Oak Ridge National Laboratory

High-dimensional maximum-entropy phase space tomography

Austin Hoover

North American Particle Accelerator Conference

Sacramento, CA, USA

August 12, 2025

Outline

- Phase space tomography and maximum-entropy inference
- Two algorithms:
  - GPSR: generative modeling approach
  - MENT: classical approach
- Conclusion

Background

Goal: infer phase space density from measured projections

$p(x)$

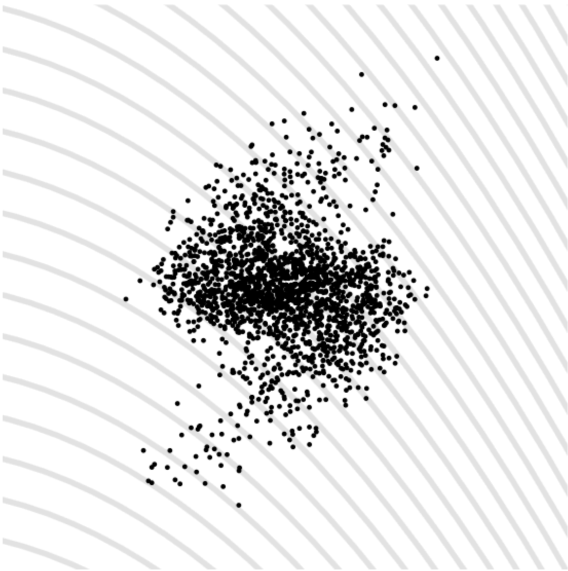

Goal: infer phase space density from measured projections

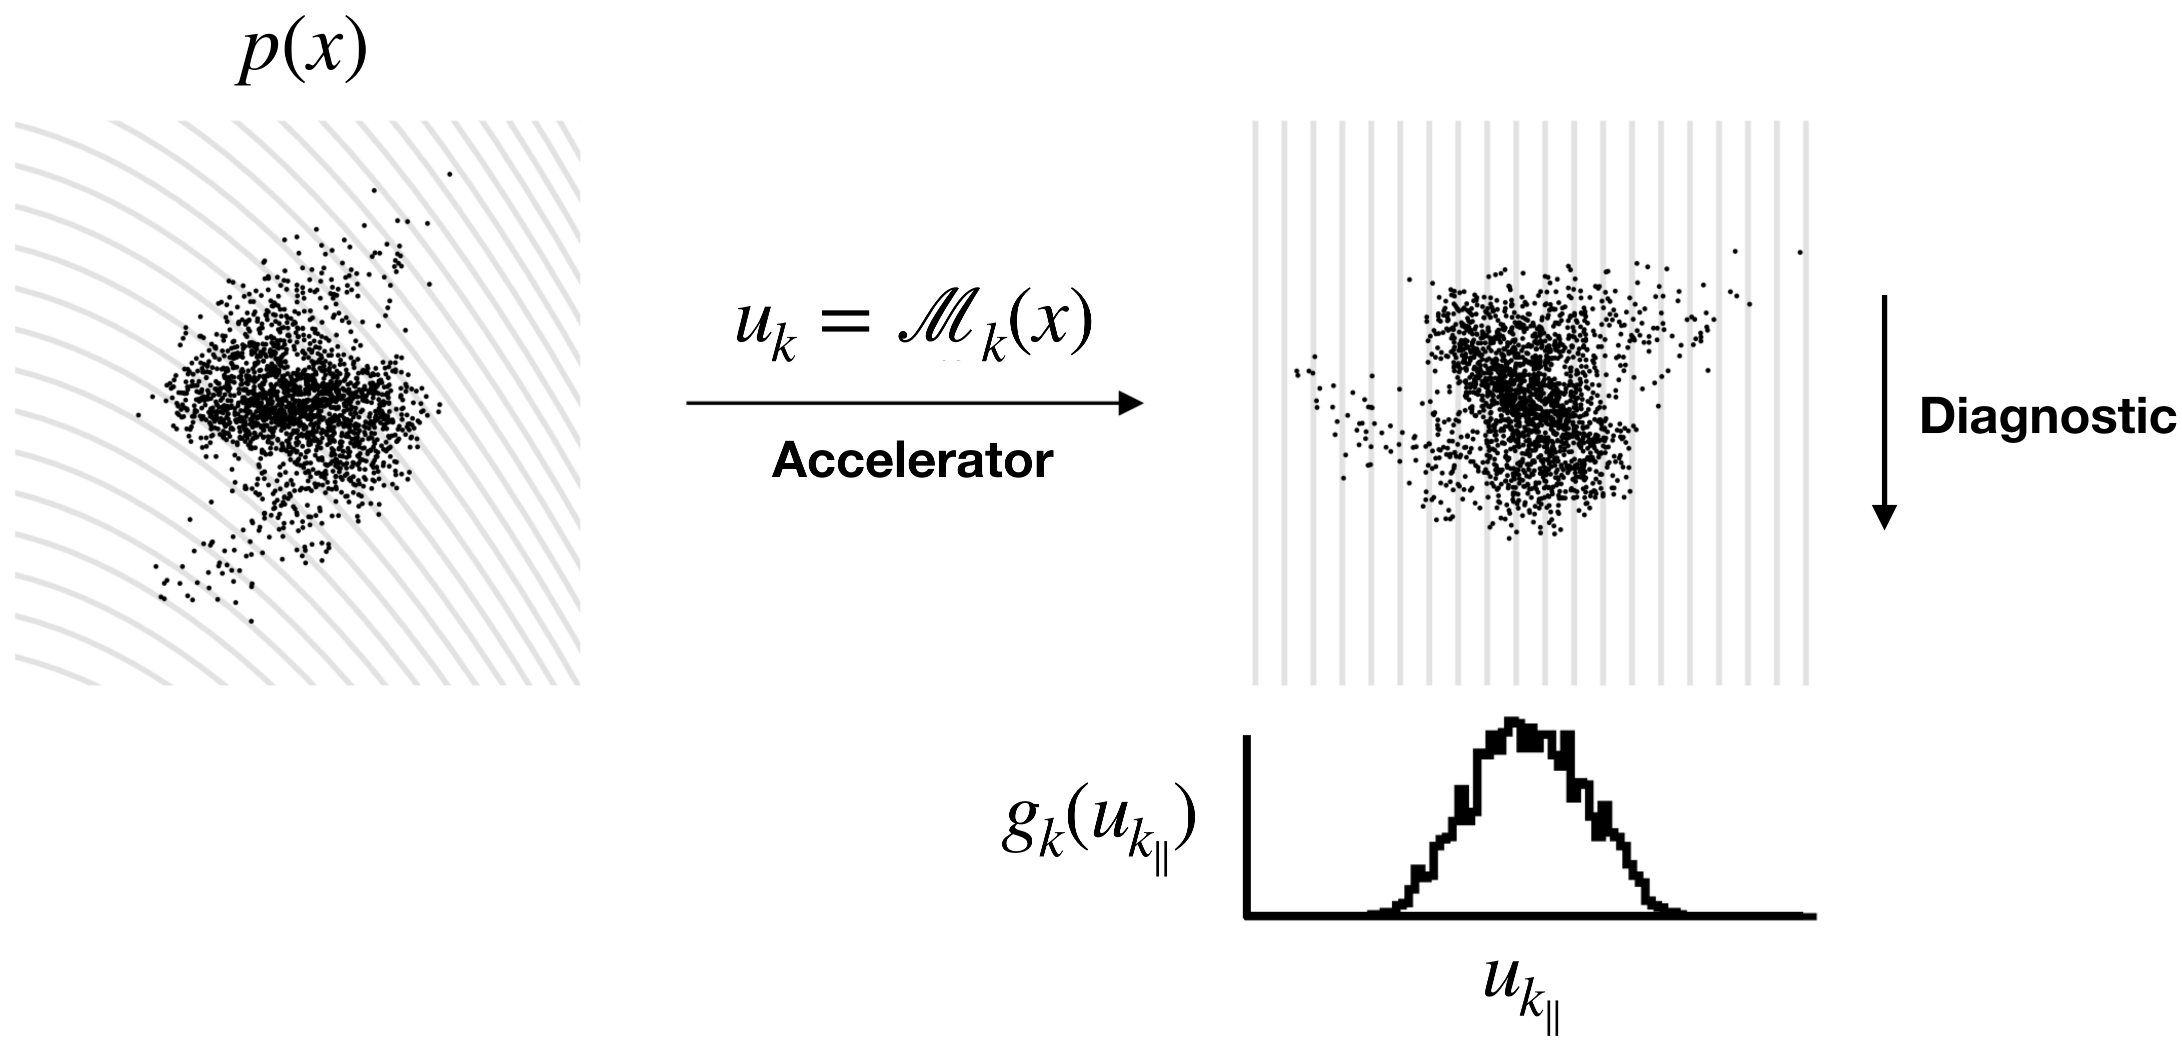

Goal: infer phase space density from measured projections

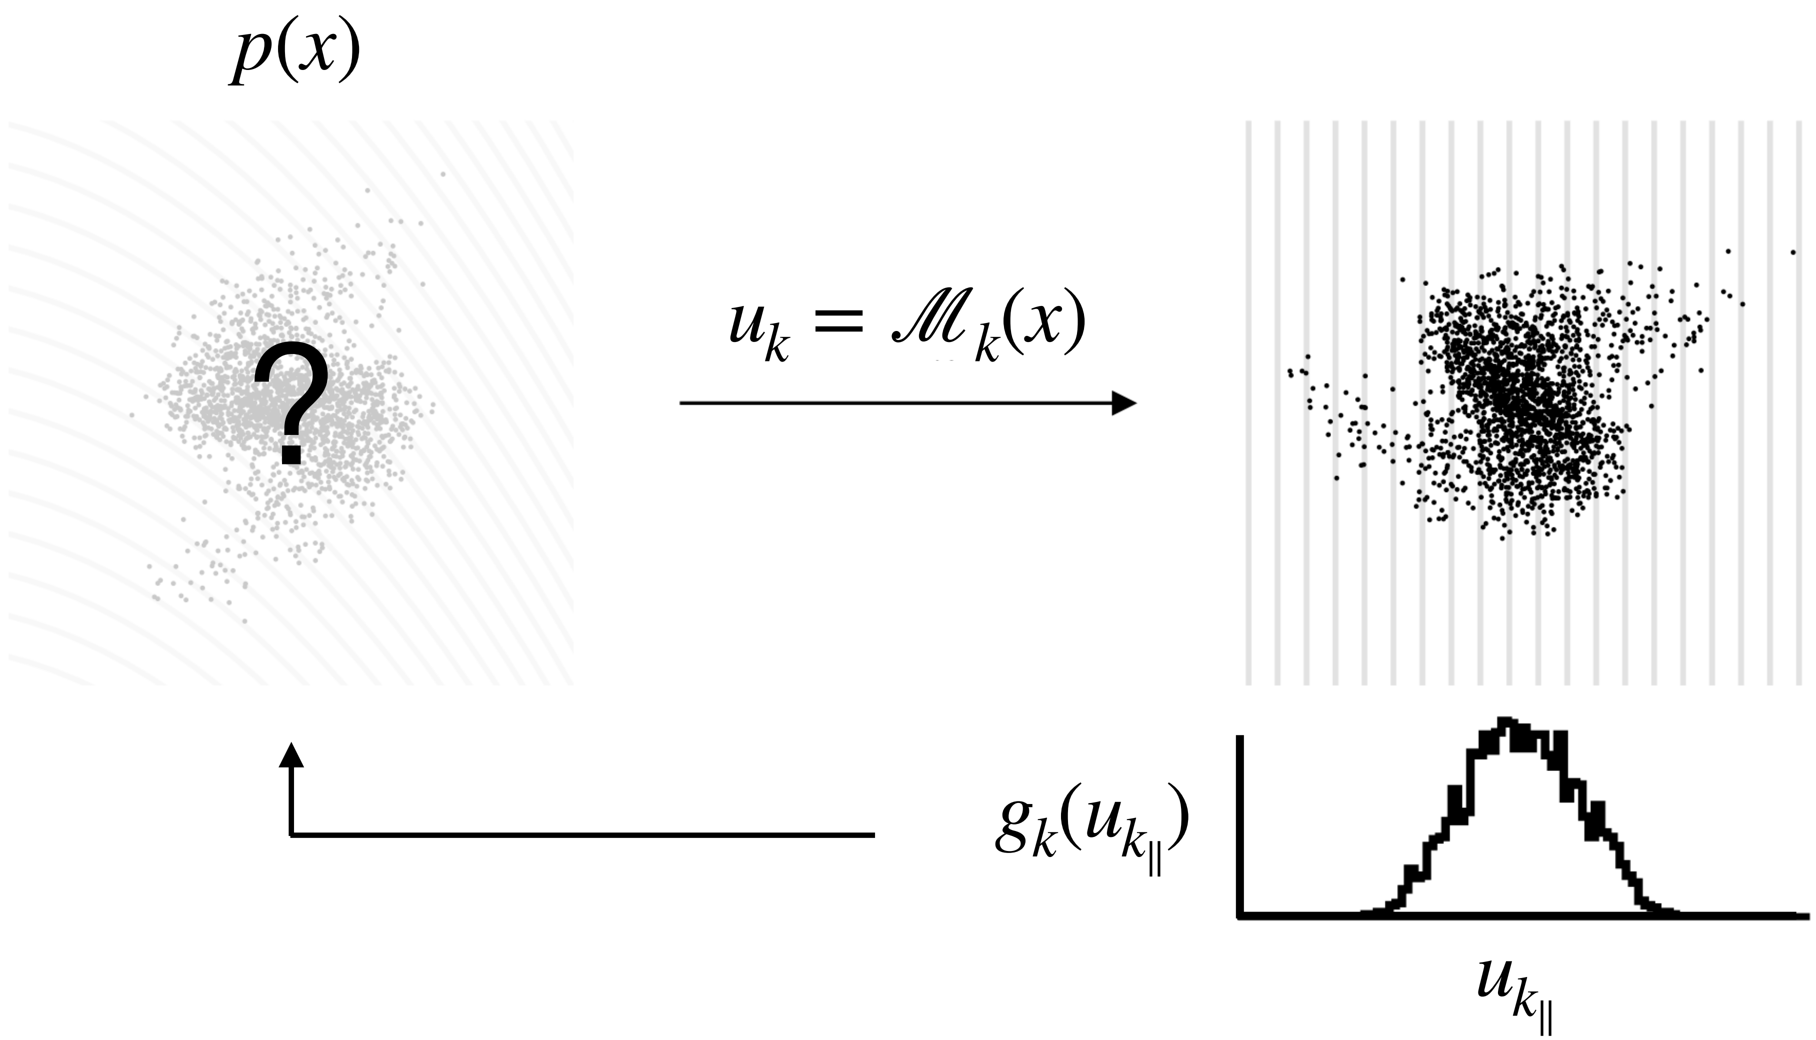

Challenge A: searching the space of distributions

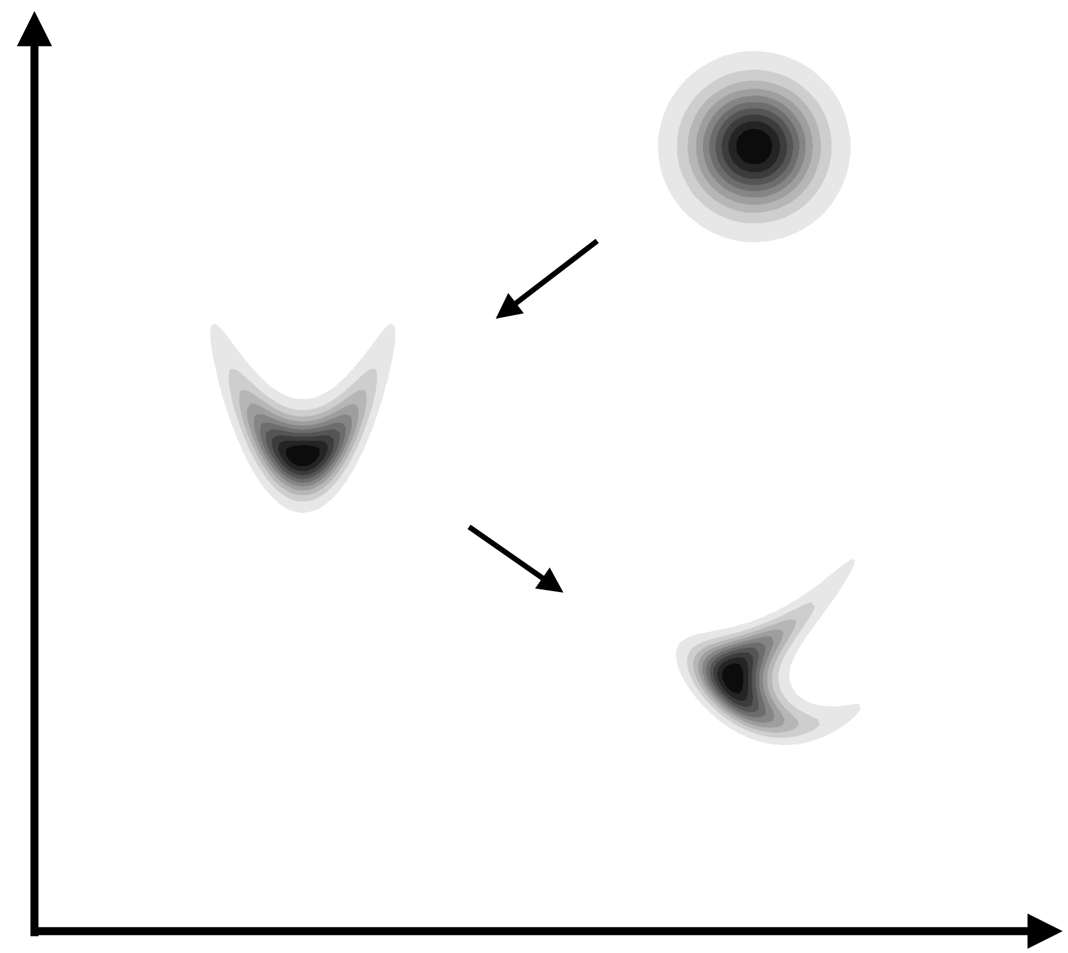

Space of distributions $p(x)$

Challenge B: solution is not unique

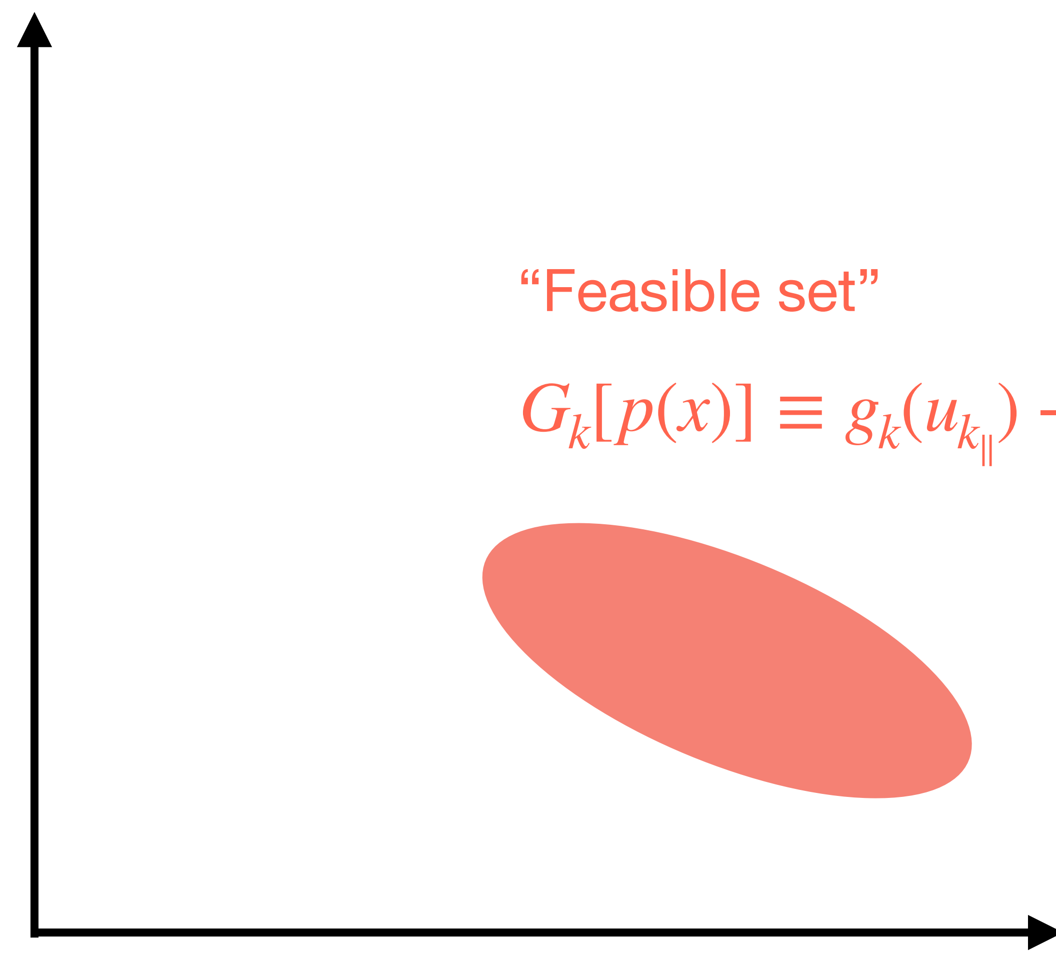

Challenge B: solution is not unique

Things I *will not* discuss:

Experimental design: how to *shrink* volume of feasible set

Uncertainty quantification: how to *calculate* volume of feasible set

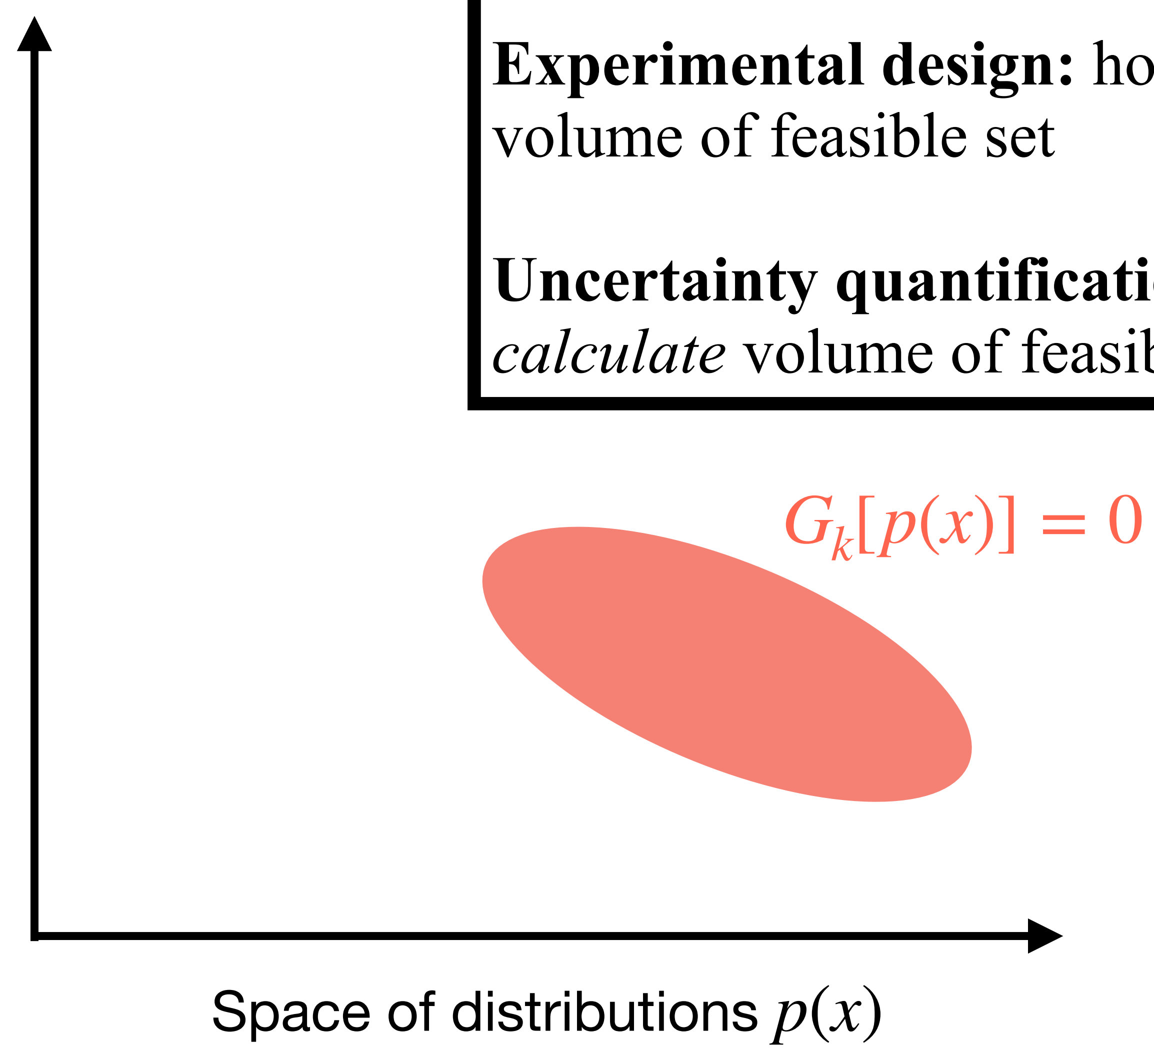

Oak Ridge National Laboratory

Challenge B: solution is not unique

Things I *will* discuss:

How to select a **single distribution** from the feasible set.

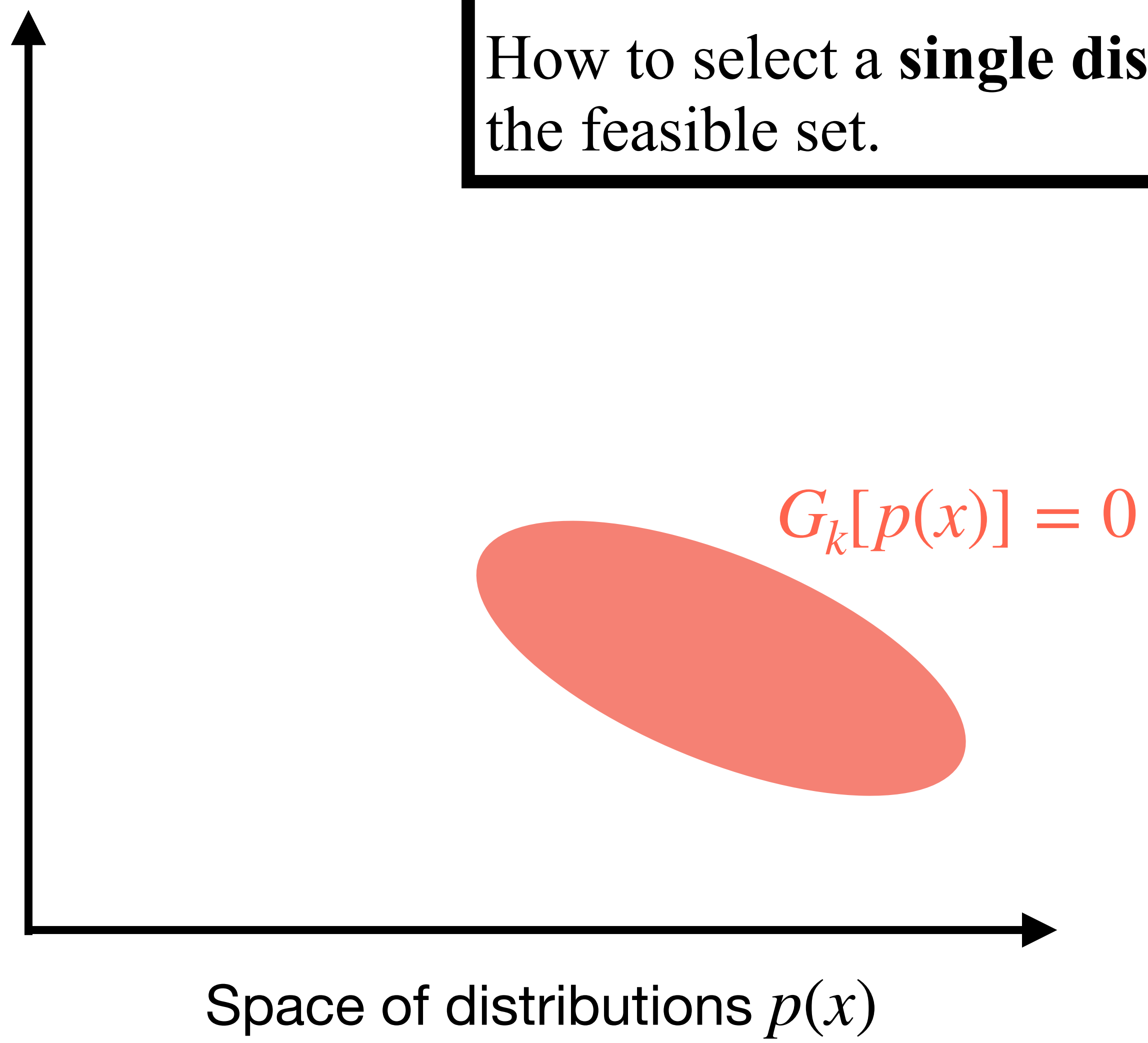

MaxEnt: rank solutions by entropy (simplicity) and choose highest ranked solution

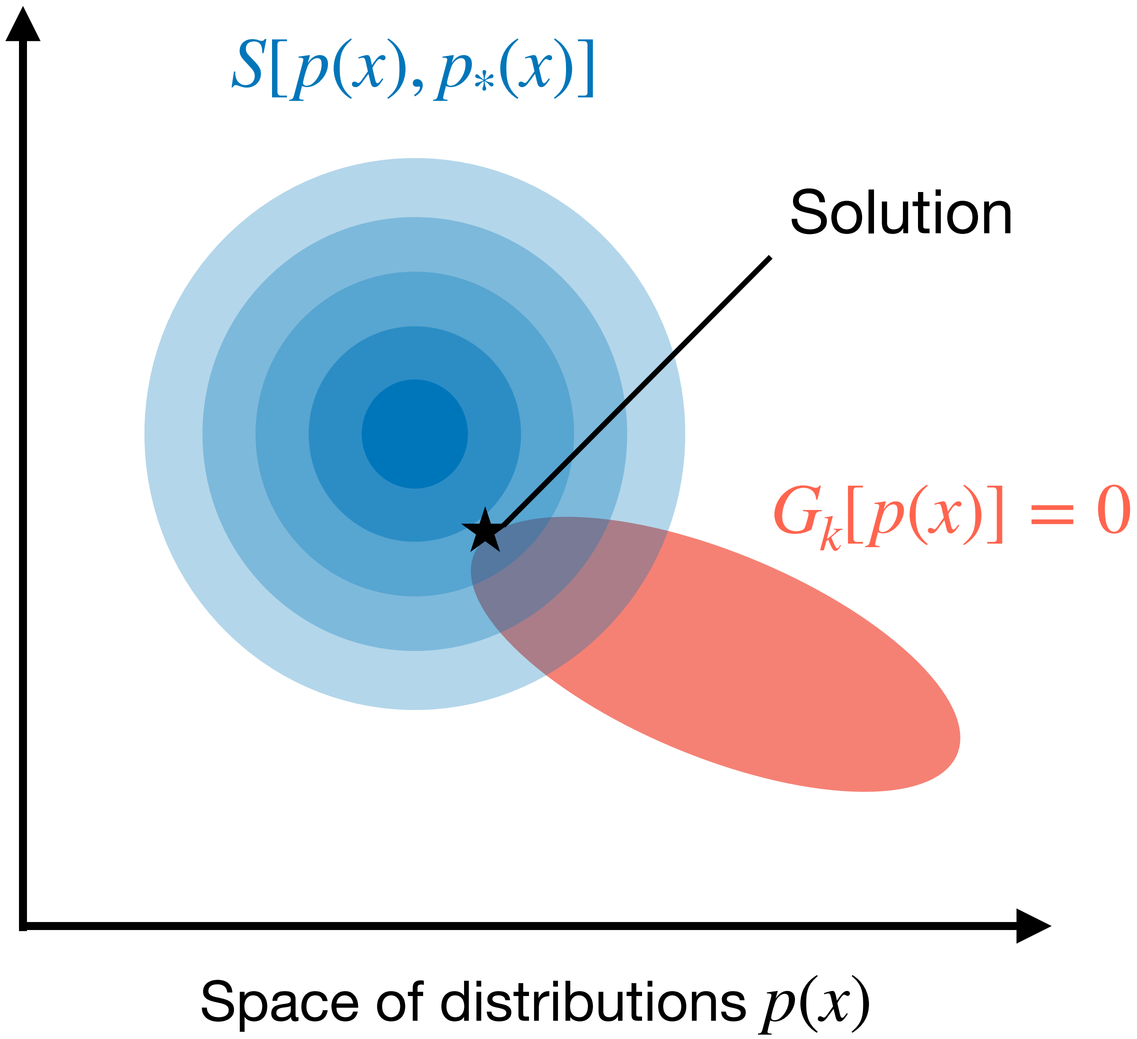

Entropy functional fixed by “Principle of Minimal Updating”

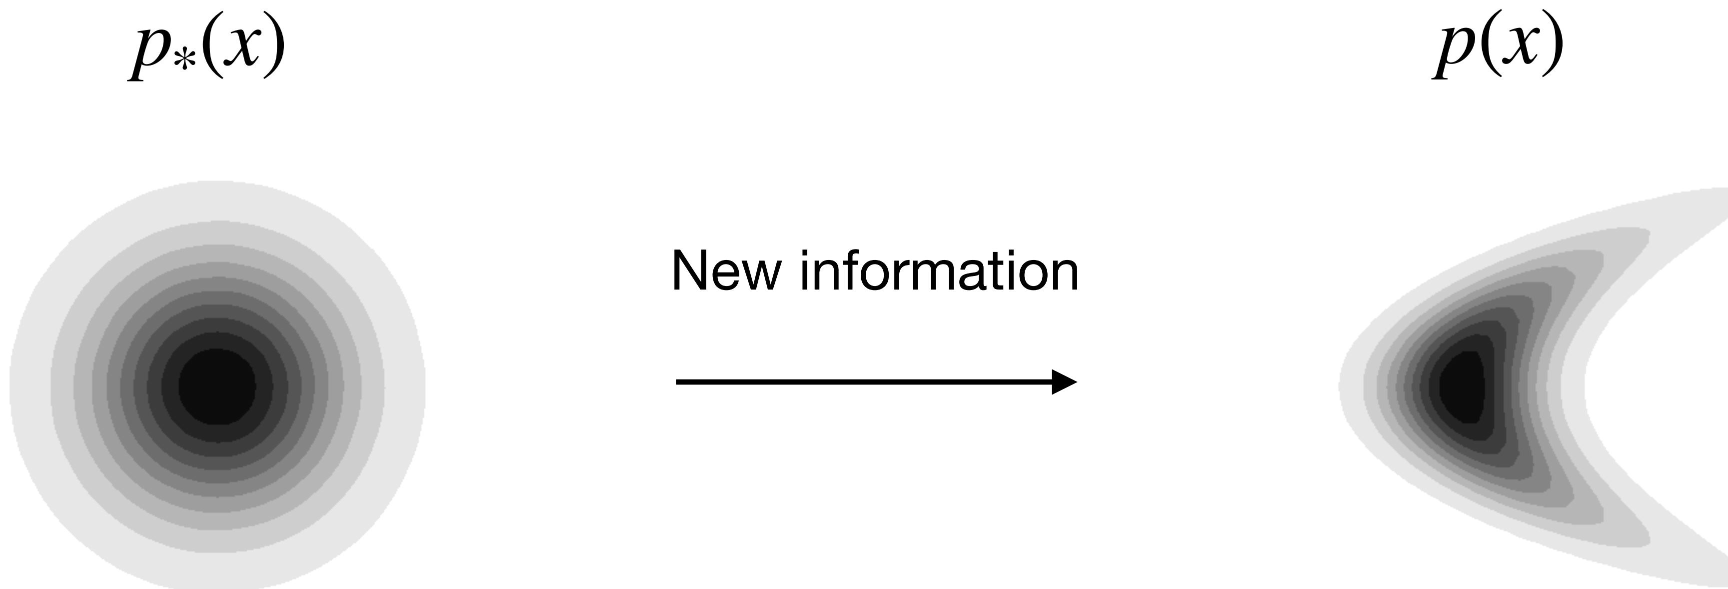

$$S[p(x), p_*(x)] = -\int p(x) \log\left(\frac{p(x)}{p_*(x)}\right) dx$$

MaxEnt solution may be incorrect, but it's the best we can do with the available information

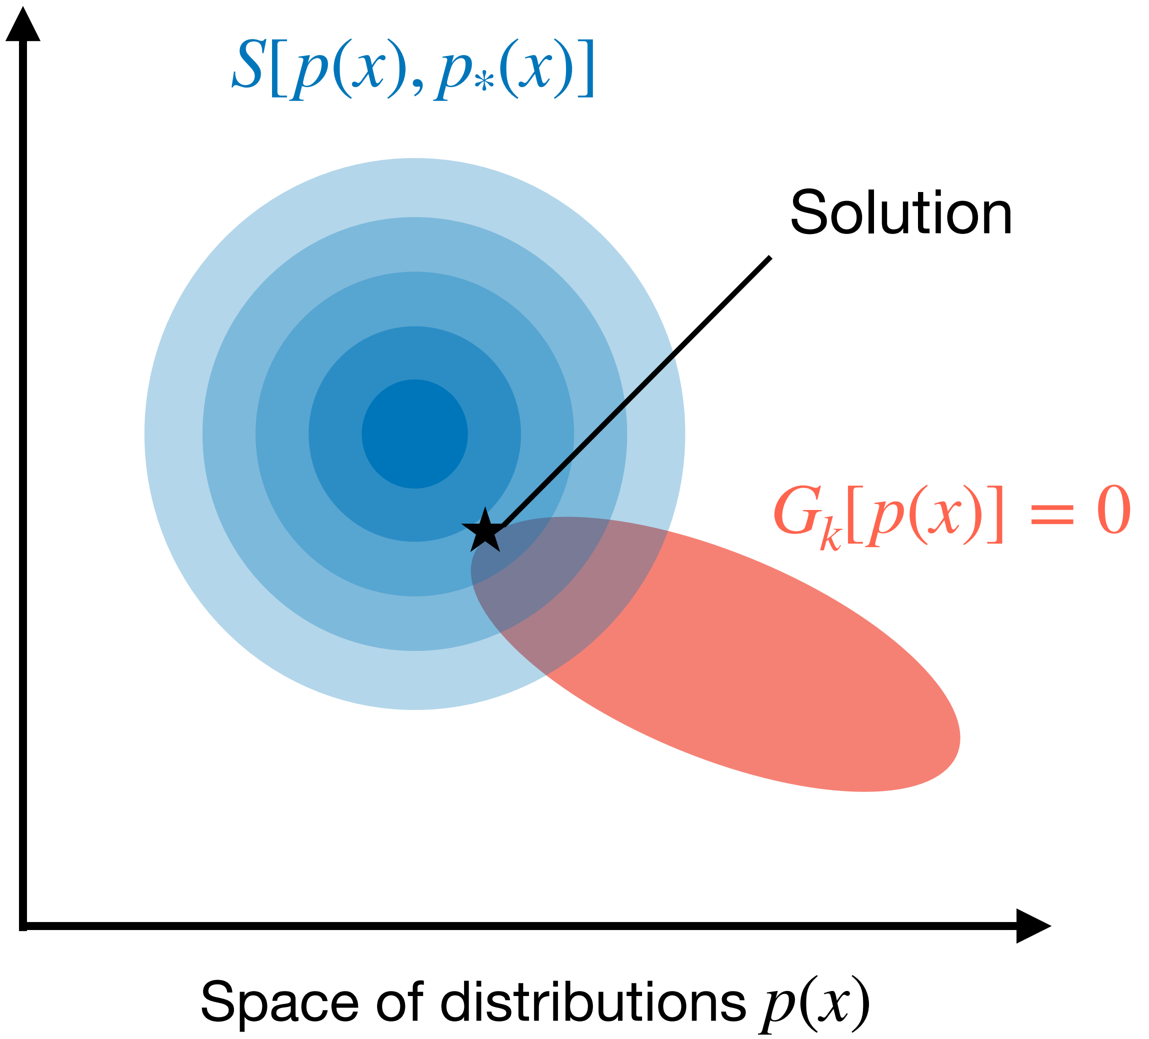

GPSR: generative modeling approach

Generative Phase Space Reconstruction (GPSR)

PHYSICAL REVIEW ACCELERATORS AND BEAMS **27,** 094601 (2024)

Editors' Suggestion

Efficient six-dimensional phase space reconstructions from experimental measurements using generative machine learning

Ryan Roussel,[1] Juan Pablo Gonzalez-Aguilera,[2] Eric Wisniewski,[3] Alexander Ody,[3] Wanming Liu,[3] John Power,[3] Young-Kee Kim,[2] and Auralee Edelen[1]

[1]*SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA*
[2]*Department of Physics and Enrico Fermi Institute, University of Chicago, Chicago, Illinois 60637, USA*
[3]*Argonne National Laboratory, Lemont, Illinois 60439, USA*

(Received 19 April 2024; accepted 12 August 2024; published 11 September 2024)

[MOP034] “Efficient 6-dimensional phase space measurements and applications to autonomous monitoring at LCLS-II”
[MOP029] “Development and Applications of Differentiable Coherent Optical Transition Radiation Simulations”
[MOP071] “Phase space reconstruction of beams affected by coherent synchrotron radiation”

Generative Phase Space Reconstruction (GPSR)

Generative model

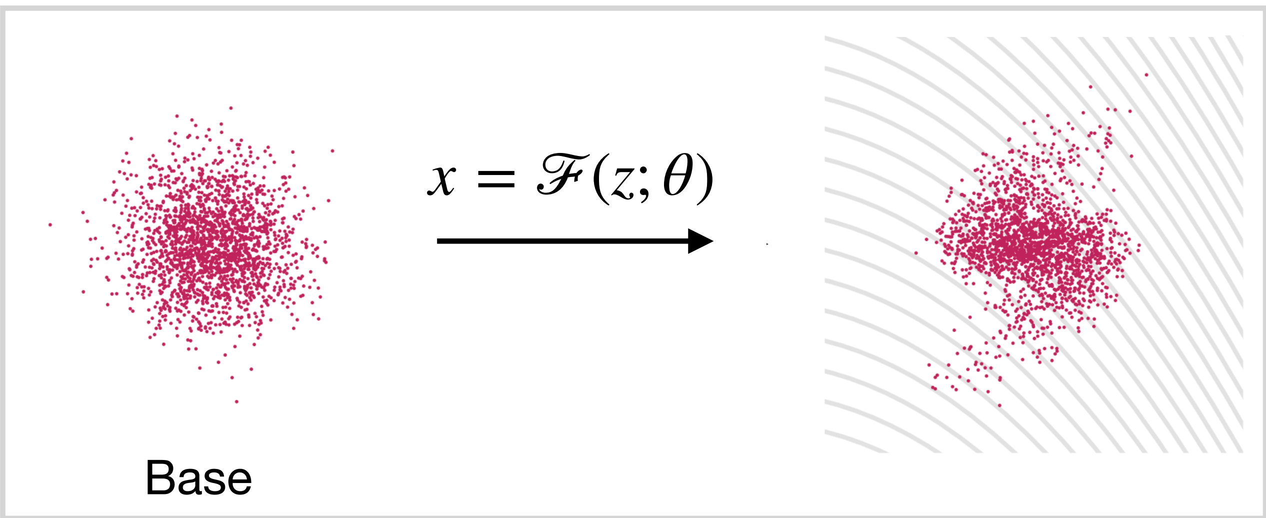

See Roussel et al., PRAB (2025) [https://doi.org/10.1103/PhysRevAccelBeams.27.094601]

Generative Phase Space Reconstruction (GPSR)

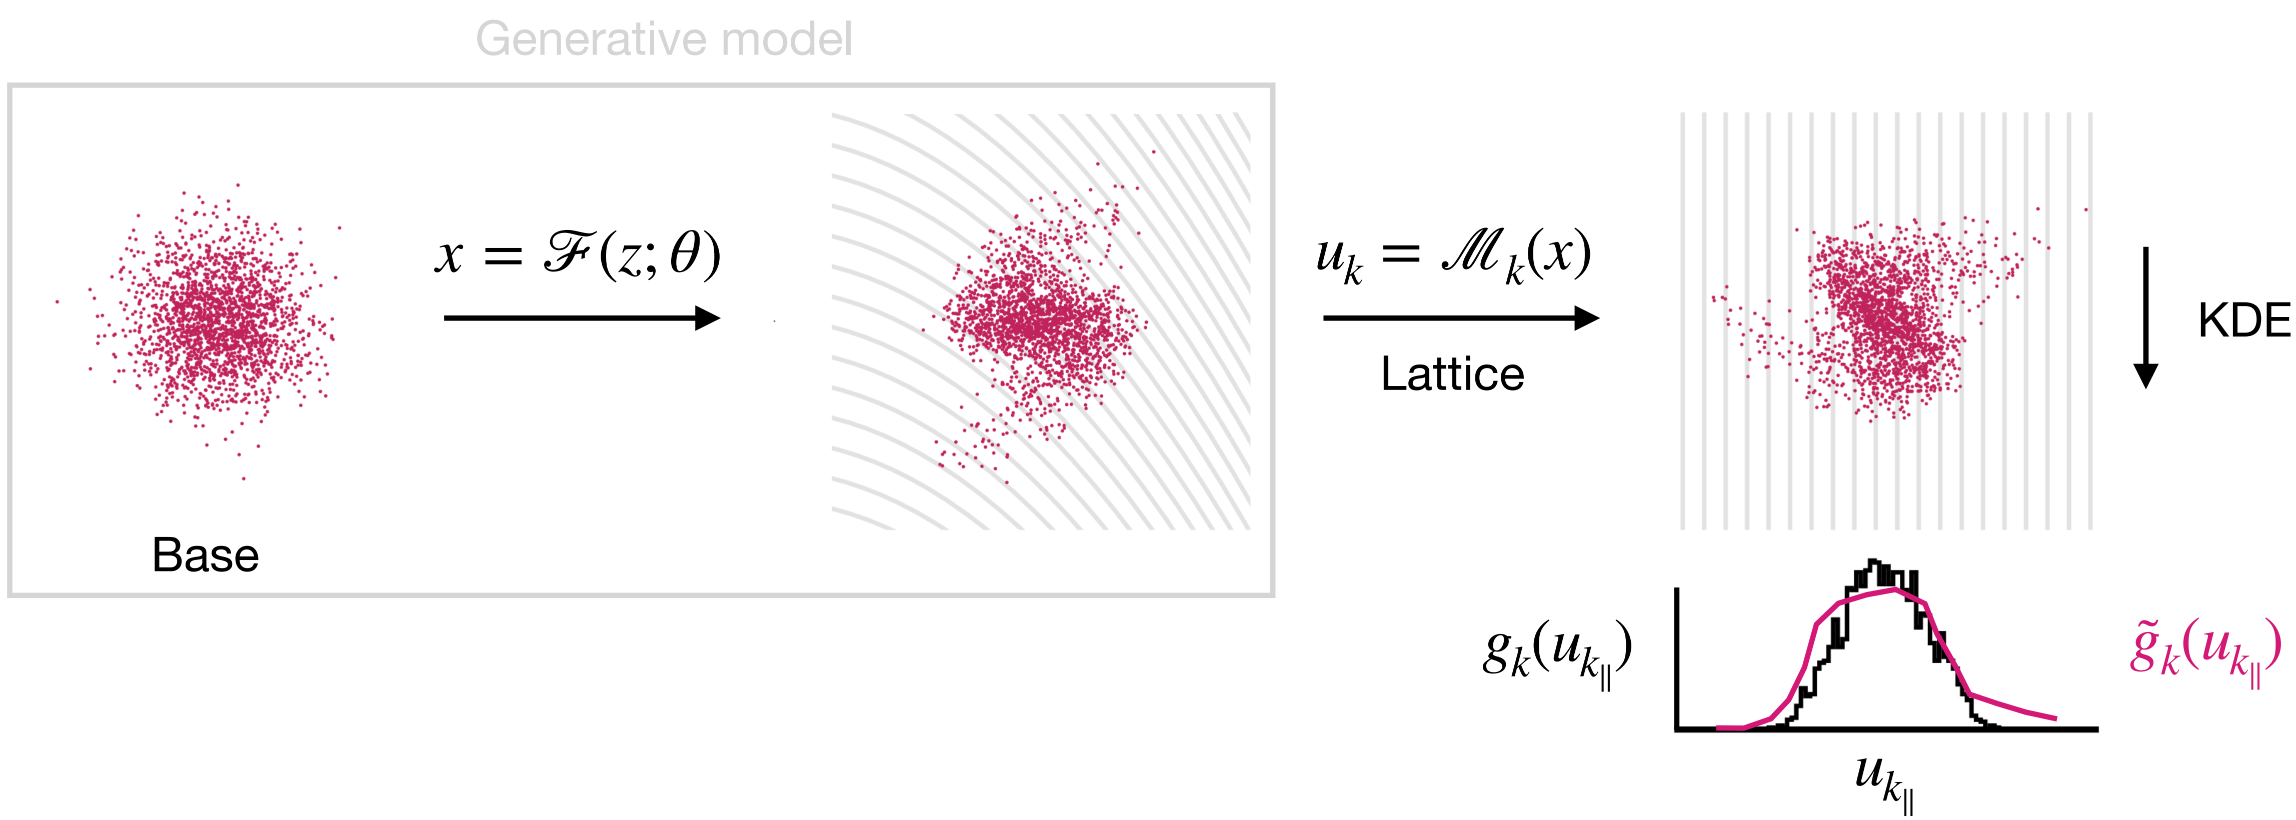

See Roussel et al., PRAB (2025) [https://doi.org/10.1103/PhysRevAccelBeams.27.094601]

Generative Phase Space Reconstruction (GPSR)

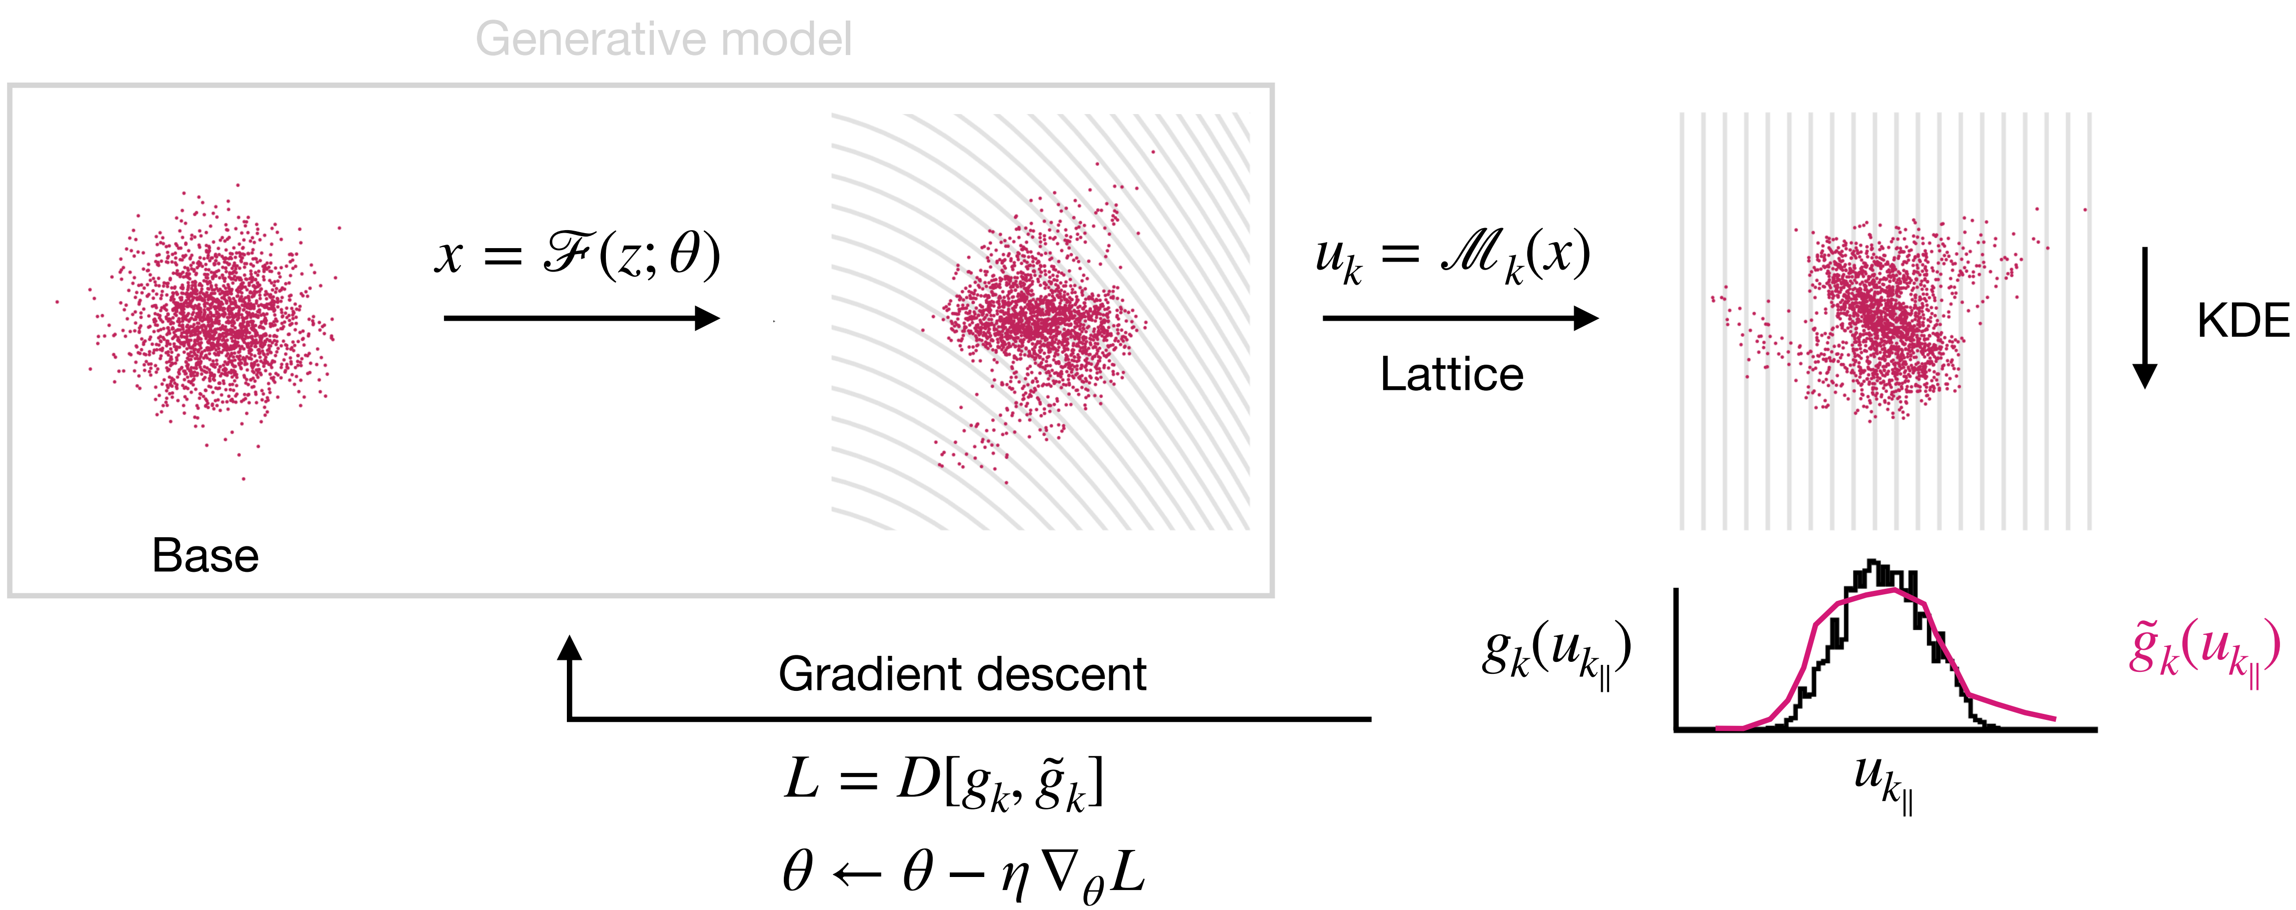

See Roussel et al., PRAB (2025) [https://doi.org/10.1103/PhysRevAccelBeams.27.094601]

Normalizing flows enable simultaneous sampling and density evaluation

$\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is differentiable, invertible map from z to x.

$$\log p(x) = \log p(z) - \log | J_{\mathcal{F}}(z) |$$

$$J_{\mathcal{F}} = \frac{dx}{dz} = \begin{bmatrix} \frac{\partial x_1}{\partial z_1} & \cdots & \frac{\partial x_1}{\partial z_N} \\ \vdots & \ddots & \vdots \\ \frac{\partial x_N}{\partial z_1} & \cdots & \frac{\partial x_N}{\partial z_N} \end{bmatrix}$$

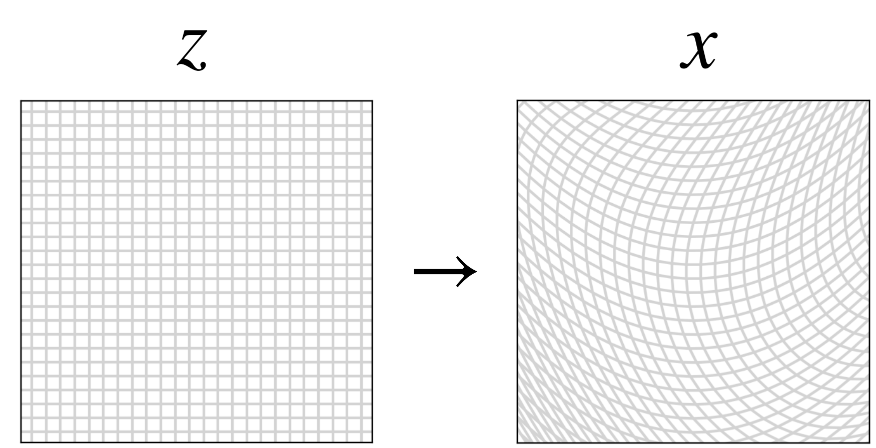

Normalizing flows can estimate entropy

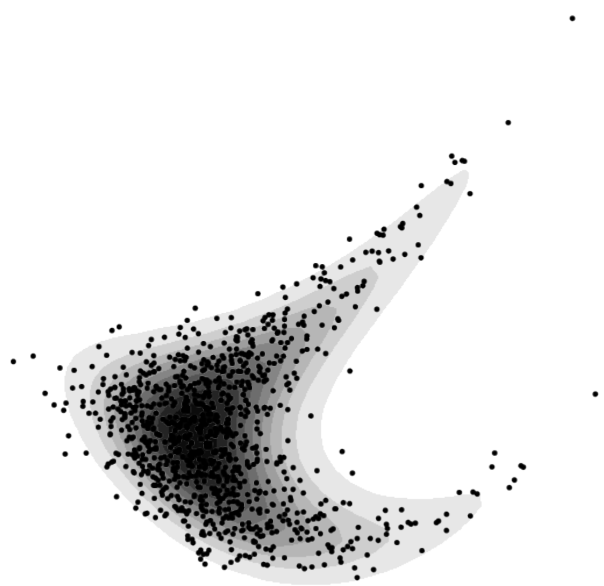

$\{x^{(i)}\} \sim p(x)$

$$S[p(x), p_*(x)] = -\int p(x) \log \left(p(x)/p_*(x)\right) dx$$

$$= -\mathbb{E}_{p(x)}\left[\log \left(p(x)/p_*(x)\right)\right]$$

$$\approx -\frac{1}{N}\sum_{i=1}^{N} \log(p(x^{(i)})/p_*(x^{(i)}))$$

OAK RIDGE National Laboratory

Normalizing flows can estimate entropy

Published as a conference paper at ICLR 2017

MAXIMUM ENTROPY FLOW NETWORKS

Gabriel Loaiza-Ganem*, Yuanjun Gao* & John P. Cunningham
Department of Statistics
Columbia University
New York, NY 10027, USA
{gl2480,yg2312,jpc2181}@columbia.edu

ABSTRACT

Maximum entropy modeling is a flexible and popular framework for formulating statistical models given partial knowledge. In this paper, rather than the traditional method of optimizing over the continuous density directly, we learn a smooth and invertible transformation that maps a simple distribution to the desired maximum entropy distribution. Doing so is nontrivial in that the objective being maximized (entropy) is a function of the density itself. By exploiting recent developments in normalizing flow networks, we cast the maximum entropy problem into a finite-dimensional constrained optimization, and solve the problem by combining stochastic optimization with the augmented Lagrangian method. Simulation results demonstrate the effectiveness of our method, and applications to finance and computer vision show the flexibility and accuracy of using maximum entropy flow networks.

Penalty method not ideal but (in my experience) can generate pretty good solutions

$$L = -S[p, p_*] + \mu D[g_k, \tilde{g}_k]$$

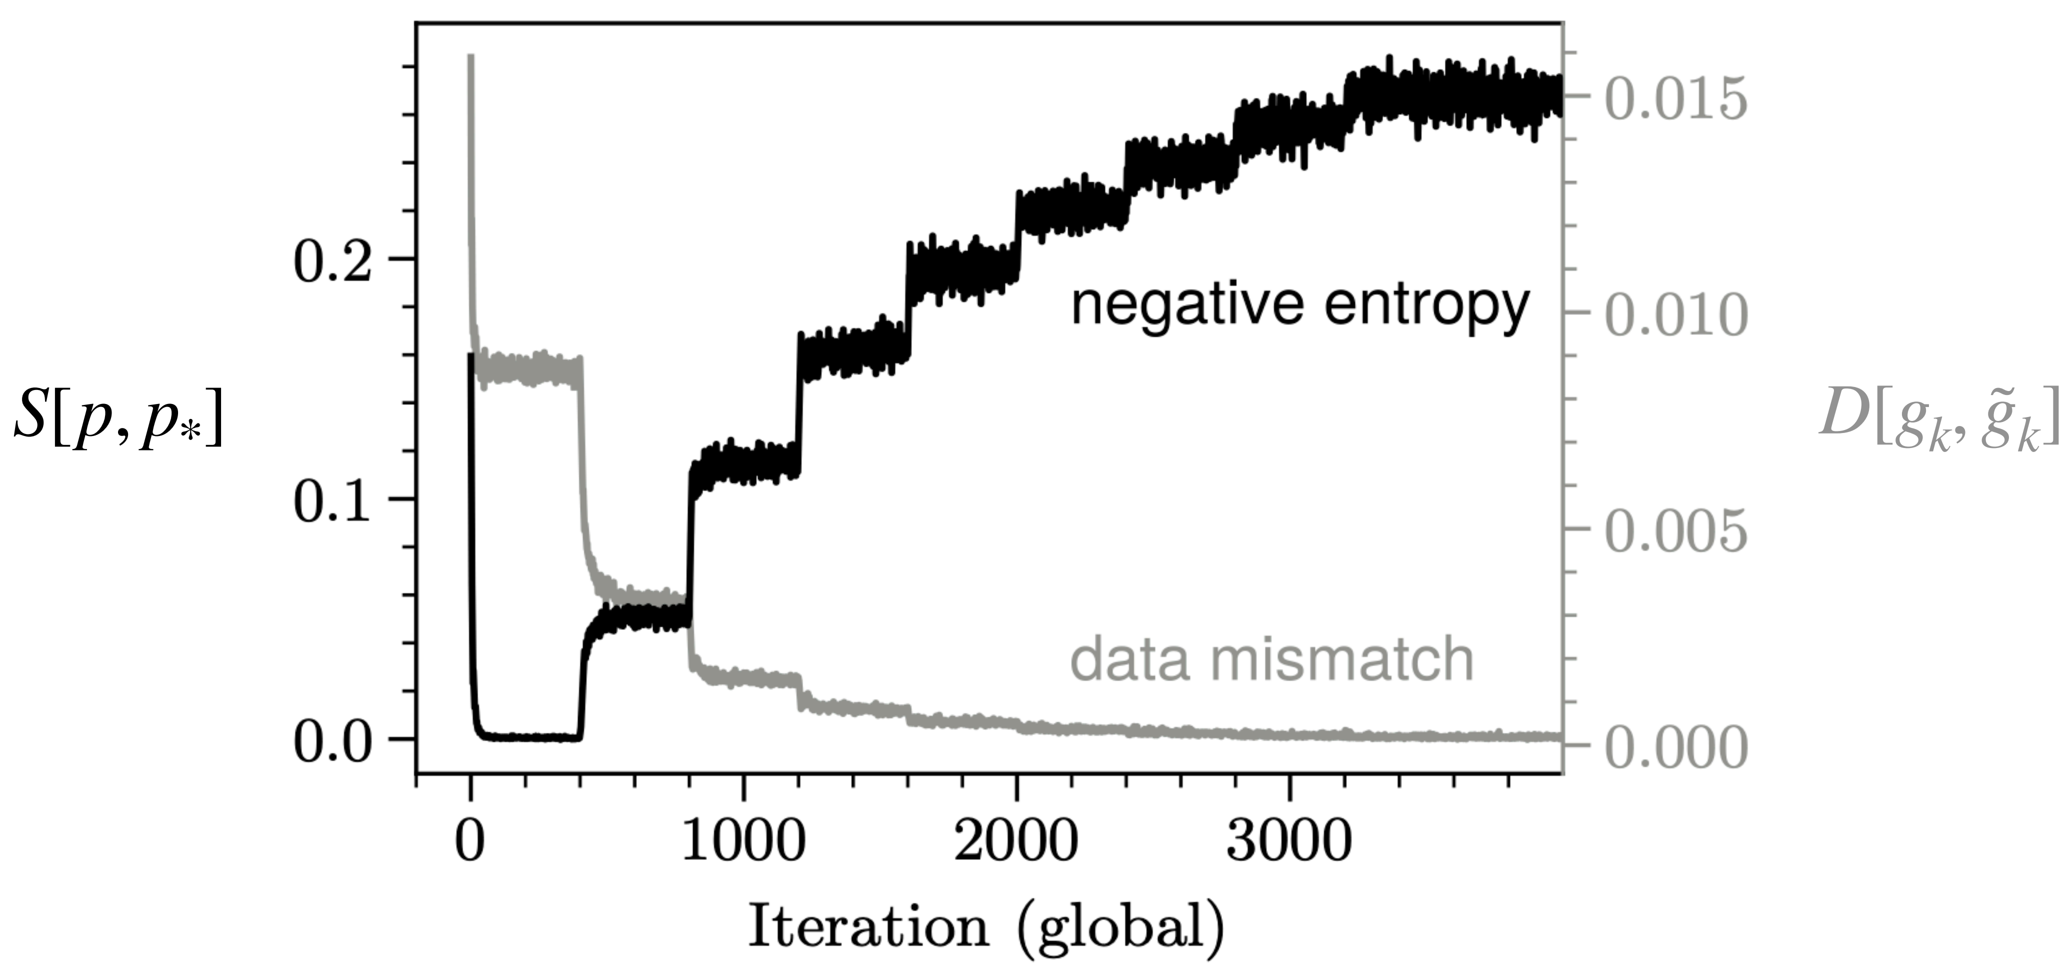

OAK RIDGE National Laboratory

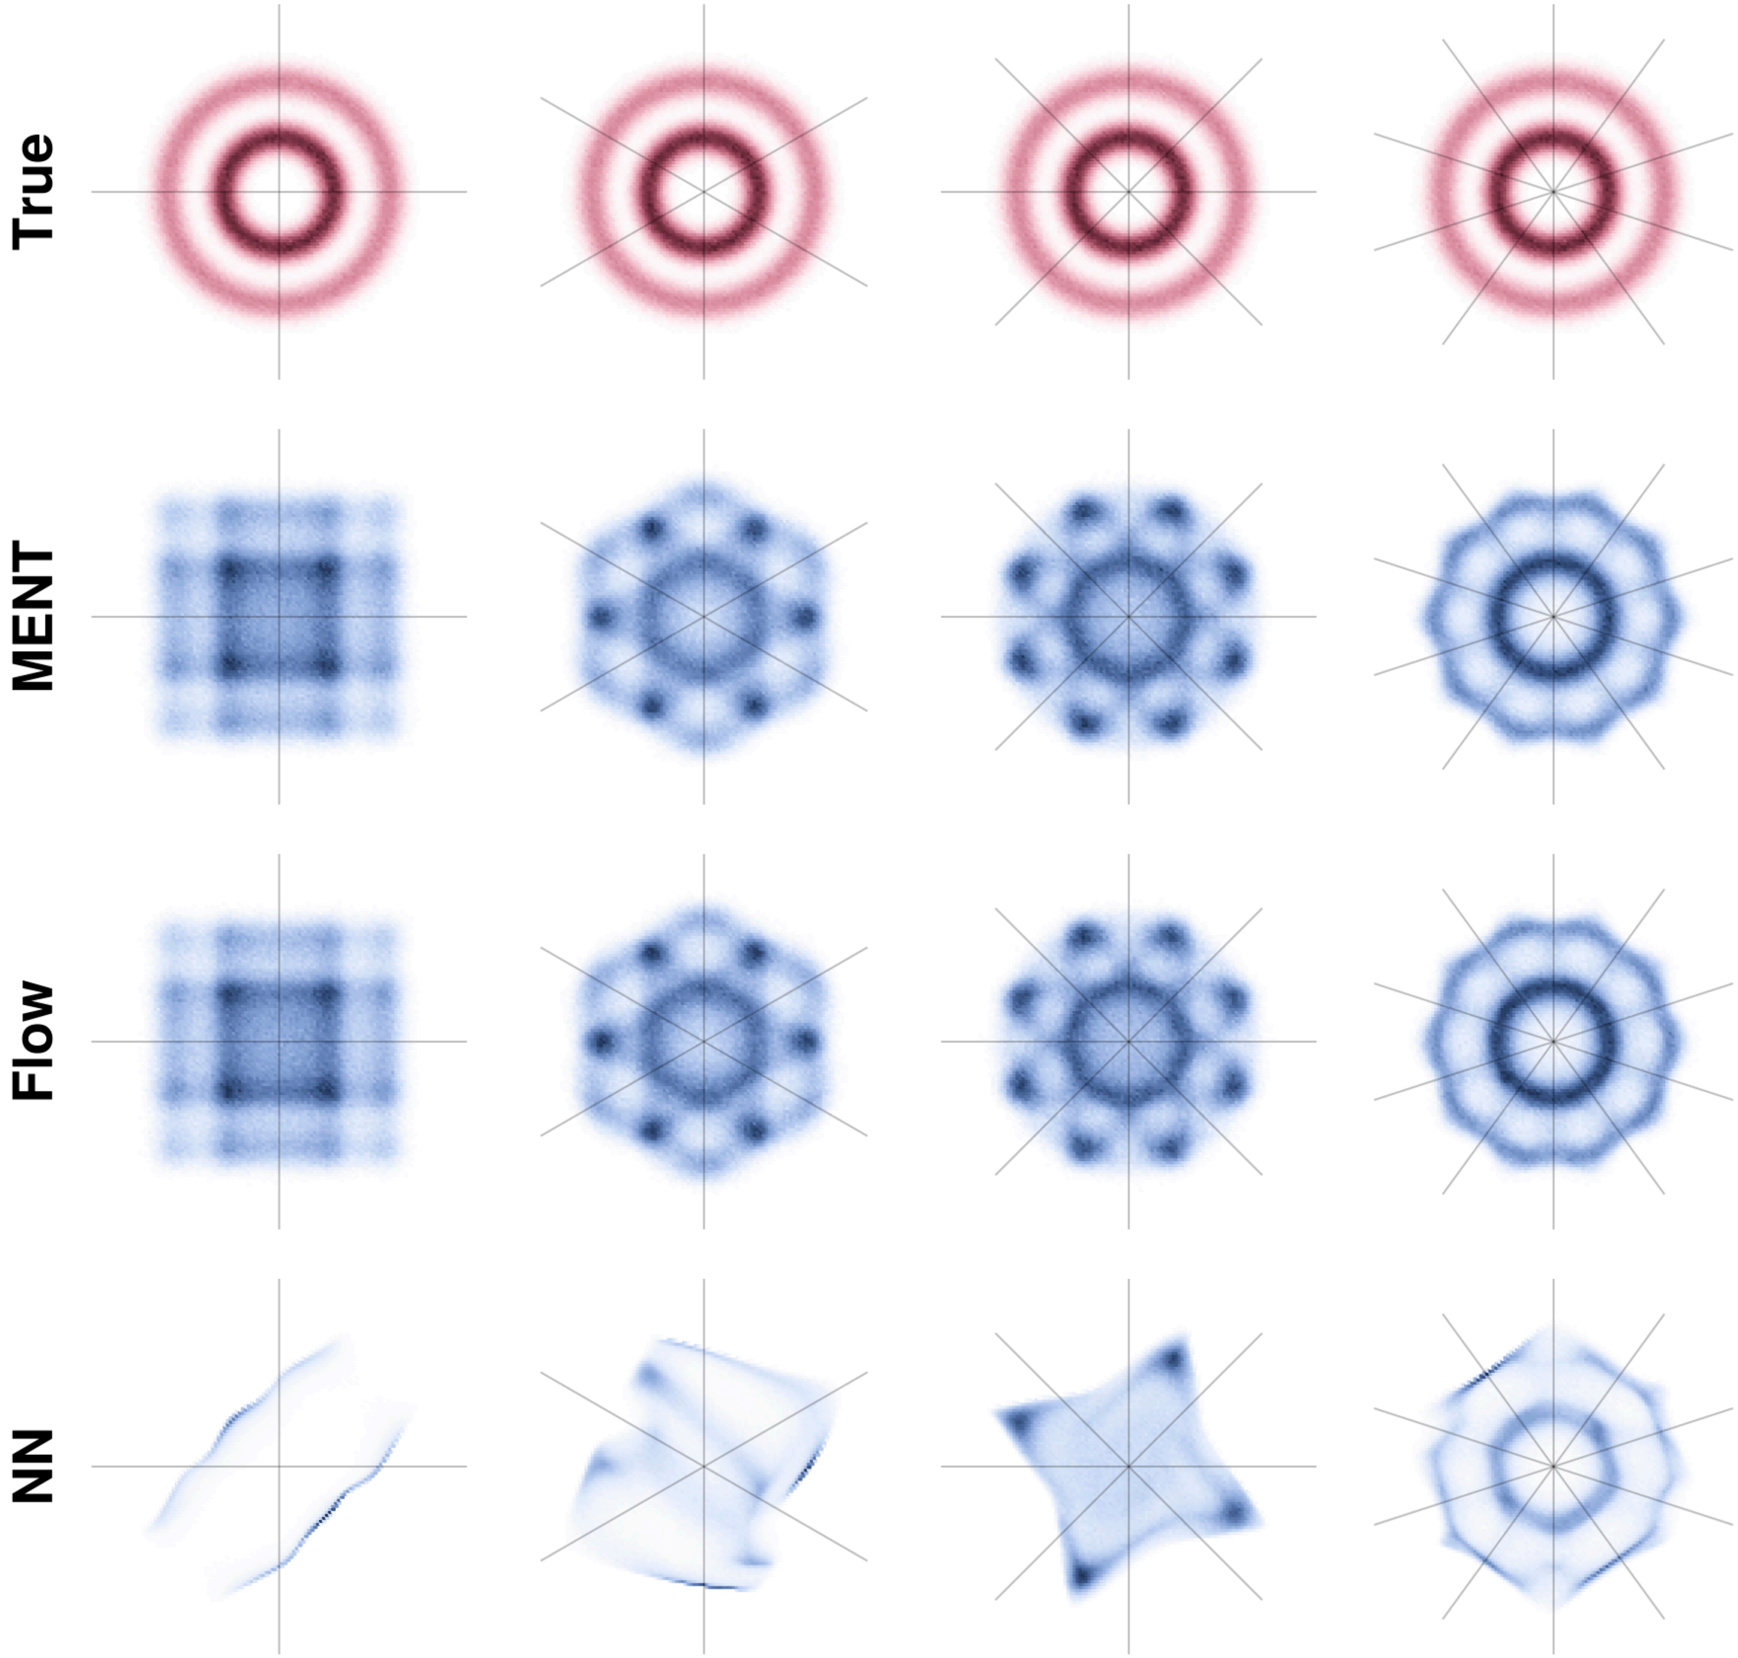
True
MENT
Flow
NN

Toy problem: 6D “rings” distribution projected along random 1D directions

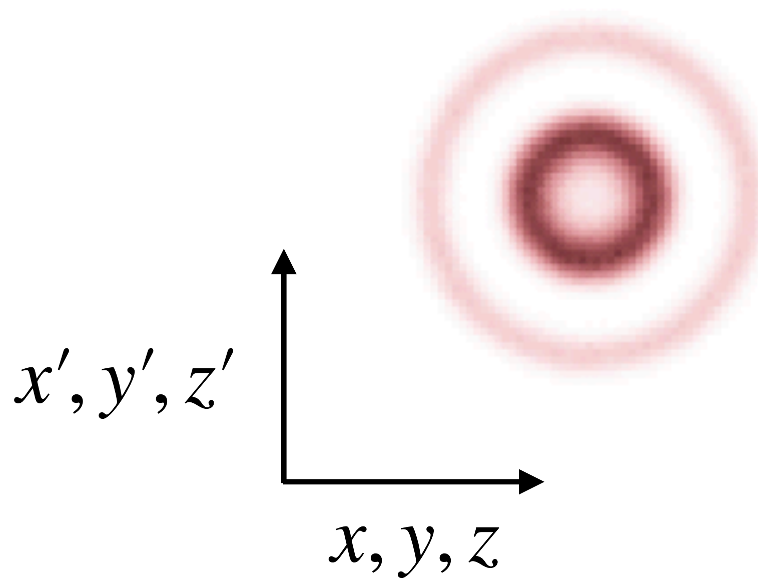

Oak Ridge National Laboratory

Toy problem: 6D "rings" distribution projected along random 1D directions

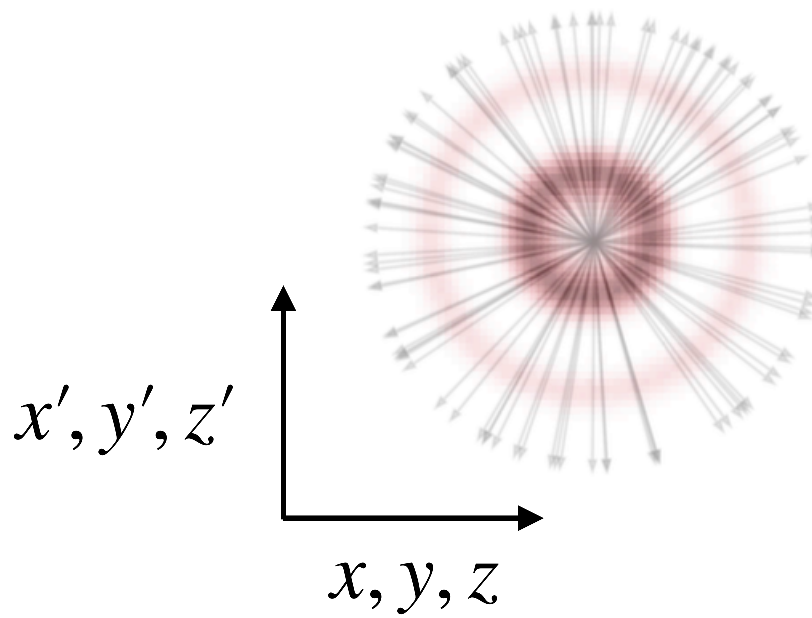

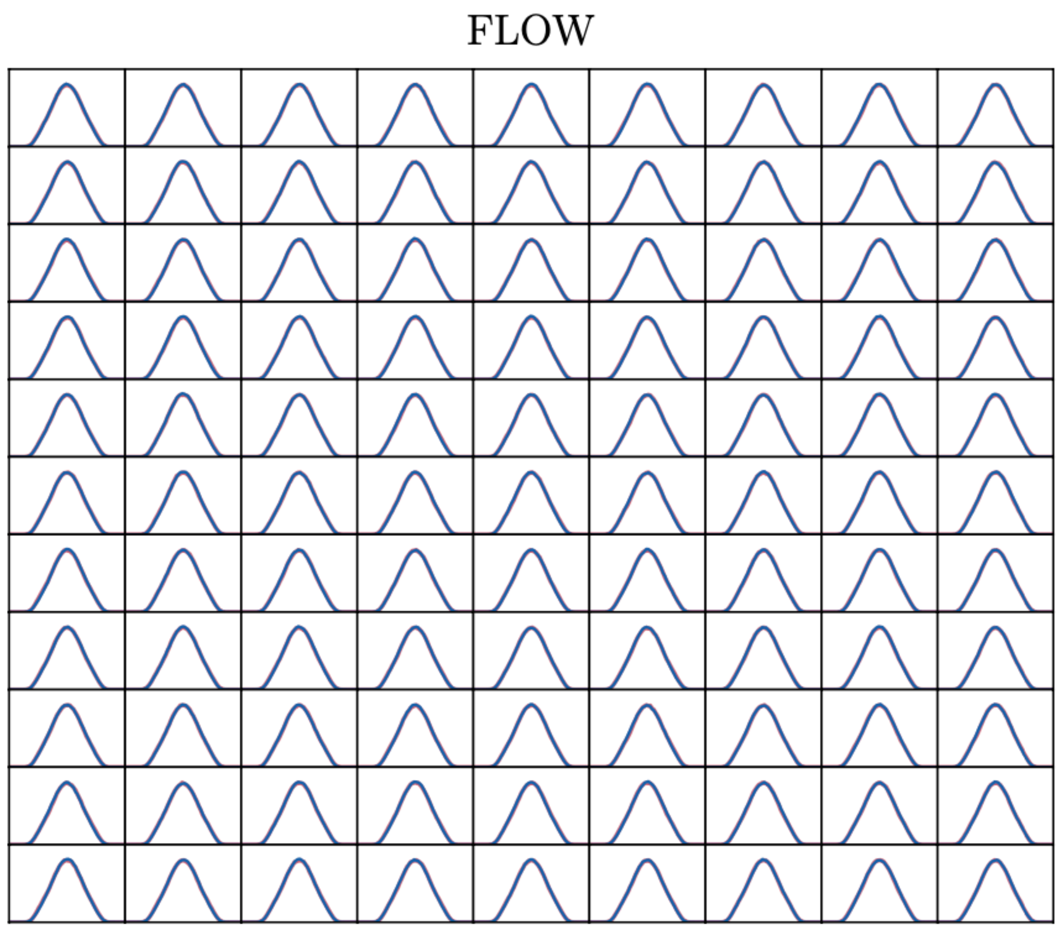
FLOW

(a)

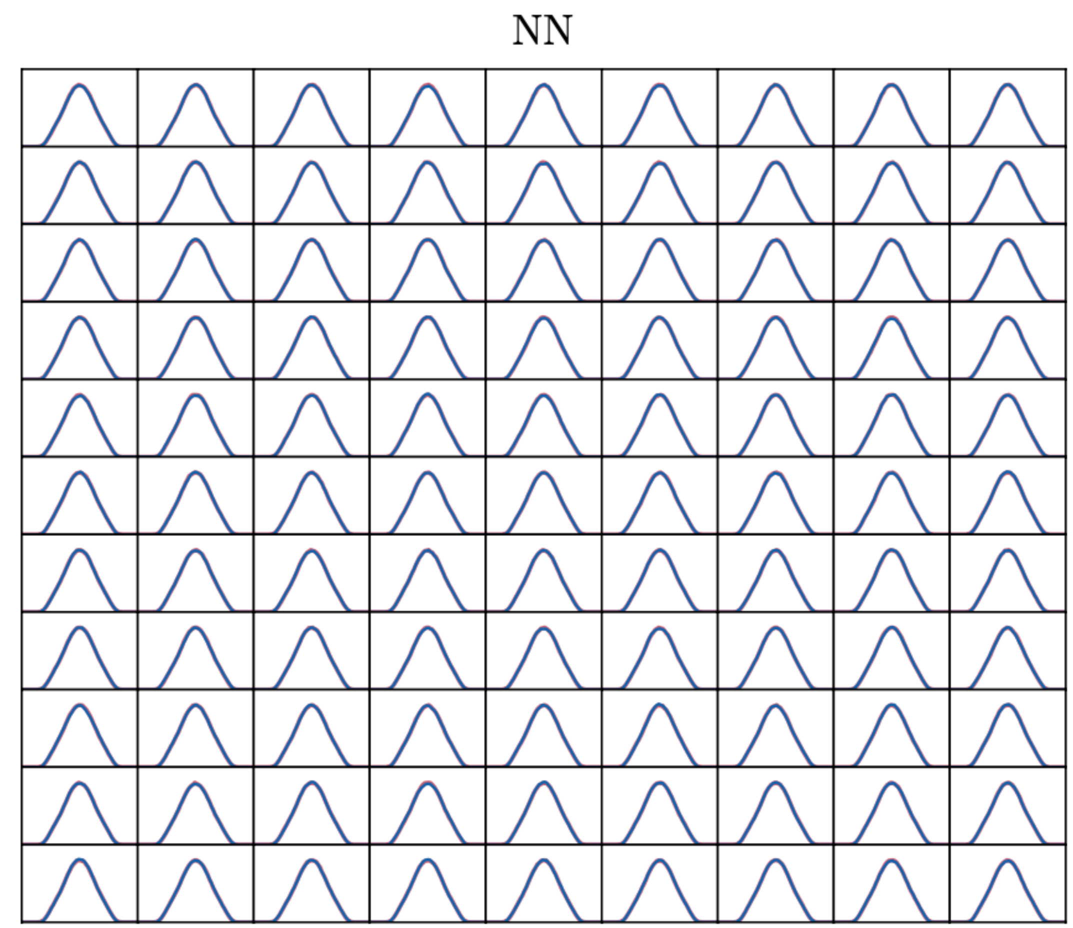
NN

(b)

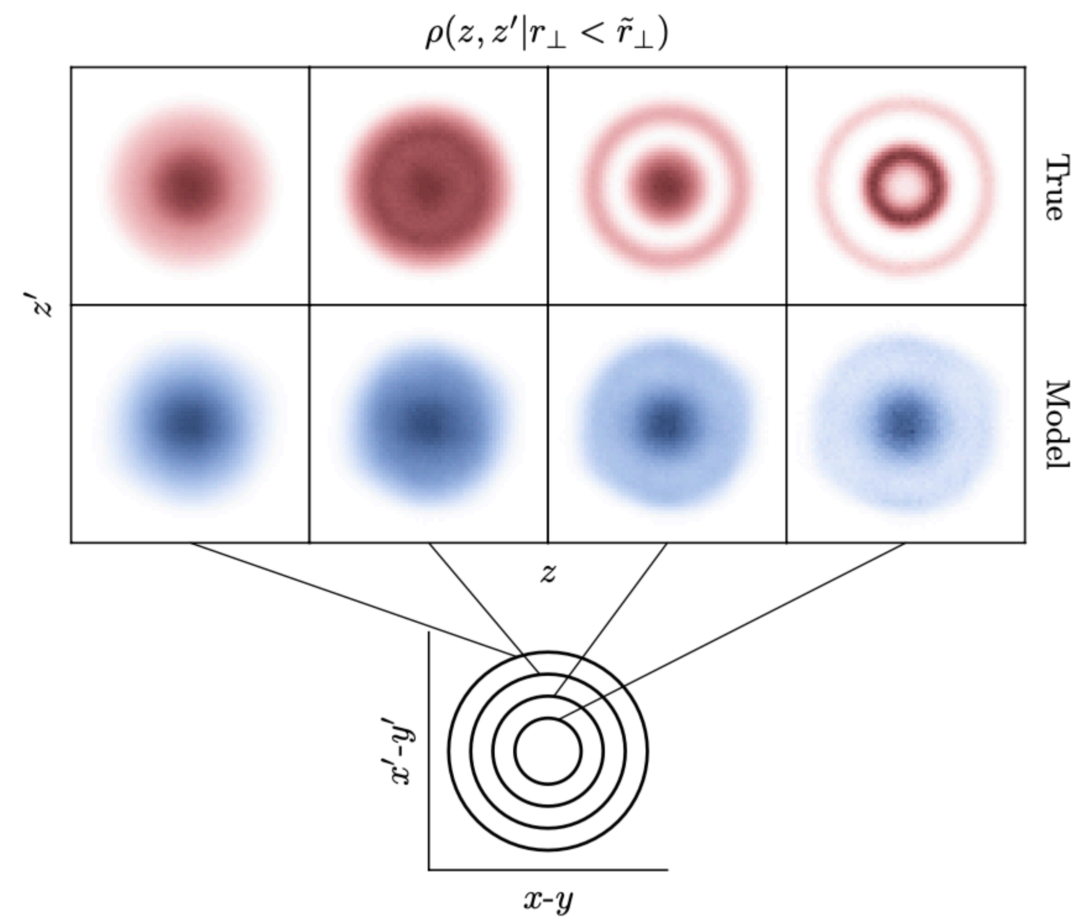
$\rho(z, z'|r_\perp < \tilde{r}_\perp)$
True
Model
z'
z
x'-y'
x-y

(c)

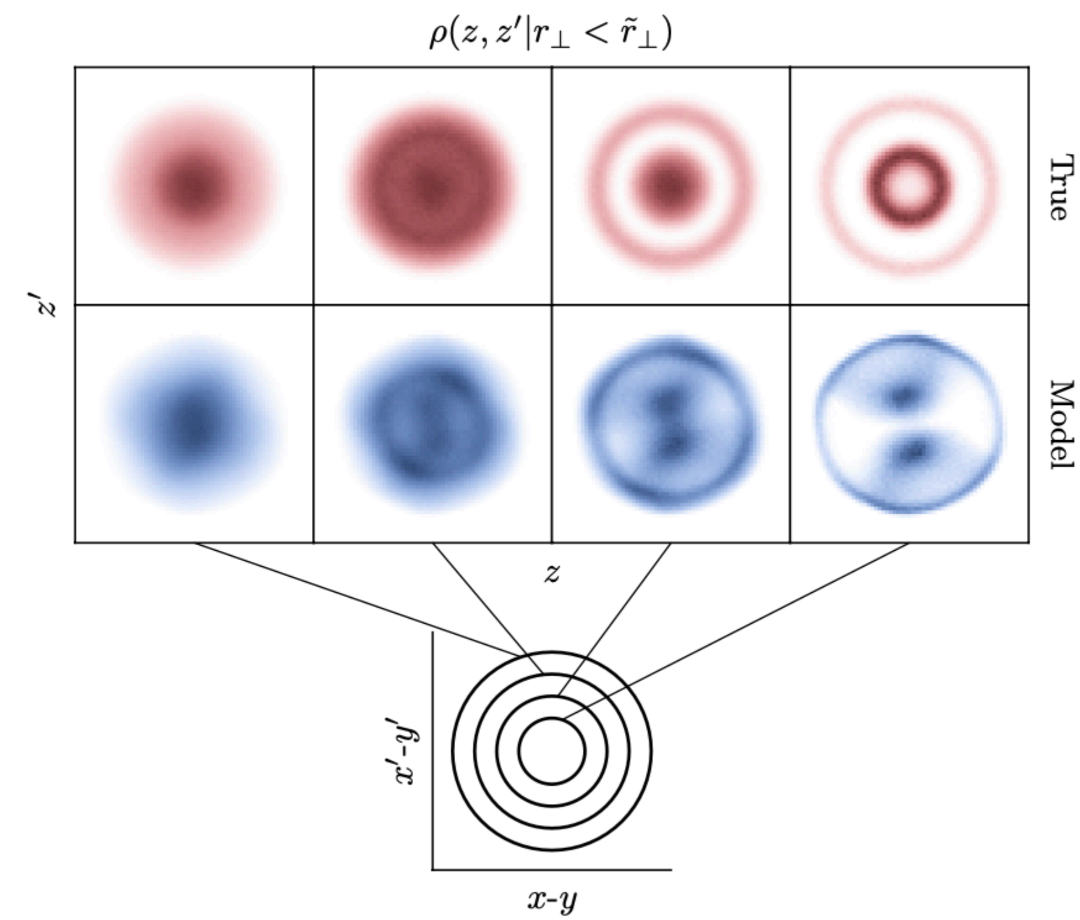
$\rho(z, z'|r_\perp < \tilde{r}_\perp)$
True
Model
z'
z
x'-y'
x-y

(d)

PHYSICAL REVIEW RESEARCH **6**, 033163 (2024)

High-dimensional maximum-entropy phase space tomography using normalizing flows

Austin Hoover *
Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830, USA

Jonathan C. Wong
Institute of Modern Physics, Chinese Academy of Sciences, Lanzhou 730000, China

(Received 31 May 2024; accepted 27 July 2024; published 12 August 2024)

Particle accelerators generate charged-particle beams with tailored distributions in six-dimensional position-momentum space (phase space). Knowledge of the phase space distribution enables model-based beam optimization and control. In the absence of direct measurements, the distribution must be tomographically reconstructed from its projections. In this paper, we highlight that such problems can be severely underdetermined and that entropy maximization is the most conservative solution strategy. We leverage *normalizing flows*—invertible generative models—to extend maximum-entropy tomography to six-dimensional phase space and perform numerical experiments to validate the model's performance. Our numerical experiments demonstrate consistency with exact two-dimensional maximum-entropy solutions and the ability to fit complicated six-dimensional distributions to large measurement sets in reasonable time.

DOI: 10.1103/PhysRevResearch.6.033163

Summary of GPSR

- Efficient method to optimize high-dimensional distributions when differentiable simulations are available.
- Flexible — different kinds of data.
- Future: more sophisticated constrained optimization
- Future: better/non-invertible flows

MENT: classical approach

Maximum-entropy PDF is exponential of product of Lagrange Multipliers

Lagrangian:

$$\Psi = S[p(x), p_*(x)] + \sum_k \int G_k[p(x)]\lambda_k(u_{k_\parallel})du_{k_\parallel}$$

Enforce $\delta\Psi = 0$ with respect to p and λ_k:

$$p(x) = p_*(x) \prod_k \exp\left[\lambda_k(u_{k_\parallel}(x))\right]$$

$$p(x) = p_*(x) \prod_k h_k(u_{k_\parallel}(x))$$

Maximum-entropy PDF is exponential of product of Lagrange Multipliers

Lagrangian:

$$\Psi = S[p(x), p_*(x)] + \sum_k \int G_k[p(x)]\lambda_k(u_{k_\parallel})du_{k_\parallel}$$

Enforce $\delta\Psi = 0$ with respect to p and λ_k:

$$p(x) = p_*(x)\prod_k \exp\left[\lambda_k(u_{k_\parallel}(x))\right]$$

$$p(x) = p_*(x)\prod_k h_k(u_{k_\parallel}(x))$$

Prior

Positive functions defined on measurement axes

OAK RIDGE National Laboratory

Nonlinear Gauss-Seidel iterations used to solve for unknown functions: **simple algorithm!**

$$h_k^{(i+1)} = h_k^{(i)} \left[(1-\omega) + \omega \left(g_k/\tilde{g}_k\right)\right]$$

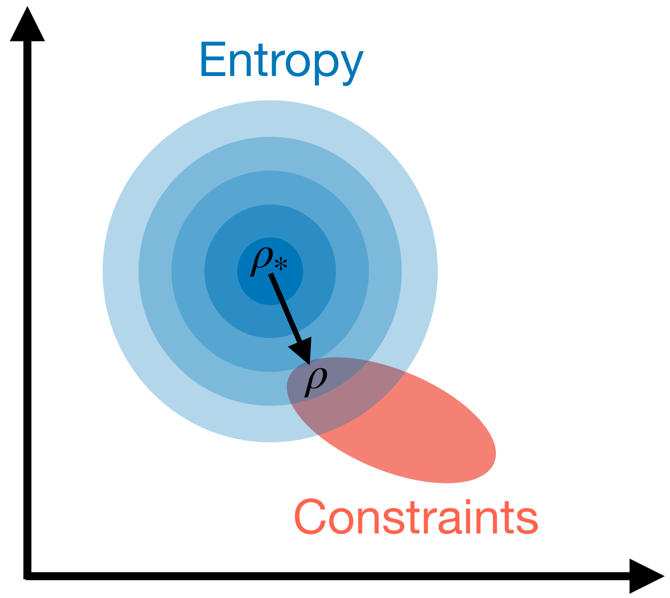

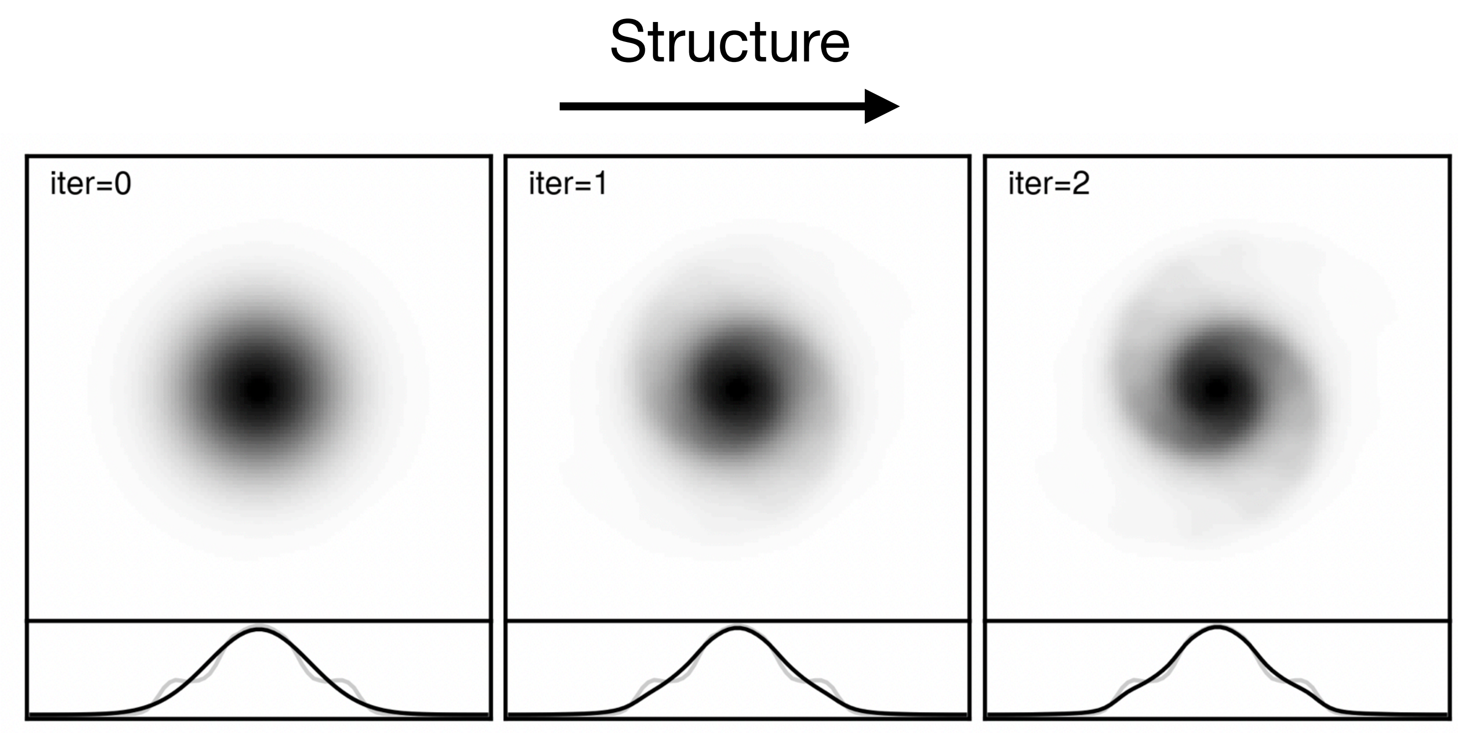

Standard implementation uses numerical integration (**reverse mode**) to compute projections

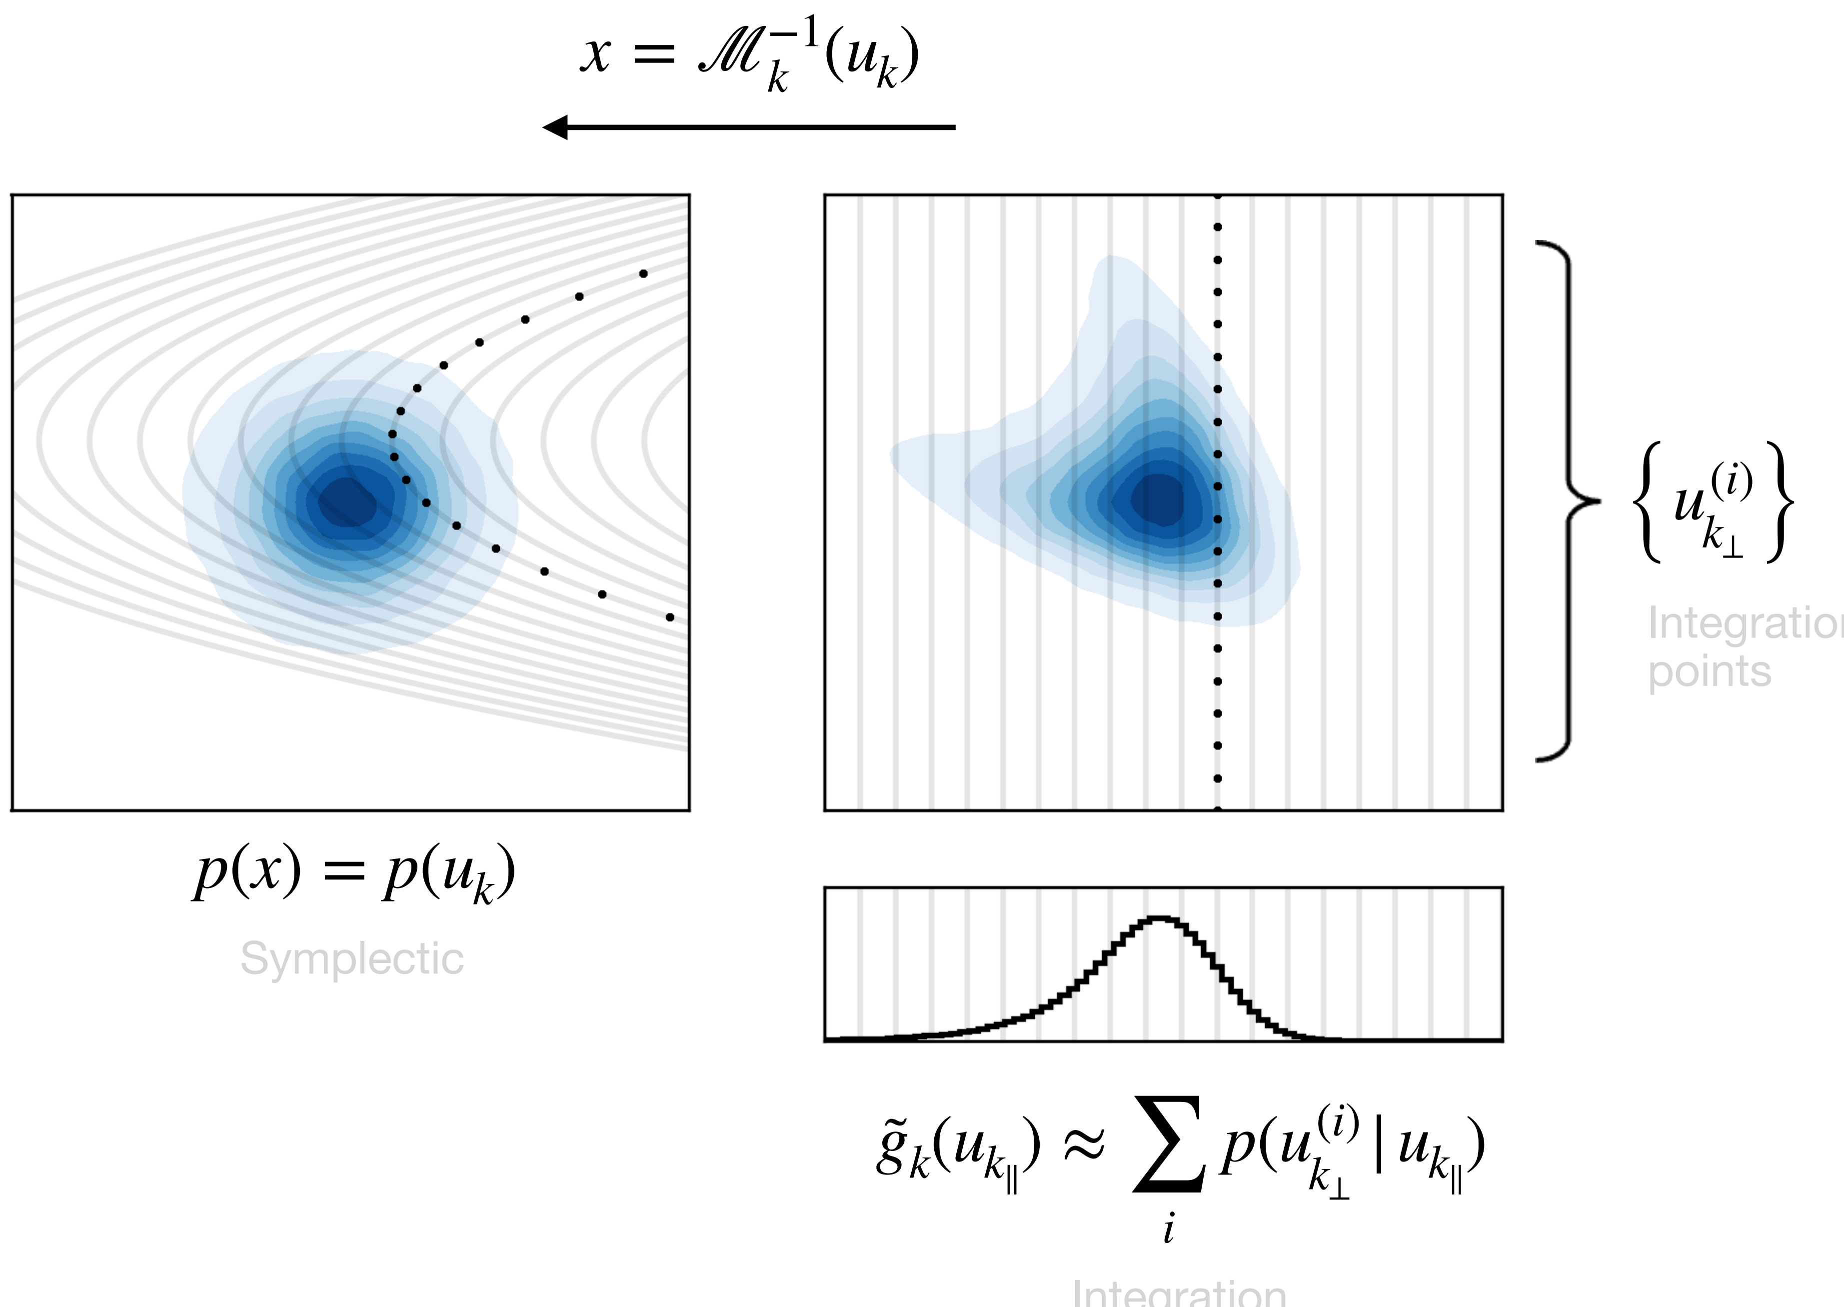

OAK RIDGE National Laboratory

I propose particle sampling (**forward mode**) for high-dimensional problems

$$u_k = \mathscr{M}_k(x)$$

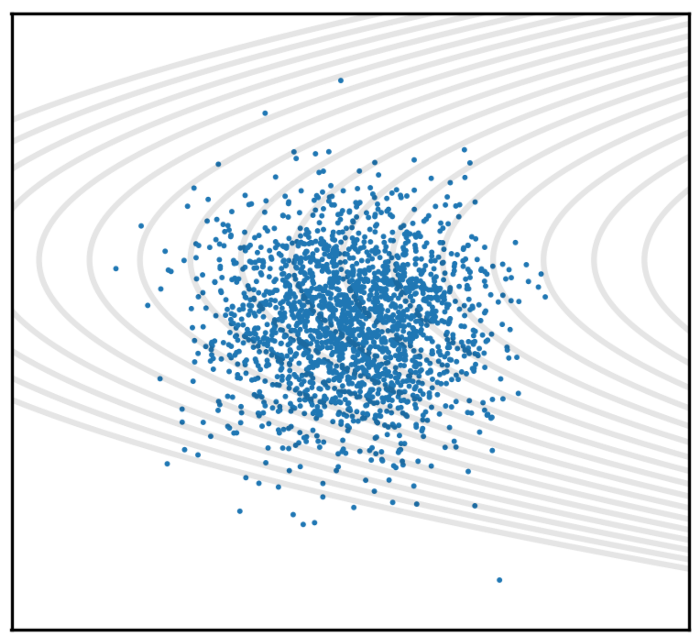

$$\{x^{(i)}\} \sim p(x)$$

Sampling

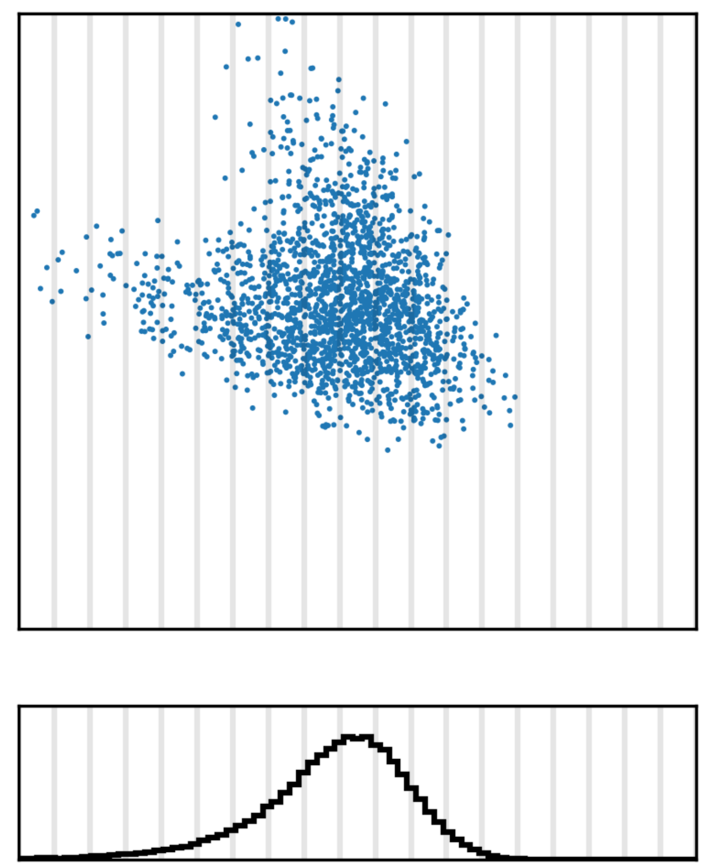

$$\tilde{g}_k(u_{k_\parallel}) \approx \text{HIST}\left\{ u_{k_\parallel}^{(i)} \right\}$$

Density estimation

Metropolis-Hastings is gridless sampling algorithm with long-run guarantees

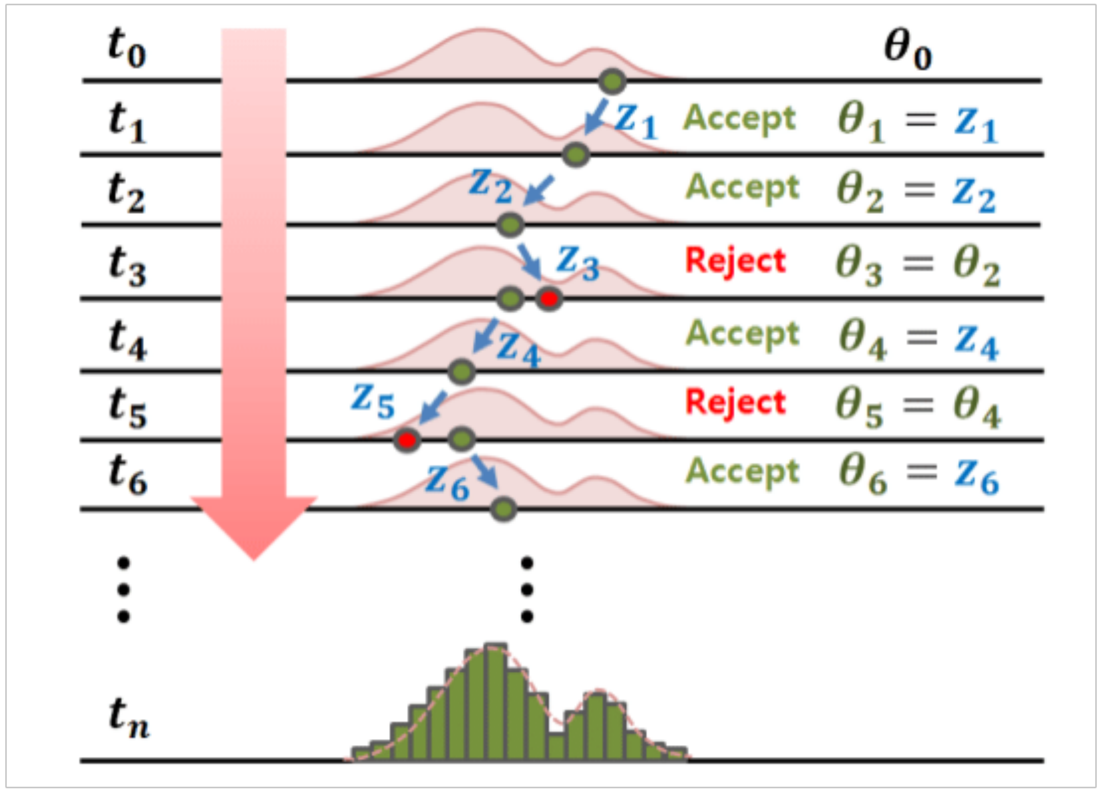

Requires hand-tuning of jumping distribution, and it's not always clear when it converges…

But let's try it.

Test: 6D Gaussian mixture distribution

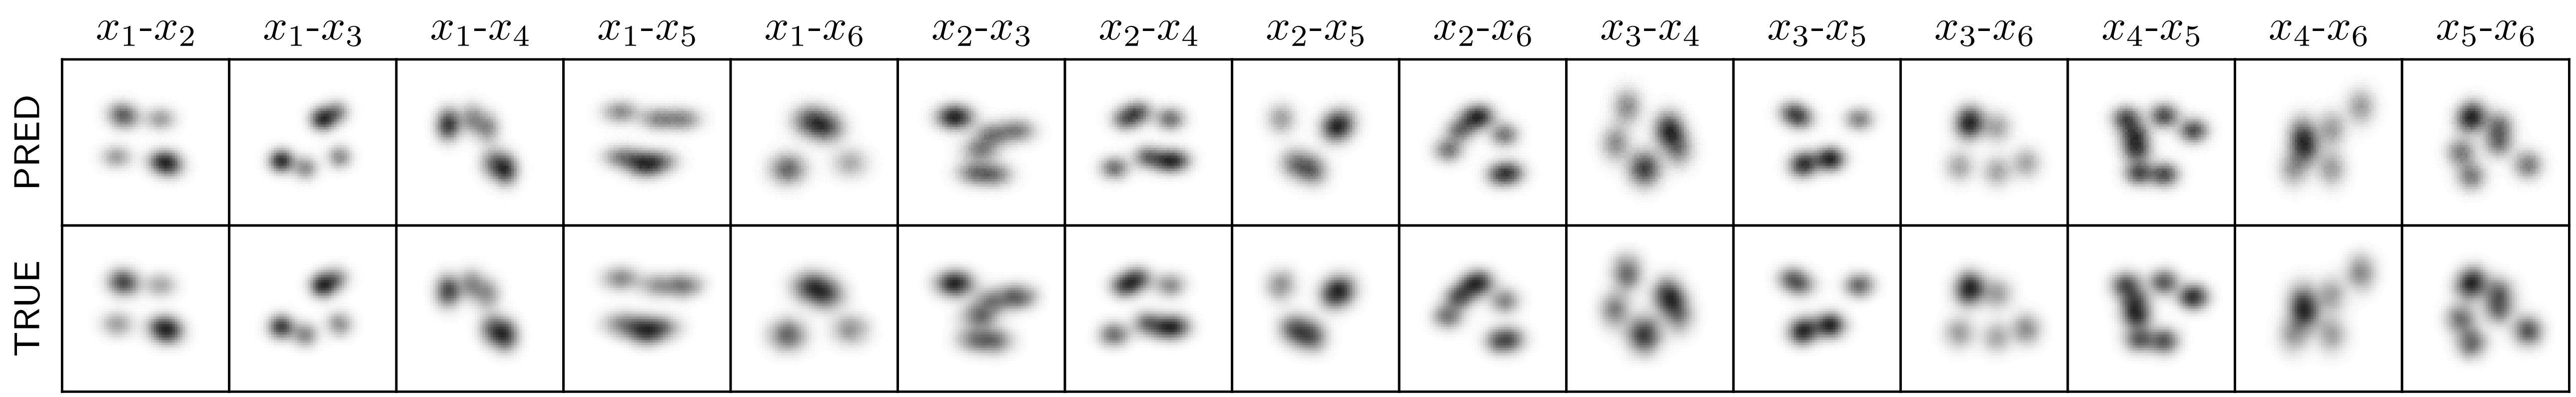

Reverse (50^4 integration, 50^2 measurement grid): **24 hour** / iteration

Forward (500,000 samples, 1000 chains) : **30 sec** / iteration

Test: 6D Argonne Wakefield Accelerator dataset

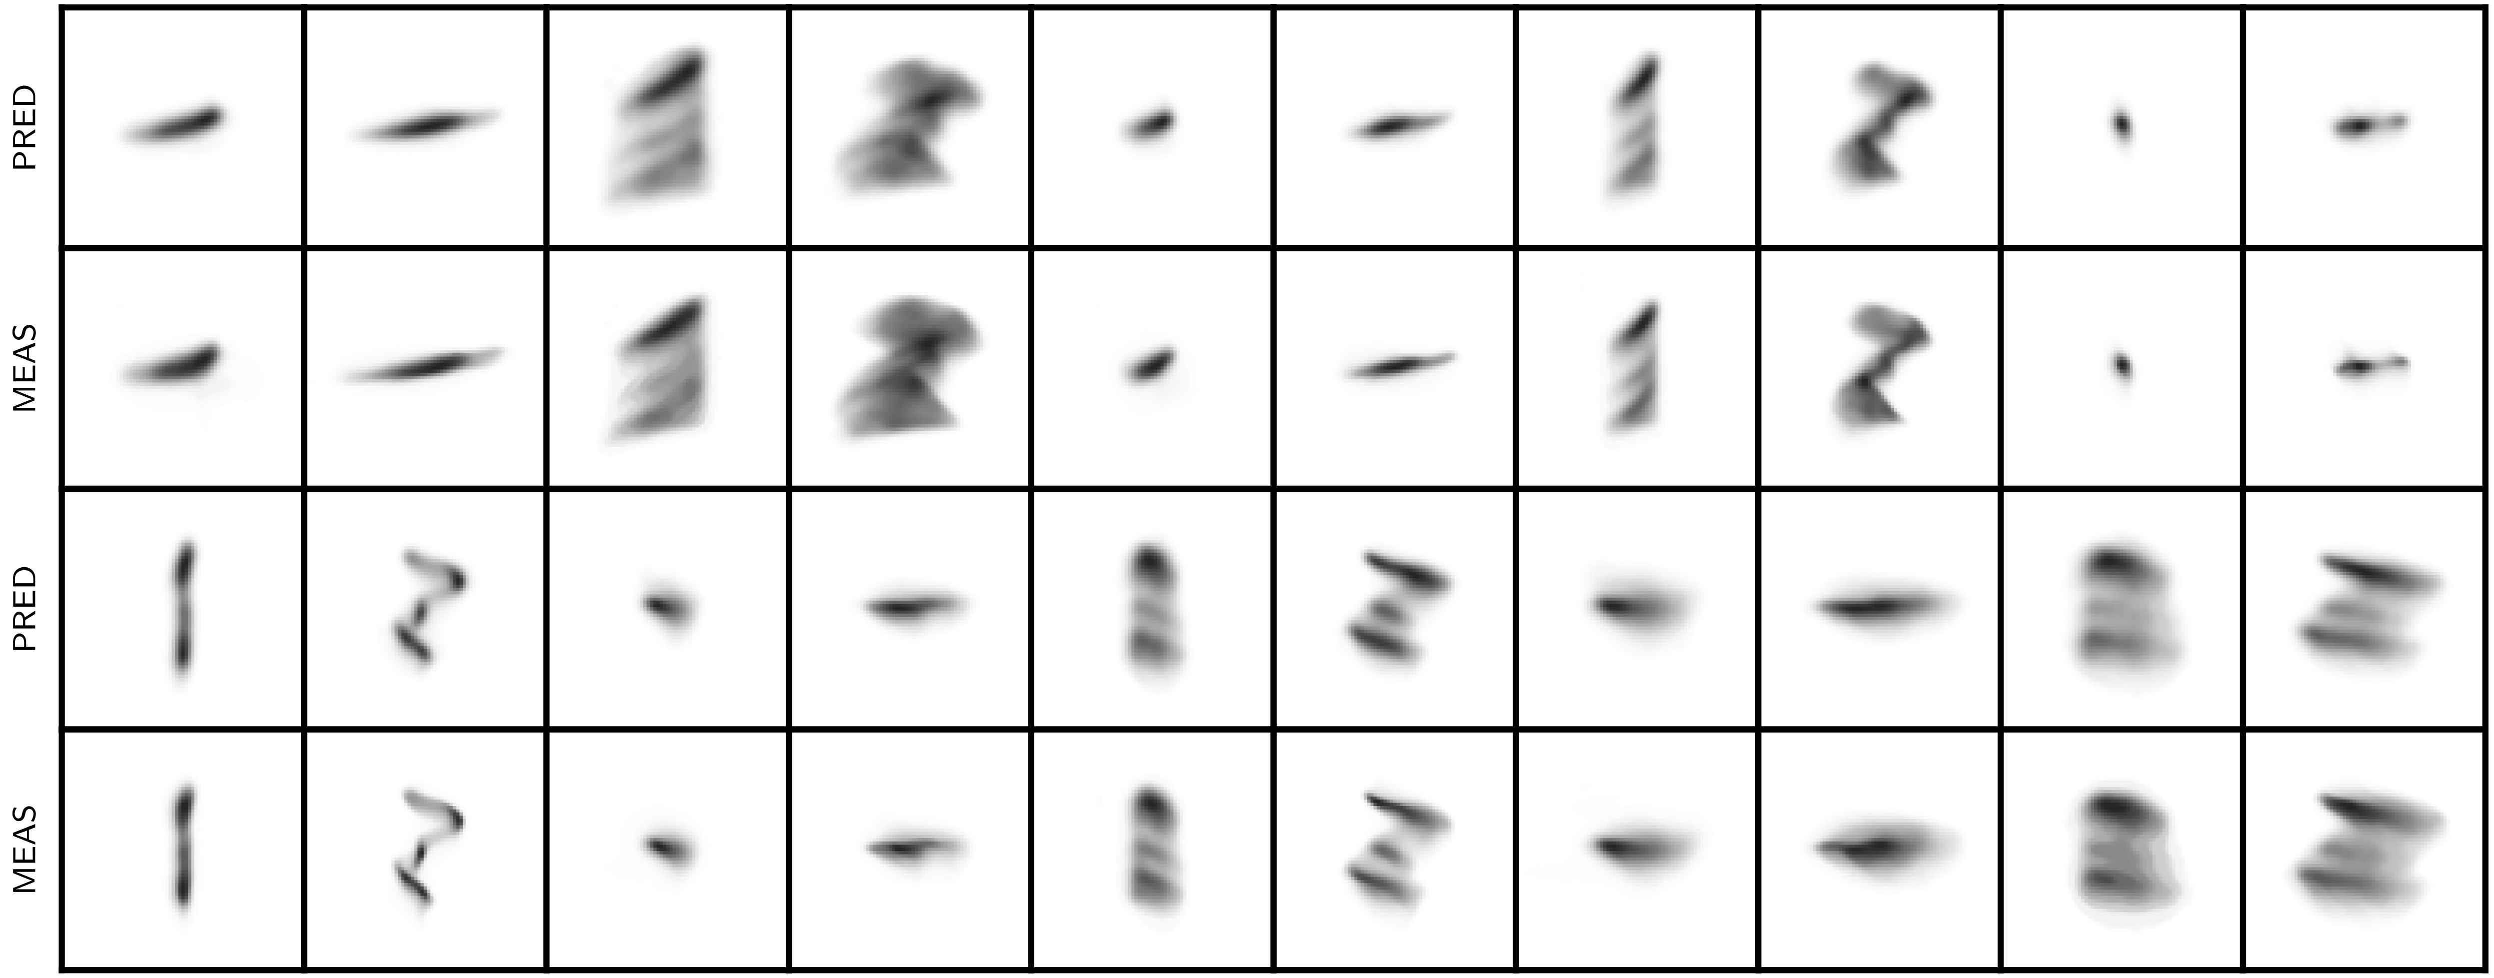

OAK RIDGE National Laboratory

https://doi.org/10.1103/PhysRevAccelBeams.27.094601

PHYSICAL REVIEW ACCELERATORS AND BEAMS **28,** L084601 (2025)

Letter | Editors' Suggestion

N-dimensional maximum-entropy tomography via particle sampling

Austin Hoover*

Oak Ridge National Laboratory, Oak Ridge, Tennessee 37830, USA

(Received 4 October 2024; accepted 14 May 2025; published 7 August 2025)

We propose a modified maximum-entropy (MENT) algorithm for six-dimensional phase space tomography. The algorithm uses particle sampling and low-dimensional density estimation to approximate large sets of high-dimensional integrals in the original MENT formulation. We implement this approach using Markov Chain Monte Carlo (MCMC) sampling techniques and demonstrate convergence of six-dimensional MENT on both synthetic and measured data.

DOI: 10.1103/zl2h-3v32

Summary of MENT

- Elegant — minimal storage, simple algorithm.
- Flexible — no restrictions on distribution or transformations.
- Future: new sampling algorithm / tuning
- Future: GPU + parallel

Conclusion

Conclusion

- Several MaxEnt algorithms running in 2D/4D/6D.
- Algorithm improvements possible
- Currently working on uncertainty quantification (UQ), improved dynamic range (DR), self-fields

Conclusion

- Several MaxEnt algorithms running in 2D/4D/6D.
- Algorithm improvements possible
- Currently working on uncertainty quantification (UQ), improved dynamic range (DR), self-fields

- https://github.com/austin-hoover/ment
- https://github.com/roussel-ryan/gpsr
- Zenodo repositories linked in papers

Conclusion

- Several MaxEnt algorithms running in 2D/4D/6D.
- Algorithm improvements possible
- Currently working on uncertainty quantification (UQ), improved dynamic range (DR), self-fields

- https://github.com/austin-hoover/ment
- https://github.com/roussel-ryan/gpsr
- Zenodo repositories linked in papers

- Thanks for listening!
- **hooveram@ornl.gov**