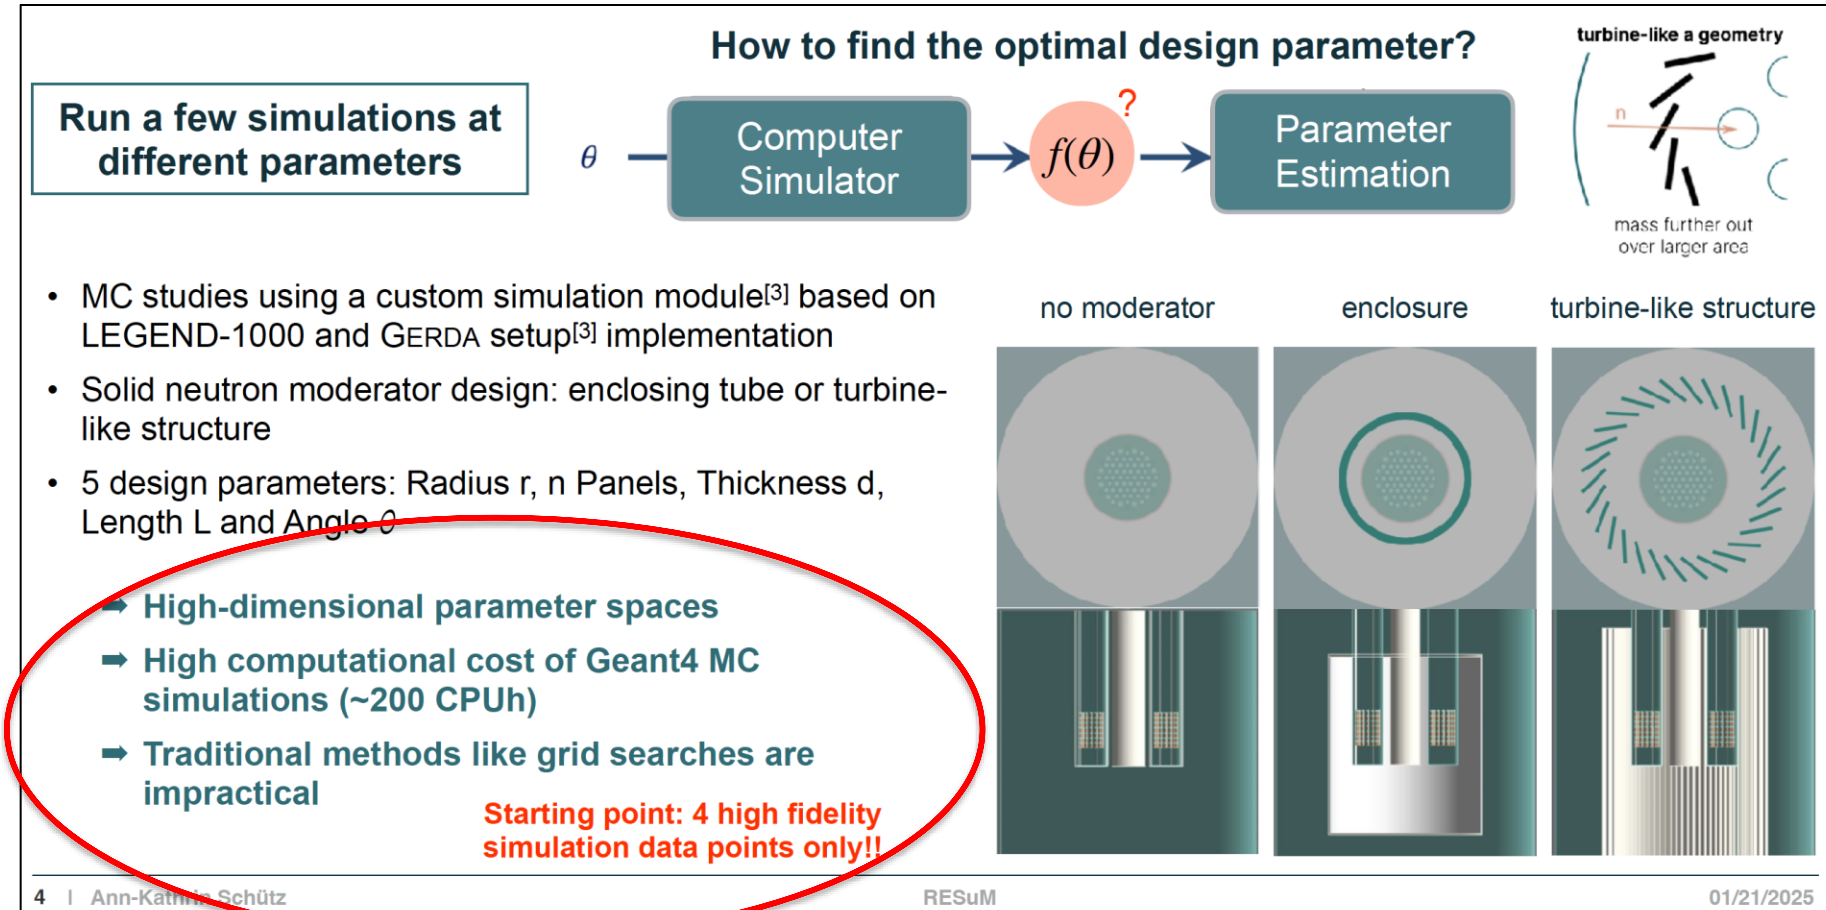

ReSUM

- Conditional Neural Process (CNP)
- Multi-Fidelity Gaussian Process (MFGP)

arXiv:2410.03873

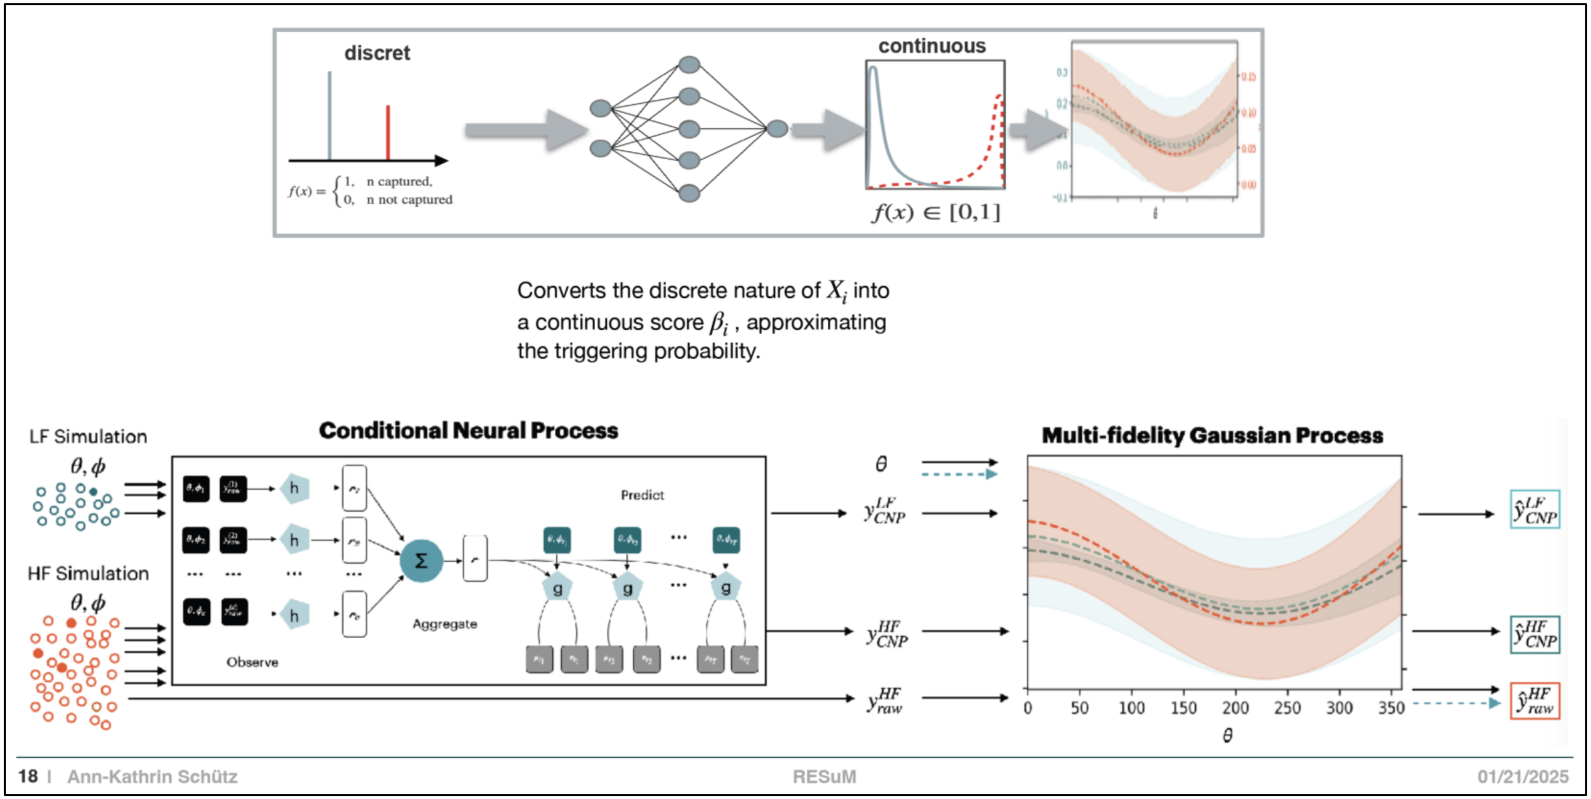

Recent calculation to optimize LEGEND neutron moderator design:
- 70% reduction in bkgd relative to baseline design (w/ UQ)
- factor ~30 gain in computational efficiency vs standard MC methods

arXiv:2508.01834

Efficient optimization of expensive black-box simulators via marginal means, with application to neutrino detector design

Hwanwoo Kim*, Simon Mak*†, Ann-Kathrin Schuetz‡, Alan Poon‡§

- ML(GP)-based interpolation: Black-Box Optimization via Marginal Means (BOMM)

- Active Learning: dynamically learn contours of parameter space to optimally choose next design points
- use in conjunction with ReSUM (WIP)

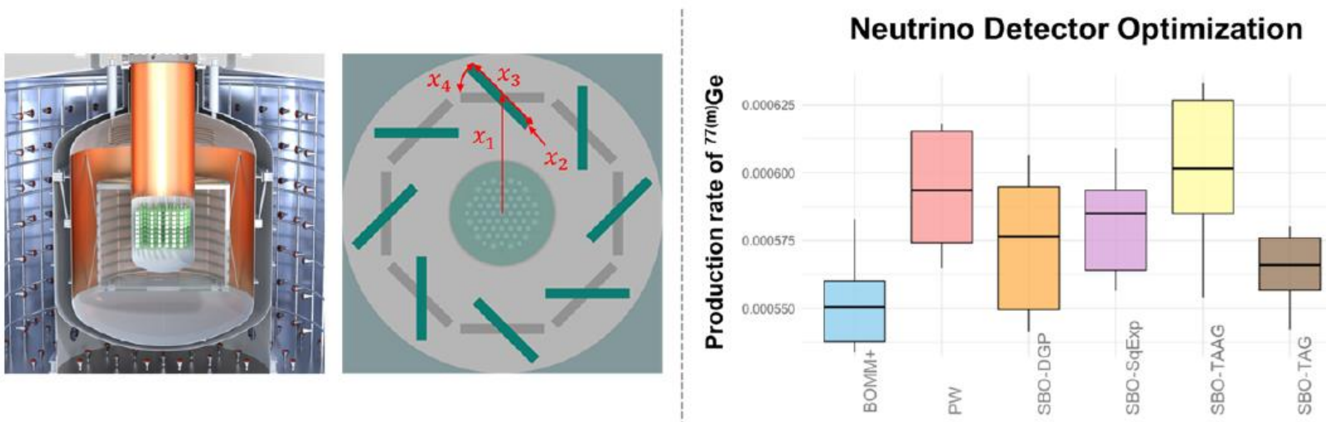

Figure 6: *[Left] The neutrino detector schematic for the LEGEND project, along with the considered neutron moderator geometry with $d = 4$ inputs. [Right] Comparison of $^{77(m)}$Ge production rates (smaller-the-better) for the selected moderator designs from each method. Boxplots show experiment variability over 10 replications.*

Selected results: QCD

(Quark-Gluon Plasma; low-x/gluon saturation)

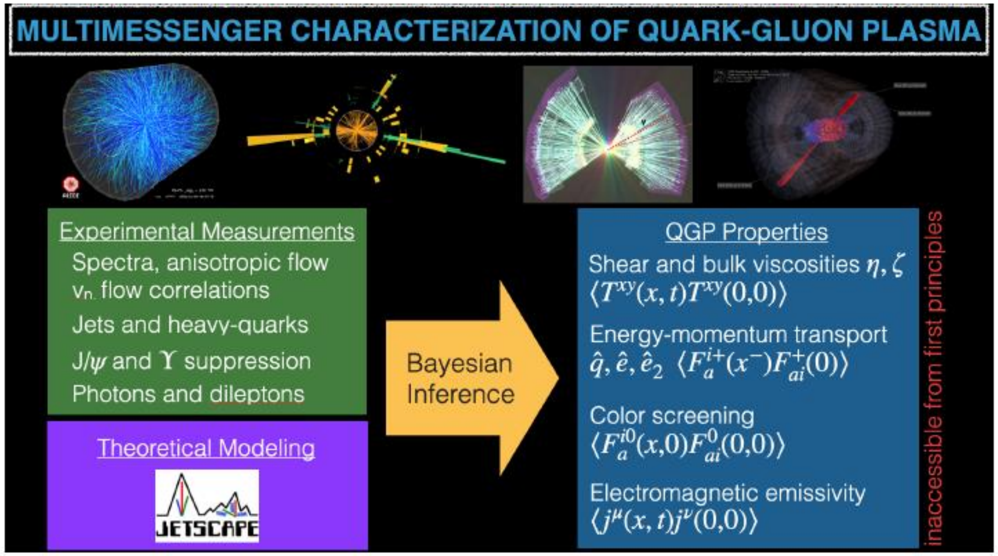

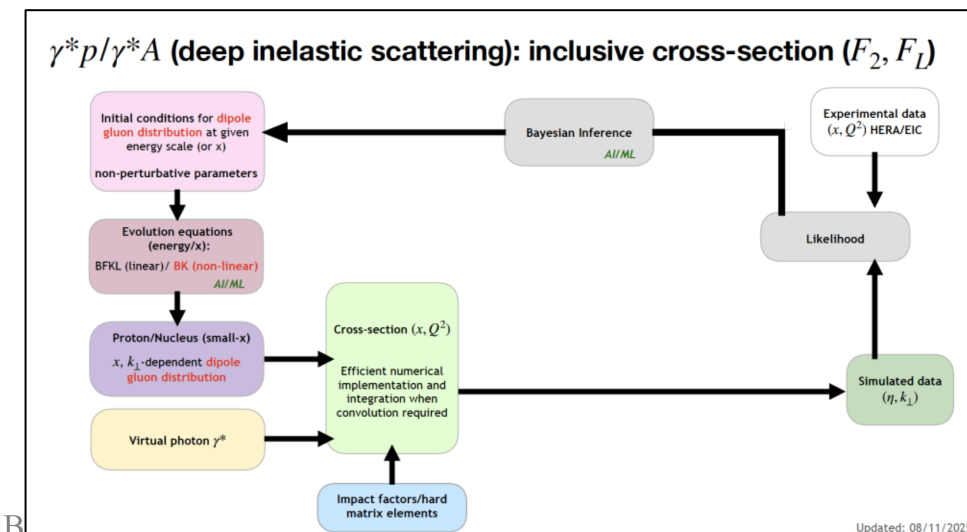

ML-enhanced Bayesian Inference for QGP studies

Rigorous comparison of complex modeling and complex heterogeneous data

- tension/agreement of model description w/ data
- constraint of QGP physical properties

Characterized by moderate number of parameters (~5-25) but very expensive forward model

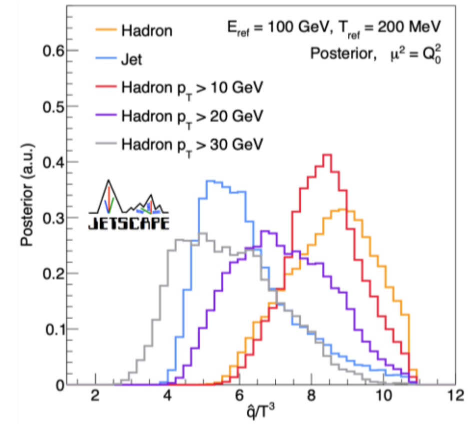

Results from recent calibration

Typical workflow

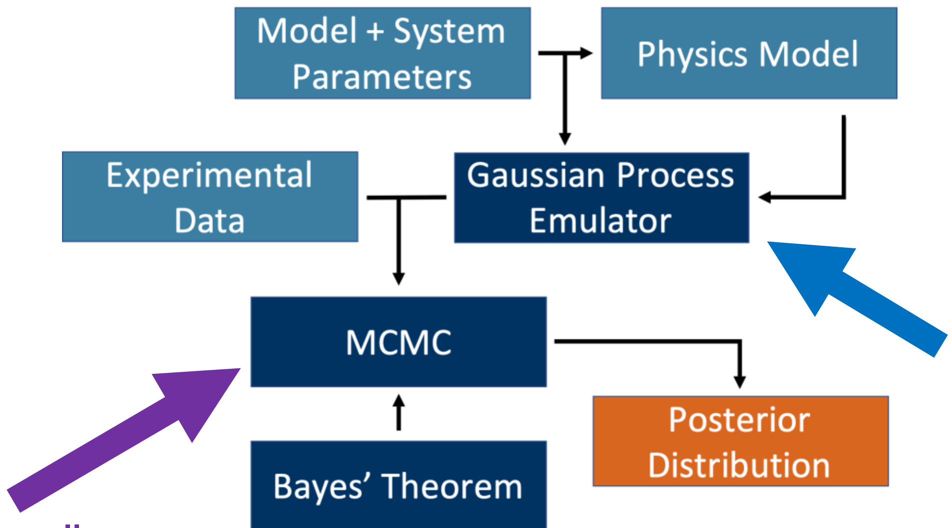

High performance emulation: multi-fidelity implementations

High performance MC sampling: active learning, normalizing flows

Released to HI community: [JETSCAPE/Bayesian]

Multi-fidelity emulation: VarP-GP

RJE, IJ, SM, PMJ; paper in preparation

Learn contours of parameter space: vary computational precision to optimize physics accuracy at minimum computational cost

Example: JETSCAPE inference calculation for jet propagation in QGP

- constrain jet transport parameters

Improved uncertainty at fixed computing budget

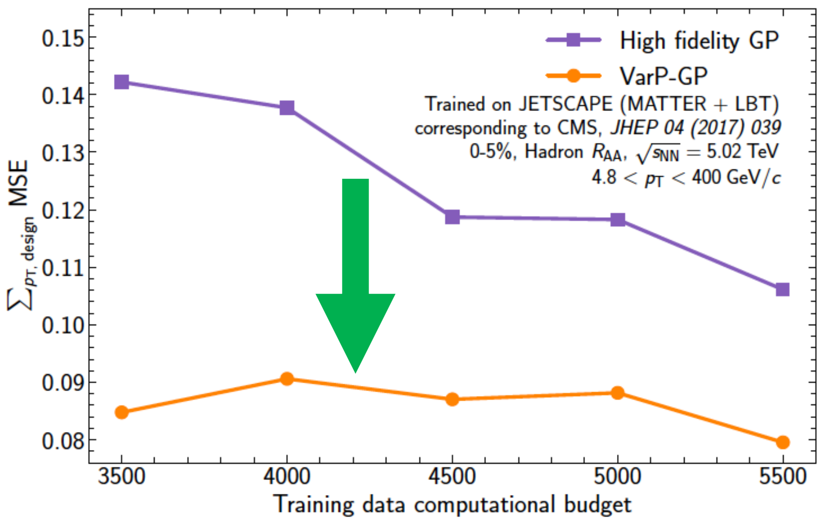

Similar physics sensitivity at reduced computing budget

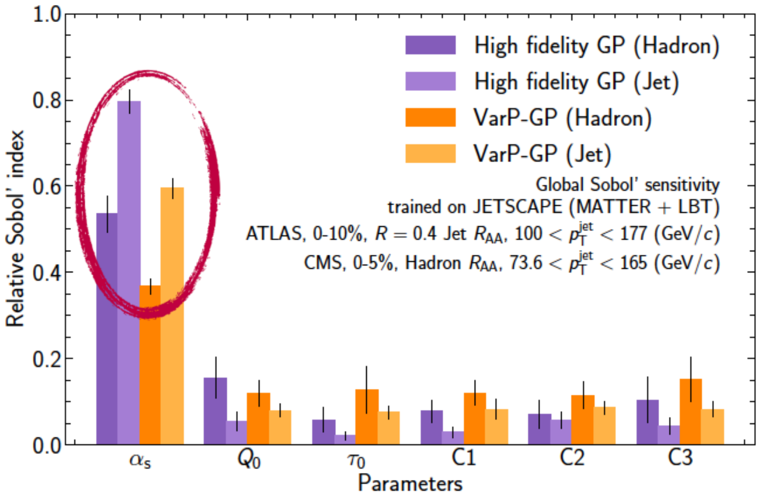

Search for gluon saturation

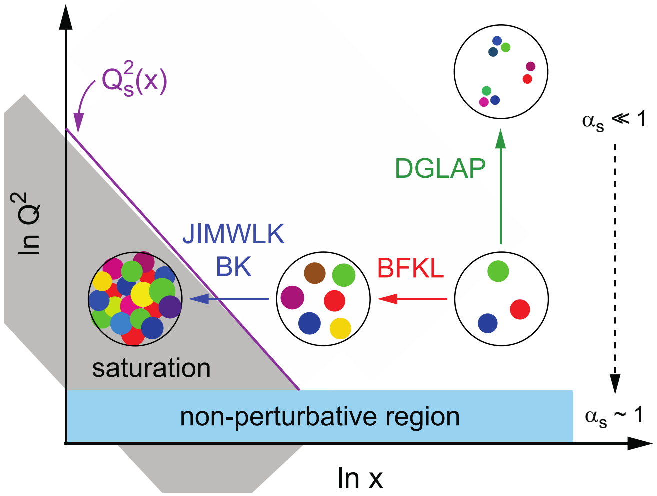

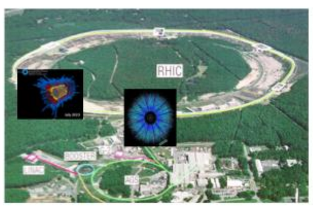

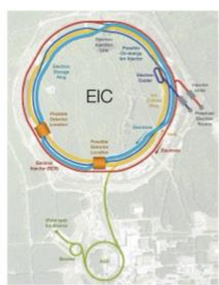

Multi-messenger program: combine measurements from e-A DIS and diffractive interactions at HERA+EIC, with forward p-A collisions at RHIC+LHC
→ ML-enhanced Bayesian Inference

Example workflows for inference

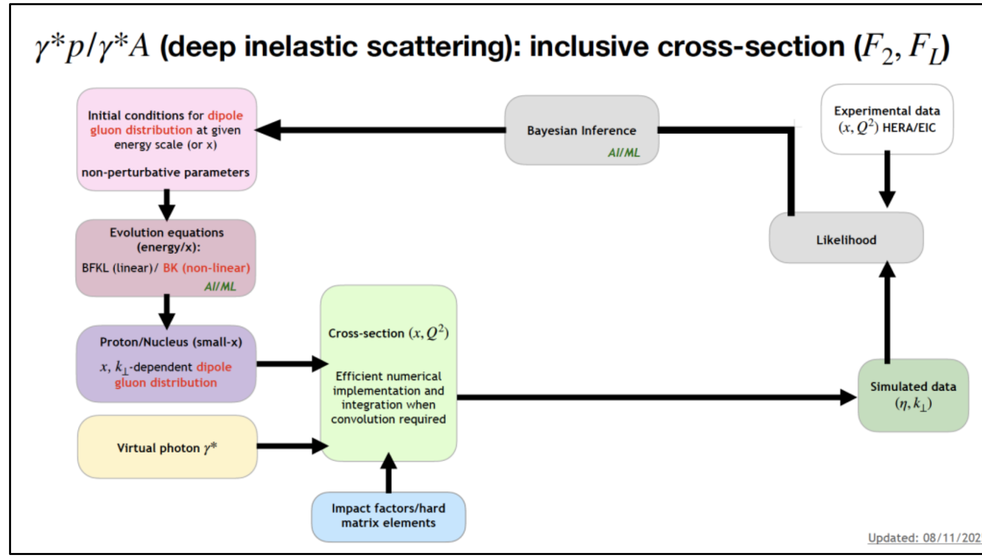

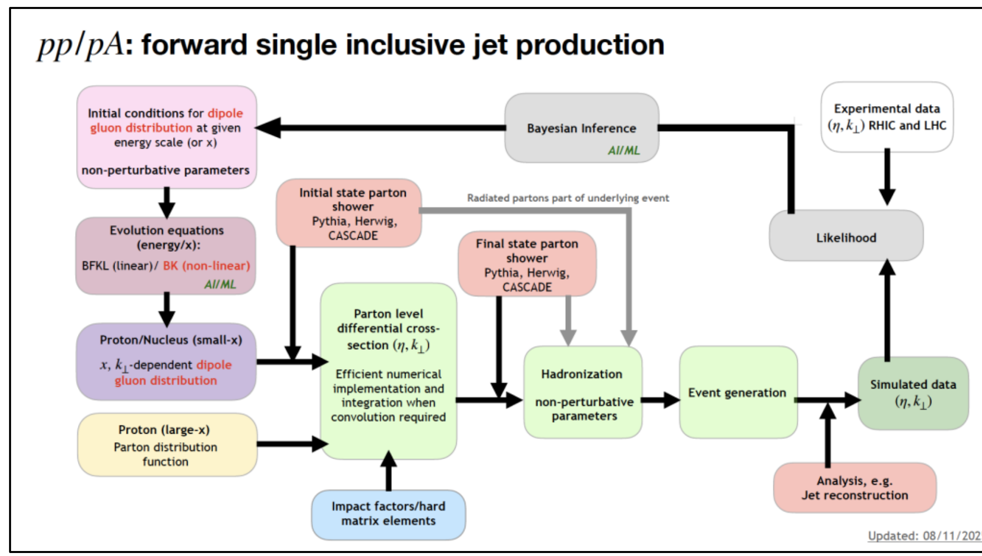

Normalizing flows for sequential inference

Generative model encoding bijective mapping between a complex (~multimodal) distribution and a simple (~Gaussian) distribution

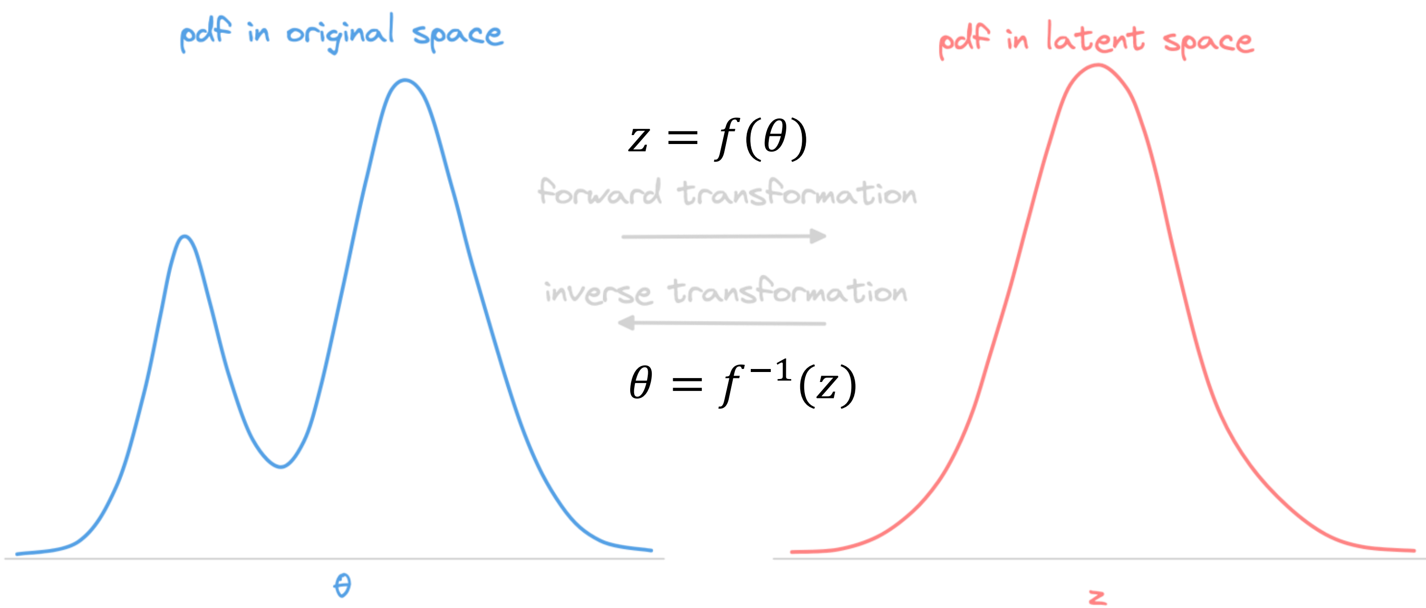

$$p_z(z) = p_\theta(f^{-1}(z)) \left| det \frac{\partial f^{-1}(z)}{\partial z} \right|$$

NF for sequential Bayesian analysis: arXiv:2310.04635

CGC/JIMWLK analysis of J/ψ photoproduction: sequential Bayesian Inference w/ NFs

Mantysaari et al., arXiv:2507.14087
Roch and Shen, arXiv:2509.14911

Data: J/ψ photoproduction from γ+p and γ+Pb

Sequential Bayesian Inference:

- calibrate first on γ+p and then γ+Pb (or vice versa)
- second-step prior is first-step posterior encoded as NF
- compare to single-step joint calibration

CGC model parameters

TABLE I: Summary of model parameters and their prior ranges. The first block corresponds to the standard setup from Refs. [19, 20], the second to JIMWLK evolution parameters, and the third to model extensions investigated in Sec. IV B. The last two columns show the estimated model MAP parameter sets from the 25 posterior sample runs with the smallest χ^2/dof compared to the experimental data.

Block	Parameter	Description	Prior range	MAP ($K \equiv 1$)	MAP (variable K)
1	m [GeV]	Infrared regulator	$[0.02, 1.2]$	0.31	0.51
	B_G [GeV^{-2}]	Proton size	$[1, 10]$	2.99	3.83
	B_q [GeV^{-2}]	Hot spot size	$[0.05, 3]$	0.07	0.30
	σ	Magnitude of Q_s fluctuations	$[0, 1.5]$	0.99	0.88
	$Q_s/(g^2\mu)$	Ratio of color charge density to saturation scale	$[0.05, 1.5]$	0.63	0.37
2	m_{JIMWLK} [GeV]	Infrared regulator	$[0.02, 1.2]$	0.12	0.16
	Λ_{QCD} [GeV]	Spatial Λ_{QCD}	$[0.0001, 0.28]$	0.0097	0.0348
3	N_q	Number of hot spots	$[0, 10]$	3 (fixed)	3 (fixed)
	ω	Modification of the proton shape	$[0, 10]$	1 (fixed)	1 (fixed)
	v_{UV} [GeV^{-1}]	Damping of high-frequency modes	$[0.0, 1.0]$	0 (fixed)	0 (fixed)
	K	Cross section scaling factor	$[0.01, 4]$	1 (fixed)	0.308
				$\chi^2/\mathrm{dof} = 5.94$	$\chi^2/\mathrm{dof} = 1.82$

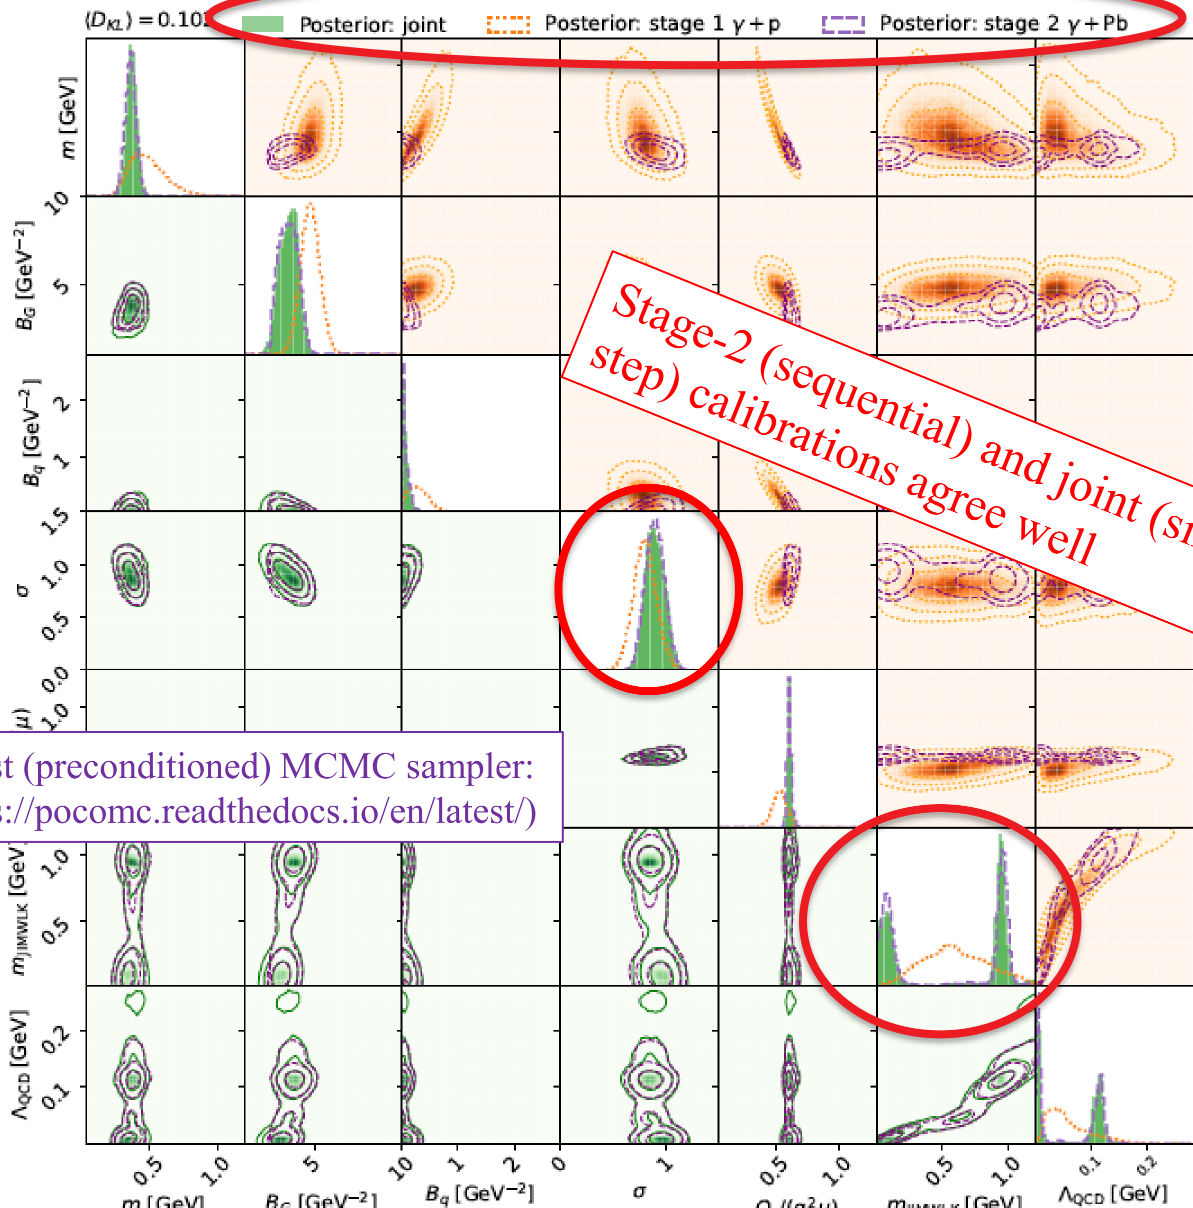

Stage-2 (sequential) and joint (single-step) calibrations agree well

Requires robust (preconditioned) MCMC sampler: pocoMC (https://pocomc.readthedocs.io/en/latest/)

Selected results: radiation imaging of the environment

Motivation
- Contamination mapping
- Radiation survey path planning
- Emergency response
- …

Goal: quantify radiation intensity distribution

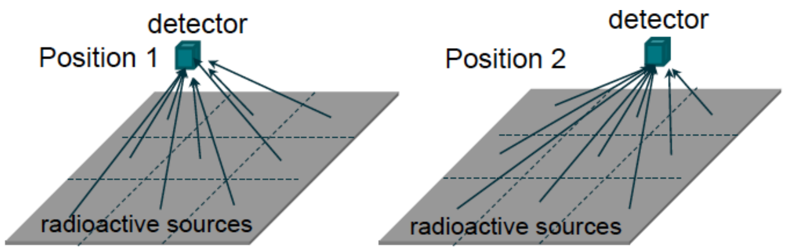

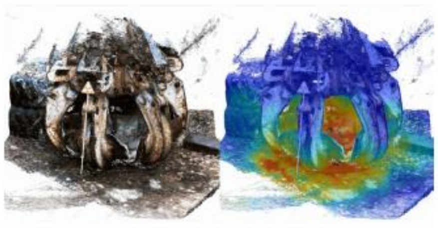

Chornobyl claw, Vetter et al. (2019)

MicroCanonical Langevin Monte Carlo (MCLMC): variant of Markov Chain Monte Carlo (MCMC) developed by Seljak & Robnik (UC Berkeley)

- Standard MCMC: computationally very expensive in high dimensions
- MCLMC is designed to scale much better at high dimensions

Deployed MCLMC for free-moving radiation image reconstruction

- computationally tractable (near-real-time) UQ for distributed radiation
- more stable convergence than traditional reconstruction methods like Maximum Likelihood Expectation Maximization (ML-EM)

Radiation image reconstruction using MCLMC outperforms traditional Hamiltonian Monte Carlo (HMC)

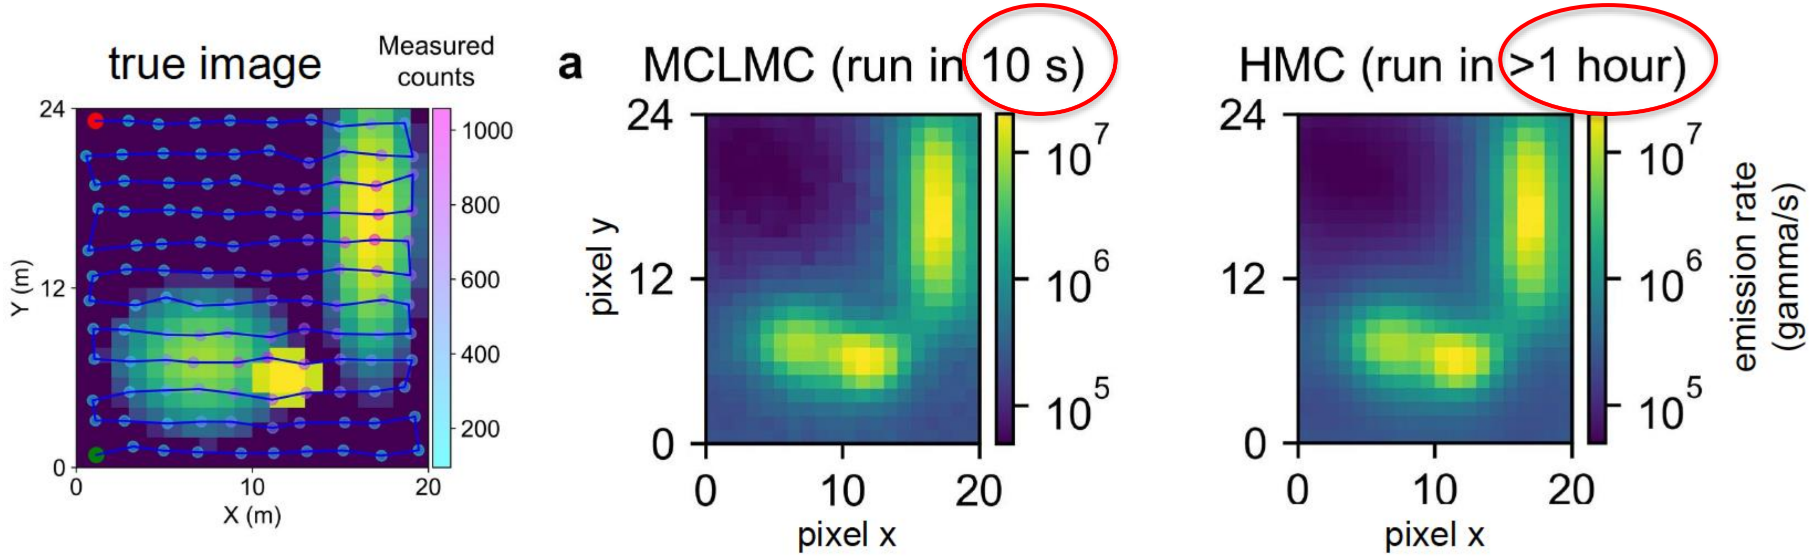

MCLMC solves the overfitting problem

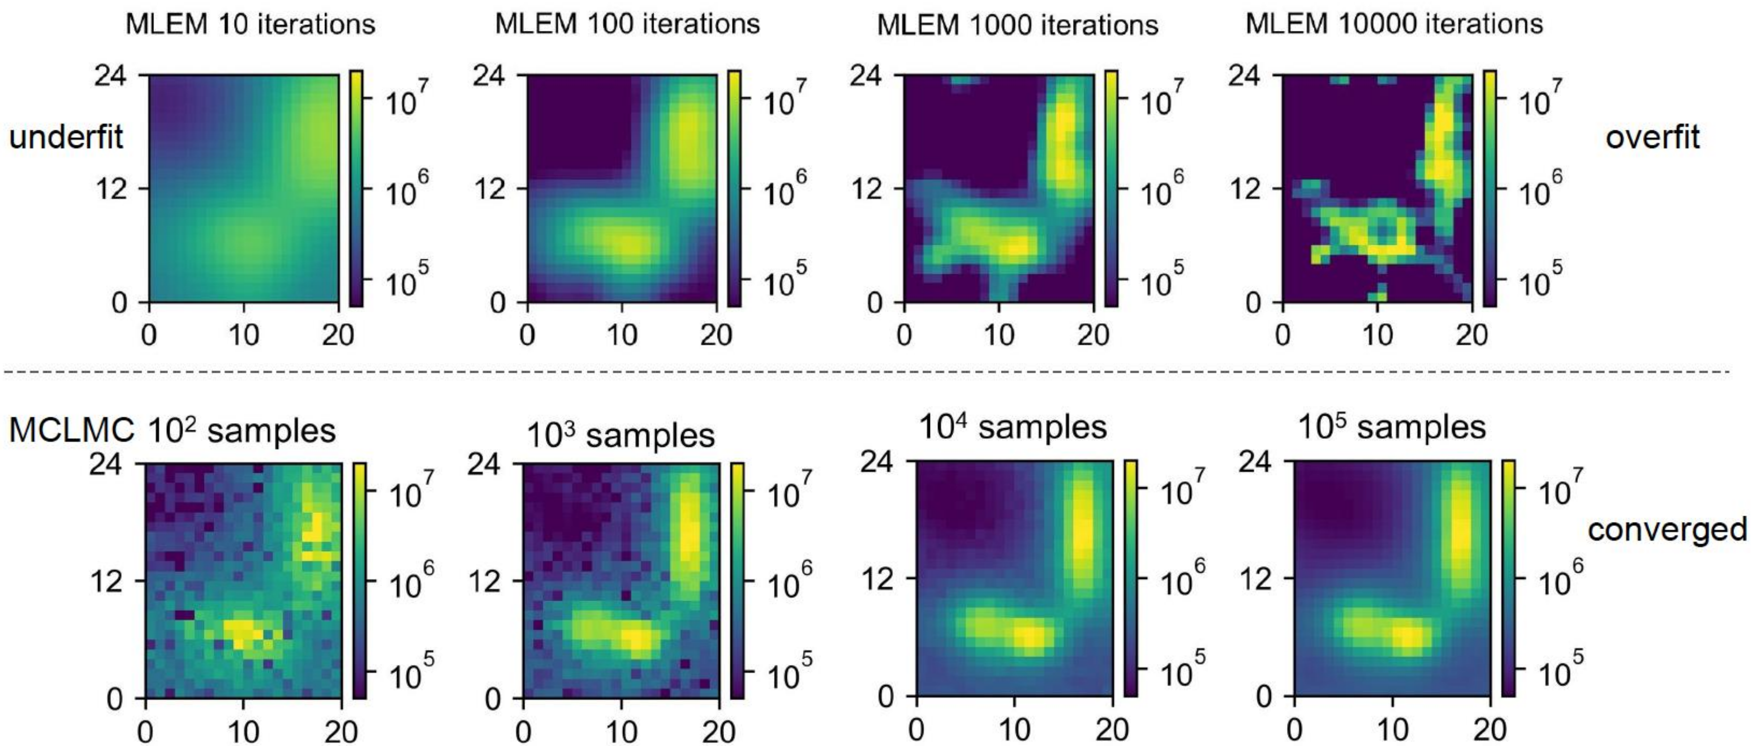

MCLMC provides UQ estimates

Traditional ML-EM only provides point estimates of the pixel/voxel intensities

Choice of prior can strongly influence the UQ magnitude

- e.g. Gaussian Process prior makes nearby pixels strongly correlated, reducing the range of emission rates consistent with the data

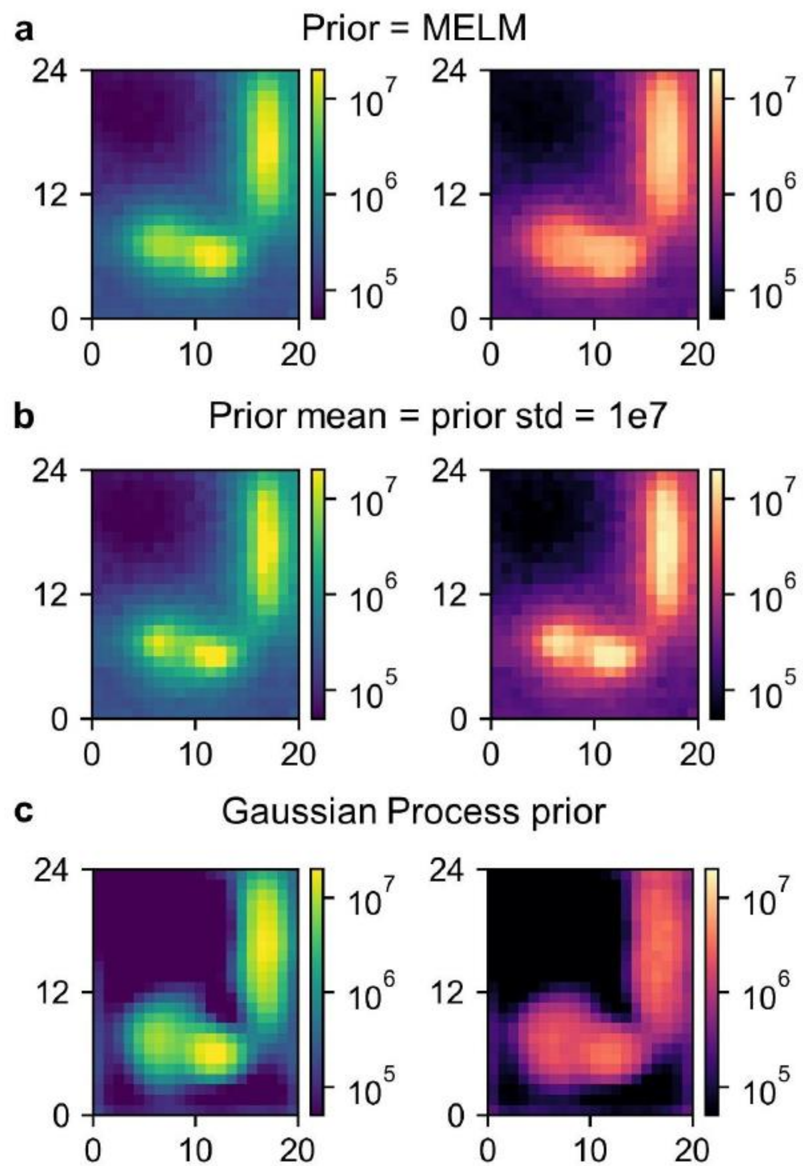

BUQ Phase 2 Budget

	Year 1 ($k)	Year 2 ($k)	Total ($k)
Allocated			
LBNL	666K	686K	1,352K
Duke	148K	152K	300K
UC Berkeley	49K	48K	95K
UC San Diego	150K	155K	305K
Wayne State	122K	126K	248K
total	**1,133K**	**1,167K**	**2,300K**

UC Berkeley also supported by carryover from BUQ Phase 1 (~$130K)

Deliverables: neutrinos

Topic	Y1 PI + staff	Y1 PD + students	Y2 PI + staff	Y2 PD + students
Neutrinos				
Bayesian background model	Li, Poon, Fujikawa, Mak	LBNL(0.5), UCSD(0.5), UCSD(students), Duke(students)		
Enhance/Benchmark RESuM			Li, Poon, Fujikawa, Mak	LBNL(0.5), UCSD(0.5), UCSD(students), Duke(students)
Advanced Sampling Techniques	Kolomensky, Poon, Fujikawa, Seljak	LBNL(0.25), UCB(0.5)	Kolomensky, Poon, Fujikawa, Seljak	LBNL(0.25), UCB(0.75)
Spectrum Modeling KATRIN	Poon	LBNL(0.25)	Poon	LBNL(0.25)
Source Design CUPID	Kolomensky	UCB (0.25)	Kolomensky	UCB (0.25)

Deliverables: Quark-Gluon Plasma

Topic	Y1 PI + staff	Y1 PD + students	Y2 PI + staff	Y2 PD + students
QGP				
Heteroskedastic GP	Shen, Jacobs, Mak	WSU (0.25), LBNL (0.20), Duke (students)		
Boundary-Safe Model Selection			Shen, Jacobs, Mak	WSU (0.25), LBNL (0.1), Duke (students)
Theory UQ	Shen	WSU (0.25)	Shen	WSU (0.25)
High-dim Analysis	Shen, Seljak, Mak	WSU (0.25), UCB (0.25), Duke (students)		
Generative AI			Shen, Jacobs, Mak	WSU (0.25), LBNL (0.1), Duke (students)
Iterative Multi-Messenger Analysis	Shen, Jacobs	WSU (0.25), LBNL (0.20)	Shen, Jacobs	WSU (0.25), LBNL (0.20)
Collider Monitoring	Jacobs, Mak	LBNL (0.1), Duke (students)	Jacobs, Mak	LBNL (0.1), Duke (students)

Deliverables: radiation mapping

Topic	Y1 PI + staff	Y1 PD + students	Y2 PI + staff	Y2 PD + students
Radiation mapping				
Data sufficiency	Vavrek, Mak	LBNL (0.25)	Vavrek, Mak	LBNL (0.25)
Langevin UQ	Vavrek, LBNL staff (0.1)		Vavrek, LBNL staff (0.1)	

BUQ-supported publications 2024

- RESuM: Rare Event Surrogate Model for Physics Detector Design
  A.-K. Schuetz, A. W. P. Poon, and A. Li
  ICLR 2025 Proceedings
  arXiv:2410.03873 [physics.ins-det]
- Model emulation and closure tests for (3+1)D relativistic heavy-ion collisions
  H. Roch, S. A. Jahan and C. Shen
  Physical Review C110, 044904 (2024)
  arXiv:2405.12019
- Bayesian analysis of (3+1)D relativistic nuclear dynamics with the RHIC beam energy scan data
  S. A. Jahan, H. Roch and C. Shen
  Physical Review C110, 054905 (2024)
  arXiv:2408.00537
- A graphical multi-fidelity Gaussian process model, with application to emulation of expensive computer simulations
  Y. Ji, S. Mak, D. Soeder, J. F. Paquet, and S. A. Bass
  Technometrics 66, 267 (2024)
- Conglomerate multi-fidelity Gaussian process modeling, with application to heavy-ion collisions
  Y. Ji, H. S. Yuchi, D. Soeder, J.-F. Paquet, S. A. Bass, V. R. Joseph, C. F. Wu, and S. Mak
  SIAM/ASA Journal of Uncertainty Quantification 12, 473 (2024)
- Stacking designs: designing experiments for multi-fidelity modeling with confidence
  C. L. Sung, Y. Ji, S. Mak, W. Wang, and T. Tang
  SIAM/ASA Journal on Uncertainty Quantification 12, 157 (2024)
- Local transfer learning Gaussian process modeling, with applications to surrogate modeling of expensive computer simulators
  X. Wang, S. Mak, J. Miller, and J. Wu
  SIAM/ASA Journal on Uncertainty Quantification, to be published 2025
  arXiv:2410.12690
- Hierarchical shrinkage Gaussian processes: applications to computer code emulation and dynamical system recovery
  T. Tang, S. Mak, and D. Dunson
  SIAM/ASA Journal on Uncertainty Quantification 12, 1085 (2024)
- eRPCA: Robust Principal Component Analysis for Exponential Family Distributions
  X. Zheng, S. Mak, L. Xie, and Y. Xie
  Statistical Analysis and Data Mining: The ASA Data Science Journal 17, e11670 (2024)

BUQ-supported publications 2025

- QuIP: Experimental design for expensive simulators with many qualitative factors via Integer Programming
  Yen-Chun Liu and Simon Mak
  arxiv:2501.14616
- Efficient optimization of expensive black-box simulators via marginal means, with application to neutrino detector design
  Hwanwoo Kim, Simon Mak, Ann-Kathrin Schuetz, Alan Poon
  arXiv:2508.01834
- Bayesian inference analysis of jet quenching using inclusive jet and hadron suppression measurements
  JETSCAPE Collaboration (R. Ehlers et al.)
  Physical Review C111 5, 054913 (2025)
- A Gaussian Process Generative Model for QCD Equation of State
  J. Gong, H. Roch and C. Shen
  Physical Review C111 4, 044912 (2025)

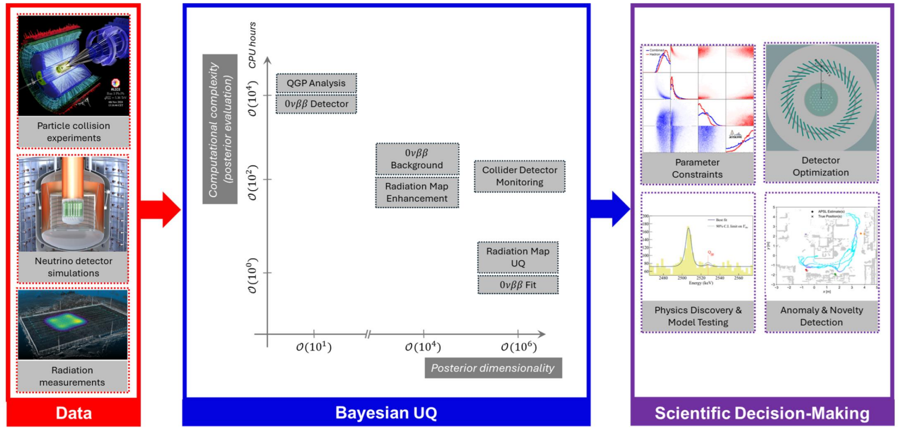

	$0\nu\beta\beta$ Detector Optimization	$0\nu\beta\beta$ Background & Fit	QGP Collider Monitoring	QGP Analysis	Radiation Map Enhancement	Radiation Map UQ
Bayesian Transfer Learning				✓	✓	
Bayesian Multi-Fidelity Learning	✓ ○			✓	✓	
Langevin Monte Carlo		✓ ○		✓		✓ ○
Bayesian Manifold Learning	○			○		
Bayesian Optimization	○					○
Boundary-Informed Surrogates	○			○		
Bayesian Image Change Detection			○			

✓: BUQ Phase 1 ○: BUQ Phase 2

Project Status: multiple ongoing projects, many using similar algorithms

WIP: understand commonalities and connections between similar analysis approaches in these diverse physics areas

Comprehensive project document in progress

Extra slides

Sequential inference

Different classes of experimental observables (“soft”; “hard”)

Sequential Bayesian Inference: posterior from “soft” calibration becomes prior for “hard” calibration

SEQUENTIAL INFERENCE

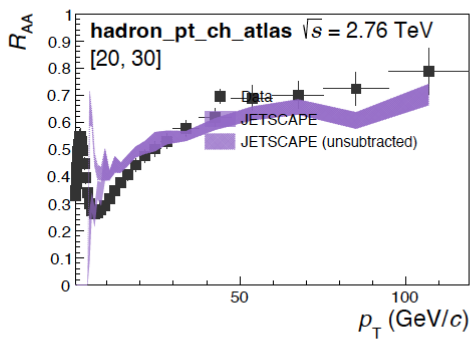

nucleon width = 0.5 fm

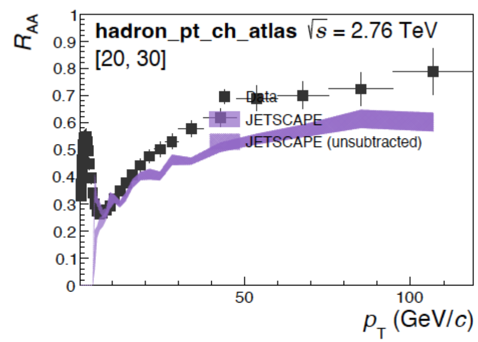

nucleon width = 1.5 fm

- Sequential inference

$$P(\theta_s \mid D_s, D_h) \propto P(D_h \mid \theta_s) \cdot P(D_s \mid \theta_s) \cdot P(\theta_s)$$

$$= \underbrace{\boxed{P(D_h \mid \theta_s)}}_{\text{second inference}} \cdot \underbrace{\boxed{\left[P(D_s \mid \theta_s) \cdot P(\theta_s)\right]}}_{\text{first inference}}$$

- Update knowledge of soft sector parameters θ_s with hard sector data D_h