

SEARCH FOR VLASOV-POISSON EQUILIBRIUM DISTRIBUTIONS USING GENERATIVE MODELS *

Austin Hoover[†]

Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA

Abstract

This work proposes the use of generative models and differentiable simulations to search for smooth equilibrium phase space distributions in arbitrary external fields. The model parameters are optimized using a two-term loss function: one term suppresses density fluctuations, while a second term pulls the solution toward some target value, for example, to encourage smoothness. As an initial demonstration of this approach, we search for a four-dimensional equilibrium distribution in a linear continuous-focusing lattice.

INTRODUCTION

This work proposes a method to search for equilibrium phase space distributions in periodic focusing channels. Equilibrium distributions are solutions to the Vlasov-Poisson (V-P) system:

$$\frac{\partial f}{\partial t} + \frac{\partial H}{\partial \mathbf{x}'} \cdot \frac{\partial}{\partial \mathbf{x}} - \frac{\partial H}{\partial \mathbf{x}} \cdot \frac{\partial f}{\partial \mathbf{x}'} = 0. \quad (1)$$

$$\frac{\partial}{\partial \mathbf{x}} \cdot \frac{\partial}{\partial \mathbf{x}} \phi(\mathbf{x}, s) = -\frac{q}{\epsilon_0} \int f(\mathbf{x}, \mathbf{x}', s) d\mathbf{x}'. \quad (2)$$

$$H(\mathbf{x}, \mathbf{x}', s) = \frac{1}{2} |\mathbf{x}'|^2 + U(\mathbf{x}, \mathbf{x}', s) + \frac{q}{mc^2 \beta^2 \gamma^3} \phi(\mathbf{x}, s). \quad (3)$$

Here f is the distribution function, H is the Hamiltonian, U is an external potential, ϕ is the self-generated electric potential, \mathbf{x} is the position, \mathbf{x}' is the momentum, s is the position on the reference trajectory, q is the particle charge, ϵ_0 is the electric constant, m is the particle mass, c is the speed of light, β is the synchronous particle velocity, and γ is the Lorentz factor.

Equilibrium distributions are functions of invariants. In continuous-focusing channels where $U(\mathbf{x}, \mathbf{x}', s) = U(\mathbf{x}, \mathbf{x}')$, the Hamiltonian is an invariant and there are an infinite number of solutions of the form $f = f(H)$. In s -dependent focusing channels, the Hamiltonian is not an invariant and solutions are more difficult to find. The Kurth distribution [1] is a solution for the special case of spherically or axially

* This manuscript has been authored by UT-Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains, and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>)

[†] hooveram@ornl.gov

symmetric linear fields. The only known solution in *arbitrary* linear fields is the Kapchinskij-Vladirskij (K-V) distribution [2], which is a function of the Courant-Snyder invariant. The K-V distribution is unstable over a range of beam and lattice parameters [3] and is generally assumed to be far from real distributions in accelerators. Furthermore, the K-V solution is only valid in two-dimensional focusing fields; it is unknown if there are solutions for (s -dependent) three-dimensional or nonlinear fields.

Lund [4] considered this problem for linear FODO channels with equal phase advances in the two planes of motion. The FODO lattice was mapped to a continuous-focusing (CF) lattice with the same average phase advance. Equilibrium distributions were generated in the CF lattice and then linearly transformed to match the periodic K-V envelope in the FODO lattice. This process generated surprisingly well-matched beam distributions. However, it is unclear how to adapt this method to (i) linearly coupled lattices, (ii) nonlinear lattices, or (iii) six-dimensional phase space distributions.

This work introduces a brute-force optimization to drive the distribution toward equilibrium in arbitrary external fields. The idea is to represent the distribution as a generative model—a neural network (NN) that transforms a simple base distribution to a more complicated distribution [5]—and to optimize the model parameters using gradient-based methods, enabled by differentiable particle-in-cell (PIC) simulations. The rest of this paper will discuss implementations of this approach and demonstrate convergence for a linear two-dimensional continuous-focusing lattice.

PROBLEM SETUP

Consider a periodic lattice of length L , such that $U(\mathbf{z}, s) = U(\mathbf{z}, s+L)$, and write the distribution function after n periods as $f_n(\mathbf{z})$, where $\mathbf{z} = (\mathbf{x}, \mathbf{x}')$ is the phase space vector. Let $\mathcal{M}f$ propagate the distribution function through one period: $f_{n+1} = \mathcal{M}f_n$. We seek a periodic distribution function $f_0 = f_1 = \mathcal{M}f_0$.

We also need to constrain the distribution emittance, since it is always possible to find a solution by scaling the emittance until space charge is negligible. The emittance is defined as $\epsilon = |\det|^{1/2}$, where $\Sigma = \langle \mathbf{z}\mathbf{z}^T \rangle$ is the covariance matrix. The algorithm should enforce $\epsilon = \tilde{\epsilon}$ for target emittance $\tilde{\epsilon}$.

We may also wish to pull the distribution toward some target \tilde{f} , either to stabilize training or to select one of multiple solutions; for example, to pull the solution away from the K-V equilibrium in a FODO channel. If there is only one




Figure 1: Example of an invertible generative model.

solution, the regularization will have no effect. (Note that there may not be *any* exact solutions.)

To satisfy these three objectives, we propose the following loss function to be minimized with respect to f , given the forward operator \mathcal{M} , target distribution \tilde{f} , and target emittance $\tilde{\epsilon}$:

$$L[f; \mathcal{M}, \tilde{f}, \tilde{\epsilon}] = \alpha_1 D_1[f, \mathcal{M}f] + \alpha_2 D_2[f, \tilde{f}] + |\epsilon - \tilde{\epsilon}|. \quad (4)$$

Here D_1 and D_2 are unspecified statistical distances or divergences. D_1 compares the initial distribution to the evolved distribution, while D_2 compares the initial distribution to the target distribution \tilde{f} . The constants $\alpha_{1,2}$ control the relative strengths of the three terms during optimization.

GENERATIVE MODEL

We represent the distribution using a generative model. The generative model is a function $\mathbf{T} : \mathbb{R}^D \rightarrow \mathbb{R}^D$ combined with a *base distribution* \hat{f} , which we take to be the unit Gaussian. The base distribution is defined in normalized coordinates $\hat{\mathbf{z}}$ which are connected to the phase space coordinates as $\mathbf{z} = T(\hat{\mathbf{z}})$. Samples from the model distribution $f(\mathbf{z})$ are obtained by pushing samples from the base distribution through the map T . (See Fig. 1.)

In addition to sampling particles, some generative models provide access to the underlying distribution function. Given a point $\hat{\mathbf{z}}$ sampled from $\hat{f}(\hat{\mathbf{z}})$ in the normalized space, the density at the corresponding point in phase space is:

$$f(\mathbf{z}) = \hat{f}(\hat{\mathbf{z}}) / |\det J_T(\hat{\mathbf{z}})|, \quad (5)$$

where $J_T(\hat{\mathbf{z}})$ is the Jacobian matrix of the transformation T at $\hat{\mathbf{z}}$. Thus, if we can calculate the Jacobian of T , we can evaluate the density at each sampled particle. Suppose instead that we want to evaluate the density $f(\mathbf{z})$ at an arbitrary \mathbf{z} . The density is given by

$$f(\mathbf{z}) = \hat{f}(\hat{\mathbf{z}}) |\det J_{T^{-1}}(\mathbf{z})|, \quad (6)$$

where $J_{T^{-1}}$ is the Jacobian matrix of the inverse transformation. Thus, we can only evaluate the density at arbitrary points if T is invertible.

We will implement T as a neural network (NN) with parameters θ . A standard NN architecture can be used if the density is not required. In low dimensions (2D, 4D, 6D), the

Jacobian of the NN transformation can be computed directly using AD to give the density at each randomly sampled point. In high dimensions, specially designed transformations and NN architectures must be used to simplify the calculation of the Jacobian. Calculating the density at *arbitrary* points is only possible with invertible NN architectures [6].

OPTIMIZATION

Return to the loss function in Eq. (4). The third term (the difference between the model and target emittance) can be estimated from a set of particles sampled from the model. The estimate is differentiable with respect to θ .

The second term $D_2[f, \tilde{f}]$ compares the model to the target distribution \tilde{f} . One option is the Kullback–Leibler (K-L) divergence:

$$D_2[f(\mathbf{z}), \tilde{f}(\mathbf{z})] = \int f(\mathbf{z}) \log[f(\mathbf{z})/\tilde{f}(\mathbf{z})] d\mathbf{z}. \quad (7)$$

The KL divergence can be estimated from N samples $\{\mathbf{z}_i\} \sim f(\mathbf{z})$:

$$D_2[f(\mathbf{z}), \tilde{f}(\mathbf{z})] \approx \frac{1}{N} \sum_{i=1}^N [\log f(\mathbf{z}_i) - \log \tilde{f}(\mathbf{z}_i)]. \quad (8)$$

The estimate in Eq. (8) requires the density $f(\mathbf{z}_i)$ at each sampled point, so the generative model does not need to be invertible. The estimate is differentiable with respect to θ .

The first term $D_1[f, \mathcal{M}f]$ compares the initial distribution to the evolved distribution. Label the initial distribution as f_0 and the evolved distribution as $f_1 = \mathcal{M}f_0$. We have access to samples $\{\mathbf{z}_i\} \sim f_0$, as well as the density at each sample. We might consider the KL divergence:

$$D_1[f_0(\mathbf{z}), f_1(\mathbf{z})] = \int f_0(\mathbf{z}) \log[f_0(\mathbf{z})/f_1(\mathbf{z})] d\mathbf{z}. \quad (9)$$

However, we do not know the evolved distribution function at each sampled initial point in phase space. What about the reverse KL divergence?

$$\begin{aligned} D_1[f_0(\mathbf{z}), f_1(\mathbf{z})] &= \int f_1(\mathbf{z}) \log[f_1(\mathbf{z})/f_0(\mathbf{z})] d\mathbf{z} \\ &= \int f_1(\mathbf{z}) \log f_1(\mathbf{z}) d\mathbf{z} - \int f_1(\mathbf{z}) \log f_0(\mathbf{z}) d\mathbf{z}. \end{aligned} \quad (10)$$

$$(11)$$

The first term in Eq. (10) is the entropy of the evolved distribution. The V-P system conserves entropy [7], so

$$\int f_1(\mathbf{z}) \log f_1(\mathbf{z}) d\mathbf{z} = \int f_0(\mathbf{z}) \log f_0(\mathbf{z}) d\mathbf{z}. \quad (12)$$

The entropy of the initial distribution can be estimated from the sampled particles [8]. The second term in Eq. (10) is the cross entropy; it can be evaluated from the *evolved* particles (samples from f_1): at each particle, we evaluate the model distribution function using Eq. (6). This is possible if the generative model is invertible. The estimate in Eq. (10) is differentiable with respect to θ .

A different approach is needed for non-invertible generative models. We propose a loss function that compares low-dimensional *projections* of the two distributions. If

the projections are 1D, we compute the average difference between projections along random directions in the high-dimensional phase space. In this work we use pairwise 2D projections: Define the 2D projection operator \mathcal{P}_{ij} :

$$\begin{aligned}\mathcal{P}_{ij}f(z_1, \dots, z_D) &= f(z_i, z_j) \\ &= \int f(z_1, \dots, z_D) \prod_{k \neq i, j} dz_k.\end{aligned}\quad (13)$$

The loss function then measures the average change in 2D projections over all $D(D-1)/2$ axes:

$$D_1[f, \mathcal{M}f] = \langle \mathcal{P}_{ij}f - \mathcal{P}_{ij}\mathcal{M}f \rangle \quad (14)$$

The expression in Eq. (14) is differentiable with respect to θ if the projections are calculated using kernel density estimation [5].

DEMONSTRATION IN CONTINUOUS-FOCUSING CHANNEL

As an initial demonstration, we applied the method to a two-dimensional linear continuous-focusing (CF) lattice. Particle tracking was performed in *Cheetah* [9], a beam dynamics code written in PyTorch. Space charge calculations used a standard integrated Green's function method to solve the Poisson equation on a 128×128 grid using 128,000 macro-particles [10].

We used a non-invertible generative model and the ADAM optimizer to minimize the loss function in Eq. (4), with regularization constants set to $\alpha_1 = 1$ and $\alpha_2 = 0$. Periodicity was enforced in the first term by the projection loss in Eq. (14). The second term was ignored because it had little effect on the solution. For the third term, rather than the 4D emittance ϵ , we constrained the 2D intrinsic emittances (also called eigenemittances) ϵ_1 and ϵ_2 . The target emittances were set to $\tilde{\epsilon}_1 = \tilde{\epsilon}_2 = 0.25$ [mm mrad] to eliminate linear cross-plane correlations in the beam.

We calculated the matched K-V envelope parameters for the specified beam intensity and energy (0.23 fractional tune depression). The model was first trained using reverse KL divergence to approach a Gaussian distribution with the same envelope parameters. Training then proceeded for 300 iterations to minimize the loss in Eq. (4). Fig. 2 plots the nonzero terms in the loss function during training, showing that a periodic solution is approached within the emittance constraints. Fig. 3 plots the period-by-period density fluctuations on the $x-x'$ axis before and after training (logarithmic scale).

CONCLUSION

We have introduced a method to drive simulated phase space distributions toward periodic equilibrium. The method uses a generative model to represent the distribution and a differentiable PIC simulation to optimize the model parameters. Loss functions have been identified to enforce periodicity in high-dimensional phase space and pull the solution toward a target value. We demonstrated convergence on a




Figure 2: Loss function during model training. The blue curve shows the average mismatch of two-dimensional projections. The orange curve shows the average distance from target emittances.




Figure 3: Projected phase space density fluctuations before and after training. A logarithmic density scale is used, cut off at 10^{-3} below peak density. The initial distribution in the top row is a Gaussian with covariance given by the matched K-V envelope.

two-dimensional linear continuous-focusing lattice at high space charge intensity.

Future work will continue examining CF lattices, including nonlinearities [11] and three-dimensional focusing, to benchmark against known solutions. Studies will then consider s -dependent lattices. Computer memory and gradient explosion are a concern for long accelerator lattices and strong space charge. If implementation issues are solved, generative models could provide new tools to study intensity limits in periodic accelerator lattices.

REFERENCES

- [1] C. E. Mitchell, K. Hwang, and R. D. Ryne, “Kurth Vlasov-Poisson Solution for a Beam in the Presence of Time-Dependent Isotropic Focusing”, in *Proc. IPAC’21*, Campinas, Brazil, May 2021, pp. 3213–3216. doi:10.18429/JACoW-IPAC2021-WEPAB248
- [2] I. Kapchinskij and V. Vladimirkij, “Limitations of proton beam current in a strong focusing linear accelerator associated with the beam space charge”, in *Proceedings of the International Conference on High Energy Accelerators and Instru-*

- mentation, p. 274, 1959. <https://s3.cern.ch/inspire-prod-files-5/526fe2f2ff6f05063bcb9a9b08da12f9>
- [3] I. Hofmann, “Stability of the Kapchinskij-Vladimirskij (K-V) distribution in long periodic transport systems”, *Part. Accel.*, vol. 13, pp. 145–178, 1983.
- [4] S. M. Lund, T. Kikuchi, and R. C. Davidson, “Generation of initial kinetic distributions for simulation of long-pulse charged particle beams with high space-charge intensity”, *Phys. Rev. Spec. Top. Accel. Beams*, vol. 12, no. 11, p. 114801, Nov. 2009. doi:10.1103/PhysRevSTAB.12.114801
- [5] R. Roussel *et al.*, “Efficient six-dimensional phase space reconstructions from experimental measurements using generative machine learning”, *Phys. Rev. Spec. Top. Accel. Beams*, vol. 27, no. 9, p. 094601, Sep. 2024. doi:10.1103/PhysRevAccelBeams.27.094601
- [6] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, “Normalizing Flows for Probabilistic Modeling and Inference”, *J. Mach. Learn. Res.*, vol. 22, no. 57, pp. 1–64, 2021.
- [7] V. Zhdankin, “Generalized Entropy Production in Collisionless Plasma Flows and Turbulence”, *Phys. Rev. X*, vol. 12, no. 3, p. 031011, Jul. 2022. doi:10.1103/PhysRevX.12.031011
- [8] A. Hoover and J. C. Wong, “High-dimensional maximum-entropy phase space tomography using normalizing flows”, *Phys. Rev. Res.*, vol. 6, no. 3, p. 033163, Aug. 2024. doi:10.1103/PhysRevResearch.6.033163
- [9] J. Kaiser, C. Xu, A. Eichler, and A. Santamaria Garcia, “Bridging the gap between machine learning and particle accelerator physics with high-speed, differentiable simulations”, *Phys. Rev. Accel. Beams*, vol. 27, no. 5, p. 054601, May 2024. doi:10.1103/PhysRevAccelBeams.27.054601
- [10] A. Huebl *et al.*, “Towards differentiable beam dynamics modeling in BLAST/ImpactX”, in *Proc. NAPAC’25*, Sacramento, California, USA, pp. 614–618, Jan. 2026. doi:10.18429/JACoW-NAPAC2025-TUP101
- [11] C. E. Mitchell, R. D. Ryne, and K. Hwang, “Spectral Galerkin solver for intense beam Vlasov equilibria in nonlinear constant focusing channels”, *Phys. Rev. E*, vol. 100, no. 5, p. 053308, Nov. 2019. doi:10.1103/PhysRevE.100.053308