

# A NEURAL SCORE-BASED METHOD FOR DETERMINISTIC COLLISIONAL PLASMA SIMULATION

Vasily Ilin<sup>1</sup> & Jingwei Hu<sup>2</sup>

<sup>1</sup>Department of Mathematics, <sup>2</sup>Department of Applied Mathematics  
University of Washington  
Seattle, WA 98195, USA  
{vilin, hujw}@uw.edu

## ABSTRACT

Plasma modeling is central to the design of nuclear fusion reactors, yet simulating collisional plasma kinetics from first principles remains a formidable computational challenge: the Vlasov-Maxwell-Landau (VML) equations couple six-dimensional phase-space transport to self-consistent electromagnetic fields through the nonlinear, nonlocal Landau collision operator. The sole deterministic particle method for the full VML system estimates the velocity score function via the blob method – a kernel-based approximation that incurs  $O(n^2)$  cost, does not scale well in dimension, and has unsatisfactory performance in low-density regimes. We replace the blob score estimator with score-based transport modeling (SBTM), in which a neural network is trained on-the-fly via implicit score matching at  $O(n)$  cost. We prove that the collision operator preserves mass, momentum, and kinetic energy for *any* score approximation, and dissipates an estimated entropy. On three canonical benchmarks – Landau damping, two-stream instability, and Weibel instability – SBTM is more accurate than the blob method, achieves correct long-time thermalization to Maxwellian equilibrium where the blob method fails, and delivers up to  $1.6\times$  faster runtime with  $2 - 4\times$  lower peak memory.

## 1 INTRODUCTION

Controlled nuclear fusion promises a virtually limitless source of clean energy, but achieving it requires confining a plasma – an ionized gas of electrons and ions – at temperatures exceeding  $10^8$  K, where charged particles undergo frequent Coulomb collisions while generating and responding to electromagnetic fields (Chen, 2016). Accurate numerical simulation of such plasmas is essential for reactor design and remains one of the grand challenges in computational physics. The Vlasov-Maxwell-Landau (VML) system provides the first-principles kinetic description (Landau, 1936; Villani, 2002), coupling six-dimensional phase-space transport to self-consistent electromagnetic fields through the Landau collision operator. Solving the VML equations is exceptionally challenging due to the high dimensionality, multiscale dynamics, and the need to preserve conservation laws and entropy structure.

Particle-in-cell (PIC) methods (Hockney & Eastwood, 1988; Birdsall & Langdon, 2018) are the dominant approach for high-dimensional kinetic equations: particles are advanced in the Lagrangian frame while fields are solved on a spatial grid. However, incorporating the Landau collision operator into PIC is nontrivial. Classical stochastic approaches (Takizuka & Abe, 1977; Bird, 1994; Manheimer et al., 1997) introduce statistical noise and do not preserve the conservation laws of the underlying equation.

A deterministic alternative is possible because the Landau collision operator can be rewritten as a *transport equation* driven by the velocity score function  $s(x, v) = \nabla_v \log f(x, v)$  (Villani, 1998), reducing the problem to estimating the score from particle data. Carrillo et al. (2020) introduced the first deterministic particle method for the homogeneous Landau equation, estimating the score via a kernel density estimate (the “blob method”). Bailo et al. (2024) extended this to the full VML system. However, the blob method computes pairwise kernel sums over all  $n$  particles in each

spatial cell, resulting in  $O(n^2)$  cost and memory, and – critically – produces inaccurate scores in low-density tail regions, causing systematic artifacts in long-time thermalization.

Score-based transport modeling (SBTM) (Boffi & Vanden-Eijnden, 2023) offers an alternative: a neural network  $s_\theta$  is trained on-the-fly to approximate the score by minimizing an implicit score matching loss (Hyvärinen, 2005), at  $O(n)$  cost per gradient step. Ilin et al. (2025a) applied SBTM to the spatially homogeneous Landau equation, demonstrating superior accuracy especially in sparse particle regimes, and Ilin (2025) established error bounds for SBTM with Coulomb collisions.

**Contribution.** We apply SBTM to the full Vlasov-Maxwell-Landau system, extending Ilin et al. (2025a) to the spatially inhomogeneous case with electromagnetic field coupling. Building on the PIC framework of Bailo et al. (2024), we replace *only* the score estimation module – requiring spatial localization via the hat kernel  $\psi_\eta$  and handling distinct velocity distributions across spatial cells with a single network. On three canonical benchmarks:

- SBTM achieves correct long-time *thermalization* to Maxwellian equilibrium, where the blob method fails due to inaccurate scores in distribution tails.
- SBTM is *stable across particle counts*: the blob method requires much larger  $n$  to match the SBTM solution.
- SBTM is *faster* ( $1.3 - 1.6\times$ ) and uses *less memory* ( $2 - 4\times$ ).
- At the semi-discrete level (continuous time), the collision operator preserves *mass, momentum, and kinetic energy* exactly for any score approximation, and dissipates an estimated entropy.

## 2 RELATED WORK

**Deterministic particle methods for collisional plasmas.** The blob method for the Landau equation was introduced by Carrillo et al. (2020), with convergence analysis in Carrillo et al. (2023) and gradient flow structure in Carrillo et al. (2024); Huang & Wang (2024a). Bailo et al. (2024) extended it to the full VML system. Random batch methods reduce cost to  $O(n^2/R)$  (Carrillo et al., 2022); Yoo et al. (2025) developed an energy-conserving integrator for Vlasov-Fokker-Planck. Our work retains the PIC framework of Bailo et al. (2024) but replaces the  $O(n^2)$  blob score estimator with a learned one.

**Score-based methods for kinetic equations.** SBTM (Boffi & Vanden-Eijnden, 2023) trains a neural network to approximate  $\nabla \log f$  on-the-fly via score matching (Hyvärinen, 2005) and has been applied to Fokker-Planck equations (Shen et al., 2022; Shen & Wang, 2023; Lu et al., 2024), density sampling (Ilin et al., 2025b), and the homogeneous Landau equation (Ilin et al., 2025a; Huang & Wang, 2024b), with error bounds in Ilin (2025). Ilin & Sushko (2026) proposed DiScoFormer, a pretrained Transformer for score estimation. All prior SBTM work on the Landau equation is restricted to the *spatially homogeneous* setting; ours is the first application to the full inhomogeneous VML system.

**Other neural approaches.** Neural operator methods (Li et al., 2021; Lu et al., 2021; Azizzadehsheli et al., 2024) learn solution maps but do not inherently preserve conservation laws. Miller et al. (2021) used a neural network to approximate the collision operator in the XGC gyrokinetic code. Our method uses the neural network only for score estimation within the structure-preserving PIC framework, retaining exact conservation (Theorem 2).

## 3 BACKGROUND AND METHOD

### 3.1 THE VLASOV-MAXWELL-LANDAU SYSTEM

The Vlasov-Maxwell-Landau system describes the evolution of a collisional, magnetized plasma. For a single electron species with an immobile neutralizing ion background of charge density  $\rho_{\text{ion}}$ ,

the distribution function  $f(t, x, v)$  with  $v \in \mathbb{R}^{d_v}$  evolves under:

$$\partial_t f + v \cdot \nabla_x f + (E + v \times B) \cdot \nabla_v f = \nu Q[f, f], \quad (1a)$$

$$Q[f, f](x, v) = \nabla_v \cdot \int_{\mathbb{R}^{d_v}} A(v-w) [f(x, w) \nabla_v f(x, v) - f(x, v) \nabla_w f(x, w)] dw, \quad (1b)$$

$$\partial_t E = \nabla \times B - J, \quad \partial_t B = -\nabla \times E, \quad \nabla \cdot E = \rho - \rho_{\text{ion}}, \quad \nabla \cdot B = 0, \quad (1c)$$

where  $E(t, x)$  and  $B(t, x)$  are the electric and magnetic fields,  $\nu \geq 0$  is the collision frequency,  $Q[f, f]$  is the Landau collision operator (Landau, 1936),  $J = \int v f dv$  is the current density, and  $\rho = \int f dv$  is the electron charge density. The collision kernel for Coulomb interactions is  $A(z) = |z|^{2-d_v} \Pi(z)$  with  $\Pi(z) = I_{d_v} - z z^\top / |z|^2$ . The collision operator can be rewritten in terms of the velocity score  $\nabla_v \log f$  (Villani, 1998; Carrillo et al., 2020):

$$Q[f, f](x, v) = \nabla_v \cdot \left\{ f(x, v) \int_{\mathbb{R}^{d_v}} A(v-w) f(x, w) [\nabla_v \log f(x, v) - \nabla_w \log f(x, w)] dw \right\}. \quad (2)$$

This form is the basis for the particle method: the collision operator requires only the score function, not the density itself.

**1D3V reduction.** We work in the 1D-in-space,  $d_v$ -dimensional-in-velocity setting on a periodic spatial domain  $\Omega = [0, L]$ . In this geometry,  $E = (E_1(x, t), E_2(x, t), 0)$  and  $B = (0, 0, B_3(x, t))$ , and Maxwell's equations reduce to:

$$\partial_t E_1 = -J_1, \quad \partial_t E_2 = -\partial_x B_3 - J_2, \quad \partial_t B_3 = -\partial_x E_2. \quad (3)$$

The constraint  $\partial_x E_1 = \rho - \rho_{\text{ion}}$  is enforced at initialization and preserved by the update.

#### EQUILIBRIUM CHARACTERIZATION

The conservation laws and entropy structure uniquely determine the long-time equilibrium, providing ground truth for our thermalization experiments.

**Theorem 1** (Global steady states). *Consider the VML system equation 1 with  $\nu > 0$ , periodic  $x \in \mathbb{T}^3$ ,  $v \in \mathbb{R}^3$ , and uniform neutralizing background  $\rho_{\text{ion}} > 0$ . Every smooth steady state  $f > 0$  with finite energy and entropy is a spatially uniform global Maxwellian  $f_\infty(v) = \rho_{\text{ion}} (2\pi T_\infty)^{-3/2} \exp(-|v - u_\infty|^2 / (2T_\infty))$  with  $E_\infty = 0$ .*

(i) **VML** (with magnetic field): Ampère's law forces  $u_\infty = 0$ , and  $B_\infty = \frac{1}{|\mathbb{T}^3|} \int B_{\text{init}} dx$  is the conserved spatial mean. The temperature is  $T_\infty = \frac{2}{3\rho_{\text{ion}} |\mathbb{T}^3|} (\mathcal{E}_0 - \frac{1}{2} |B_\infty|^2 |\mathbb{T}^3|)$ .

(ii) **VPL** (electrostatic,  $B = 0$ ): The system is Galilean invariant, permitting nonzero drift  $u_\infty = \iint v f_{\text{init}} dv dx / (\rho_{\text{ion}} |\mathbb{T}^3|)$ , with  $T_\infty = \frac{2}{3\rho_{\text{ion}} |\mathbb{T}^3|} (\mathcal{E}_0 - \frac{1}{2} \rho_{\text{ion}} |u_\infty|^2 |\mathbb{T}^3|)$ .

Here  $\mathcal{E}_0$  is the initial total energy. Proofs are given in Appendix A.

### 3.2 PARTICLE METHOD AND COLLISION OPERATOR

Following Bailo et al. (2024), we represent the distribution as  $n$  weighted particles:

$$f_n(t, x, v) \approx \sum_{p=1}^n w_p \delta(x - x_p(t)) \delta(v - v_p(t)). \quad (4)$$

The particle positions and velocities evolve according to the coupled ODE system:

$$\frac{dx_p}{dt} = v_{1,p} \pmod{L}, \quad (5)$$

$$\frac{dv_p}{dt} = (E(x_p) + v_p \times B(x_p)) - \nu \sum_{q=1}^n w_q \psi_\eta(x_p - x_q) A(v_p - v_q) [s(x_p, v_p) - s(x_q, v_q)], \quad (6)$$

where  $s(x, v) = \nabla_v \log f(x, v)$  is the velocity score function,  $\psi_\eta$  is a hat kernel for spatial localization with support  $\eta = L/M$  ( $M$  is the number of grid cells), and the first term is the Lorentz force.

The spatial hat kernel restricts collisions to particles within the same grid cell, reducing the per-step collision cost from  $O(n^2)$  to  $O(n^2/M)$ . The score is the key quantity that must be estimated from particle data; the fields  $E$  and  $B$  are computed via the standard PIC field solve on the spatial grid.

**Structural properties.** The particle system equation 5–equation 6 inherits the conservation laws of the continuous VML equations. These are most naturally seen through the *weak formulation*: for any test function  $\psi(x, v)$ ,

$$\frac{d}{dt} \sum_p w_p \psi(x_p, v_p) = \underbrace{\sum_p w_p \nabla_x \psi \cdot v_p}_{C_x} + \underbrace{\sum_p w_p \nabla_v \psi \cdot (E + v_p \times B)}_{C_v} + \underbrace{\sum_p w_p \nabla_v \psi \cdot F_p^{\text{coll}}}_{D}, \quad (7)$$

where  $F_p^{\text{coll}} = -\nu \sum_q w_q \psi_\eta(x_p - x_q) A(v_p - v_q) [s_p - s_q]$  is the collision force.

**Theorem 2** (Conservation and entropy properties). *Consider the particle system equation 5–equation 6 in continuous time with  $\psi_\eta$  symmetric. For any score approximation  $s$ :*

- (i) **Mass:**  $\frac{d}{dt} \sum_p w_p = 0$ .
- (ii) **Momentum:**  $\frac{d}{dt} \sum_p w_p v_p = \sum_p w_p (E + v_p \times B)$ . *The collision operator does not change the total momentum.*
- (iii) **Total energy:**  $\frac{d}{dt} \left[ \sum_p \frac{1}{2} w_p |v_p|^2 + \frac{1}{2} \int (|E|^2 + |B|^2) dx \right] = 0$  *at the semi-discrete level.*
- (iv) **Estimated entropy dissipation:** *the estimated entropy production rate  $\sum_p w_p s_p \cdot F_p^{\text{coll}}$  is non-positive for any  $s$ .*

These hold for *any* score approximation (blob or SBTM) at the semi-discrete level; in the homogeneous case, even under forward Euler (Ilin et al., 2025a, Proposition 3.1). In the fully discrete VML setting, conservation incurs  $O(\Delta t)$  drift; empirically, SBTM exhibits better total energy conservation than blob across all benchmarks. An energy-conserving integrator is developed in Yoo et al. (2025). See Appendix B for proofs (following Bailo et al. (2024)) and Ilin (2025) for a posteriori error bounds.

### 3.3 SCORE ESTIMATION

**Blob method.** The blob method (Carrillo et al., 2019; 2020; Bailo et al., 2024) estimates the score via a regularized kernel density estimate. In all experiments, we use only the leading KDE term (the  $O(\varepsilon)$  entropy-dissipation correction of Bailo et al. (2024) has negligible effect on trajectories):

$$s_{\text{KDE}}(x, v) = \nabla_v \log \left( \sum_q K_h(v - v_q) \psi_\eta(x - x_q) \right), \quad (8)$$

where  $K_h$  is a Gaussian kernel with bandwidth  $h$  via Silverman’s rule. This is twice as fast as computing the full blob score, while yielding empirically identical trajectories. Computing equation 8 requires  $O(n^2)$  pairwise evaluations and  $O(n^2)$  memory within each spatial cell.

**SBTM.** Score-based transport modeling (Boffi & Vanden-Eijnden, 2023; Ilin et al., 2025a) trains a neural network  $s_\theta: \mathbb{R}^{1+d_v} \rightarrow \mathbb{R}^{d_v}$  to approximate the score by minimizing the implicit score matching (ISM) loss (Hyvärinen, 2005):

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{p=1}^n \left[ |s_\theta(x_p, v_p)|^2 + 2 \nabla_v \cdot s_\theta(x_p, v_p) \right]. \quad (9)$$

The divergence  $\nabla_v \cdot s_\theta$  is estimated via Hutchinson’s trace estimator (Hutchinson, 1989) with Rademacher random vectors:

$$\nabla_v \cdot s_\theta(x, v) \approx z^\top (\nabla_v s_\theta(x, v)) z, \quad z \sim \text{Rademacher}(\pm 1)^{d_v}, \quad (10)$$

where  $(\nabla_v s_\theta) z$  is a single forward-mode autodiff pass. Each particle draws an independent  $z$ ; the ISM loss averages over  $n$  particles, giving  $O(1/\sqrt{n})$  relative error ( $\sim 0.1\%$  at  $n = 10^6$ ). Hutchinson is  $\sim 3\times$  faster than the exact trace ( $d_v$  passes) at  $d_v = 3$ .

**Algorithm 1** SBTM-PIC for the Vlasov-Maxwell-Landau equations

---

**Require:** Particles  $(x_p, v_p)_{p=1}^n$  with weights  $w_p$ , fields  $E, B$ , time step  $\Delta t$ , final time  $t_{\text{final}}$ , collision frequency  $\nu$ , ISM steps  $K$ , spatial kernel  $\psi_\eta$ , collision kernel  $A$

- 1: Pretrain  $s_\theta$  on the analytic score of the initial condition
- 2:  $t \leftarrow 0$
- 3: **while**  $t < t_{\text{final}}$  **do**
- 4:   Interpolate fields:  $E_p \leftarrow E(x_p), B_p \leftarrow B(x_p)$
- 5:    $v_p \leftarrow v_p + \Delta t (E_p + v_p \times B_p)$  *(Lorentz push)*
- 6:    $x_p \leftarrow x_p + \Delta t v_{1,p} \pmod{L}$  *(position advance)*
- 7:   Deposit charge/current onto grid; update  $E, B$  via Maxwell’s equations
- 8:   **for**  $k = 1, \dots, K$  **do**
- 9:     Draw  $z \sim \text{Rademacher}(\pm 1)^{d_v}$
- 10:      $\theta \leftarrow \theta - \alpha \nabla_\theta \left[ \frac{1}{n} \sum_p (|s_\theta(x_p, v_p)|^2 + 2 z^\top (\nabla_v s_\theta(x_p, v_p)) z) \right]$  *(ISM)*
- 11:   **end for**
- 12:    $v_p \leftarrow v_p - \Delta t \nu \sum_q w_q \psi_\eta(x_p - x_q) A(v_p - v_q) [s_\theta(x_p, v_p) - s_\theta(x_q, v_q)]$  *(collision)*
- 13:    $t \leftarrow t + \Delta t$
- 14: **end while**

---

**Architecture and training.** The score network  $s_\theta: \mathbb{R}^{1+d_v} \rightarrow \mathbb{R}^{d_v}$  is a two-hidden-layer MLP with softsign activation:  $(1 + d_v) \rightarrow H \rightarrow H \rightarrow d_v$ , with  $H = 256$  ( $\approx 68\text{K}$  parameters) or  $512$  ( $\approx 267\text{K}$ ) for Weibel. At  $t = 0$ , the network is pretrained on the known analytic score. At each subsequent time step,  $K = 100$  AdamW gradient steps ( $\text{lr} = 2 \times 10^{-4}$ ) on  $\mathcal{L}(\theta)$  are performed; the choice of  $K$  is not sensitive in the range  $20 \leq K \leq 400$ .

**Algorithm.** The full SBTM-PIC method is summarized in Algorithm 1.

The blob method uses exactly the same algorithm, but replaces the ISM training (lines 9 – 11) with the KDE computation equation 8.

## 4 NUMERICAL EXPERIMENTS

### 4.1 SETUP

All experiments use the 1D- $d_v$ V setting ( $d_v \in \{2, 3\}$ ) with periodic boundary conditions. The spatial domain  $[0, L]$  is discretized into  $M$  grid cells with spacing  $\eta = L/M$ . Time integration uses forward Euler. The Maxwell solver uses centered second-order finite differences. The code is implemented in JAX (Bradbury et al., 2018) with double precision and will be made publicly available upon publication. We compare two score estimation methods – blob (KDE) and SBTM – using the same PIC framework. The *only* difference between the two methods is how the score  $\nabla_v \log f$  is estimated.

### 4.2 TWO-STREAM INSTABILITY

**Setup.** Two counter-streaming Maxwellian beams with a small spatial perturbation:

$$f_0(x, v) = \frac{1 + \alpha \cos(kx)}{2L} [\mathcal{N}(v_1; c, 1) + \mathcal{N}(v_1; -c, 1)] \prod_{j \geq 2} \mathcal{N}(v_j; 0, 1), \quad (11)$$

with beam speed  $c = 2.4$ , wavenumber  $k = 1/5$ , perturbation  $\alpha = 1/200$ , domain  $L = 10\pi$ ,  $M = 100$ ,  $\Delta t = 0.05$ ,  $d_v = 3$ ,  $\nu = 0.24$ , final time  $T = 50$ , and  $K = 100$  ISM steps per time step. We test  $n \in \{5 \times 10^5, 10^6, 3 \times 10^6\}$ .

**Phase space dynamics.** Figure 1 compares  $(x, v_1)$  phase space snapshots. The vortex should dissipate by  $t = 50$  at  $\nu = 0.24$ . SBTM achieves complete dissipation at  $n = 10^6$ , unchanged at  $n = 3 \times 10^6$ . The blob method retains residual vortex structure at  $n = 10^6$  and requires  $n = 3 \times 10^6$  to approach the SBTM solution.

**Estimated entropy production and total energy.** Figure 2 shows estimated entropy production and total energy. SBTM dissipates estimated entropy consistently at all particle counts (curves

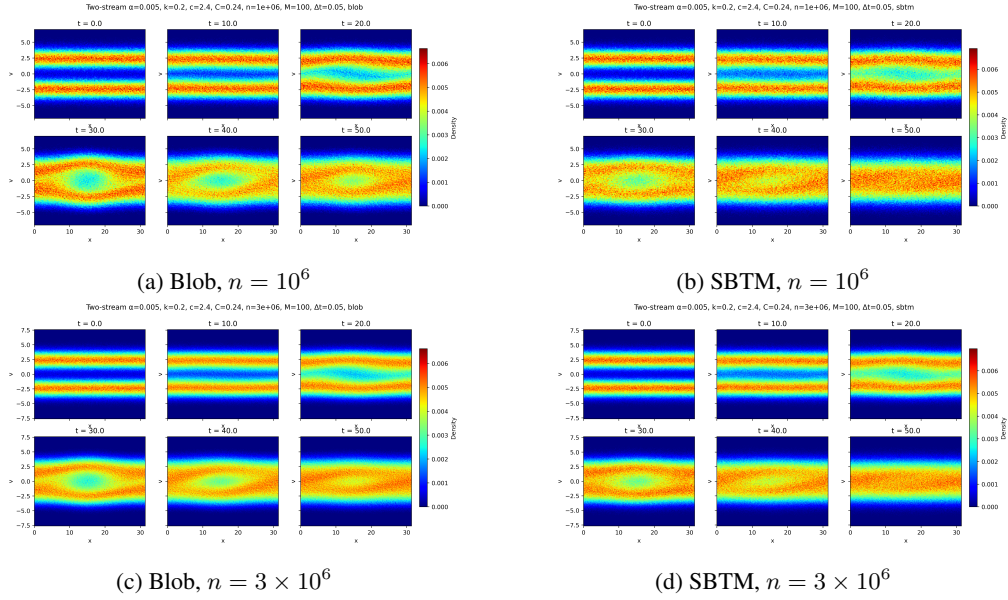


Figure 1: Two-stream instability:  $(x, v_1)$  phase space at  $\nu = 0.24$ . Top:  $n = 10^6$ ; bottom:  $n = 3 \times 10^6$ . At  $n = 10^6$ , SBTM achieves full vortex dissipation while the blob method retains residual structure. At  $n = 3 \times 10^6$ , the blob solution approaches SBTM. Additional particle counts in Appendix D.

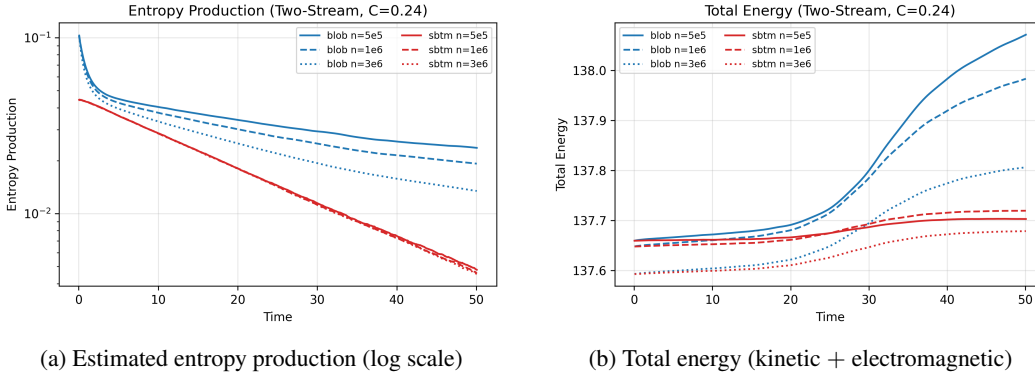


Figure 2: Two-stream instability at  $\nu = 0.24$ : estimated entropy production and total energy across particle counts. SBTM (red) dissipates estimated entropy consistently and maintains near-constant total energy. The blob method (blue) overestimates estimated entropy production and exhibits upward energy drift.

nearly coincide), while the blob method systematically overestimates. SBTM also maintains near-constant total energy, while the blob method exhibits upward drift – noisy score estimates break the cancellation guaranteed by Theorem 2(iii).

**Score estimation quality.** Figure 3 shows score quiver plots at  $n = 10^6$ . At  $t = 0$ , where the true score is known analytically, SBTM achieves low MSE while the blob method has substantially higher MSE. At  $t = T$ , SBTM scores are smooth and physically consistent, while the blob method produces noisy scores in low-density tail regions.

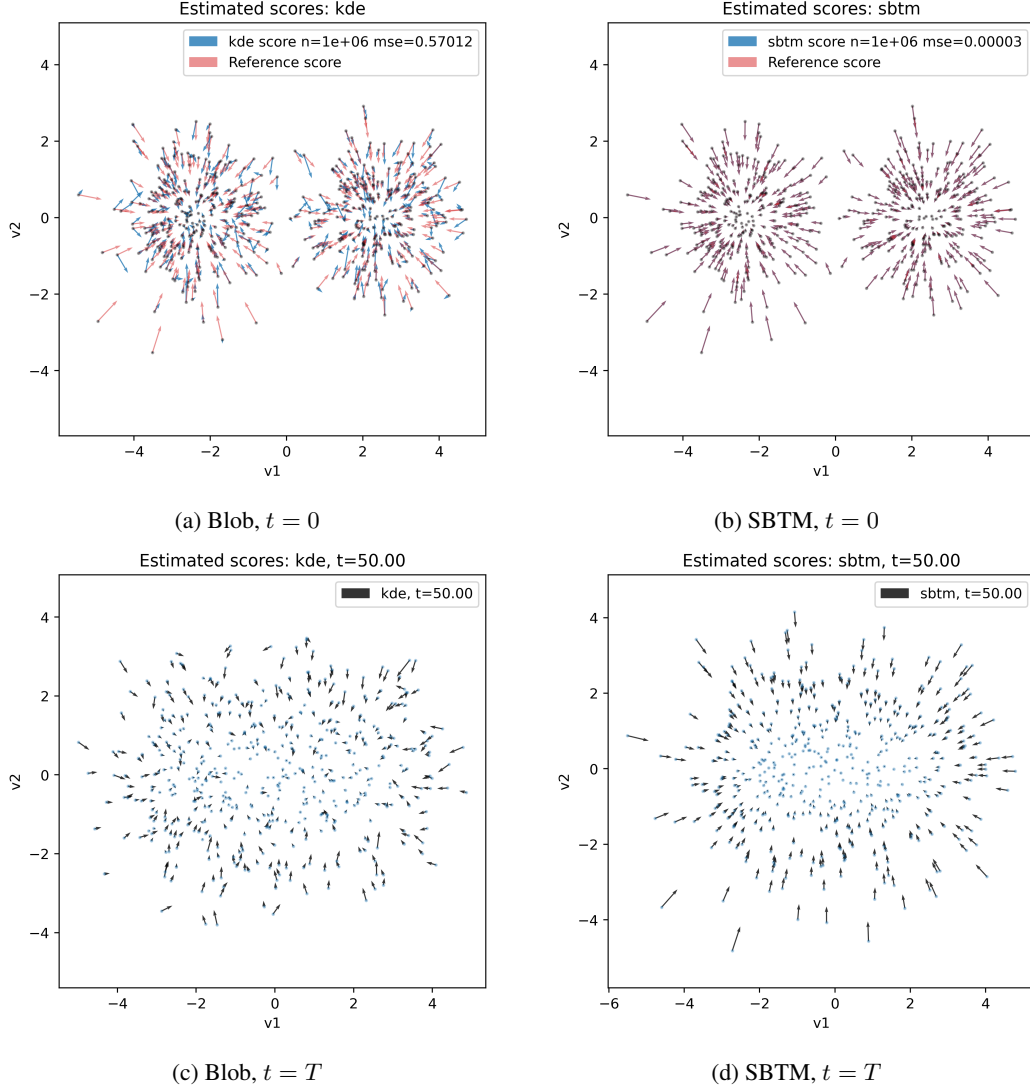


Figure 3: Two-stream instability: score quiver plots at  $\nu = 0.24$ ,  $n = 10^6$ . Top: initial time (true score in red; MSE in legend). Bottom: final time. SBTM produces smooth, physically consistent scores; the blob method has high MSE, particularly in low-density tail regions.

### 4.3 WEIBEL INSTABILITY

**Setup.** An anisotropic bi-Maxwellian with a small magnetic perturbation:

$$f_0(x, v) \propto \exp\left(-\frac{v_1^2}{\beta}\right) \left[ \exp\left(-\frac{(v_2 - c)^2}{\beta}\right) + \exp\left(-\frac{(v_2 + c)^2}{\beta}\right) \right] \prod_{j \geq 3} \exp\left(-\frac{v_j^2}{\beta}\right), \quad (12)$$

with  $c = 0.3$ ,  $\beta = 0.01$ ,  $B_3(0, x) = \alpha_B \sin(kx)$ ,  $\alpha_B = 10^{-3}$ ,  $k = 1/5$ ,  $M = 100$ ,  $\Delta t = 0.1$ ,  $d_v = 3$ ,  $n = 10^6$ ,  $T = 125$ , and  $K = 100$  ISM steps per time step. We sweep  $\nu \in \{10^{-4}, 2 \times 10^{-4}, 4 \times 10^{-4}, 8 \times 10^{-4}\}$ .

**Velocity distribution and thermalization.** By Theorem 1(i), the collision operator should drive the bimodal distribution toward a single Maxwellian. Figure 4 shows the  $v_2$ -marginal at  $\nu = 8 \times 10^{-4}$ . SBTM converges to the expected Maxwellian with smooth Gaussian tails, while the blob method produces sharp, unphysical cutoffs at  $v_2 \approx \pm 0.55$  (linear-scale plots in Appendix E).




Figure 4: Weibel instability:  $v_2$ -marginal density evolution (log scale) at  $\nu = 8 \times 10^{-4}$ . The dashed black curve is a Gaussian/Maxwellian. SBTM converges to it with smooth Gaussian tails, while the blob method exhibits sharp, unphysical cutoffs at  $v_2 \approx \pm 0.55$ .

**Phase space and scores.** Phase space snapshots (Appendix E) confirm that the blob method’s advantage over SBTM diminishes as collisions strengthen: at  $\nu = 8 \times 10^{-4}$ , blob produces unrealistic non-Gaussian densities while SBTM maintains proper Gaussian shape. Score quiver plots (Appendix E) show SBTM produces smooth scores with low MSE, while blob is noisy in the tails.

**Entropy and thermalization.** At equilibrium, entropy production must vanish. SBTM correctly decays toward zero at all  $\nu$ , while blob plateaus at nonzero values (Appendix E). Table 1 quantifies thermalization via the  $L^2$  distance of the  $v_2$ -marginal to a Gaussian at  $t = 125$ . SBTM is 1.2–3.5 $\times$  closer to the Maxwellian predicted by Theorem 1(i).

Table 1:  $L^2$  distance of  $v_2$ -marginal to Gaussian at  $t = 125$  ( $n = 10^6$ ,  $d_v = 3$ ).

$\nu$	Blob	SBTM	Ratio
$2 \times 10^{-4}$	0.325	0.265	1.2 $\times$
$4 \times 10^{-4}$	0.134	0.053	2.5 $\times$
$8 \times 10^{-4}$	0.070	0.020	3.4 $\times$

#### 4.4 LANDAU DAMPING

**Setup.** A small spatial perturbation on a Maxwellian equilibrium:

$$f_0(x, v) = \frac{1 + \alpha \cos(kx)}{L} \cdot \frac{1}{(2\pi)^{d_v/2}} \exp\left(-\frac{|v|^2}{2}\right), \quad (13)$$

with wavenumber  $k = 0.5$  and domain  $L = 2\pi/k$ . We use  $\alpha = 0.1$ ,  $d_v = 3$ ,  $M = 100$ ,  $\Delta t = 0.02$ ,  $\nu = 0.4$ ,  $T = 15$ , and  $K = 100$ , sweeping  $n \in \{5 \times 10^5, 10^6, 3 \times 10^6\}$ .

**Electric field decay.** A small perturbation decays exponentially. At  $\nu = 0.4$  – high enough that linear theory is no longer accurate – the blob method’s damping rate varies with  $n$ , while SBTM produces a consistent rate across all three particle counts. Table 2 reports fitted rates. Linear theory predicts  $\gamma = -0.151$  (collisionless) and  $\gamma = -0.258$  (collisional), though these are not expected to be accurate at  $\nu = 0.4$ . The SBTM rate is approximately constant across  $n$  and falls between the two predictions. The blob rate increases with  $n$ , approaching SBTM only at  $n = 3 \times 10^6$ .

Table 2: Fitted damping rates ( $\nu = 0.4$ ,  $\alpha = 0.1$ ,  $T = 15$ ). Linear theory:  $-0.151$  (collisionless),  $-0.258$  (collisional).

$n$	Blob	SBTM
$5 \times 10^5$	-0.121	-0.208
$10^6$	-0.133	-0.211
$3 \times 10^6$	-0.156	-0.212

SBTM also shows superior estimated entropy production and total energy conservation (Appendix C). Figure 5 shows velocity-space distributions at  $\nu = 1.0$ : SBTM produces smooth Gaussian tails, while blob exhibits a sharp cutoff at  $|v_2| \approx 3$ . Additional collision frequencies in Appendix C.

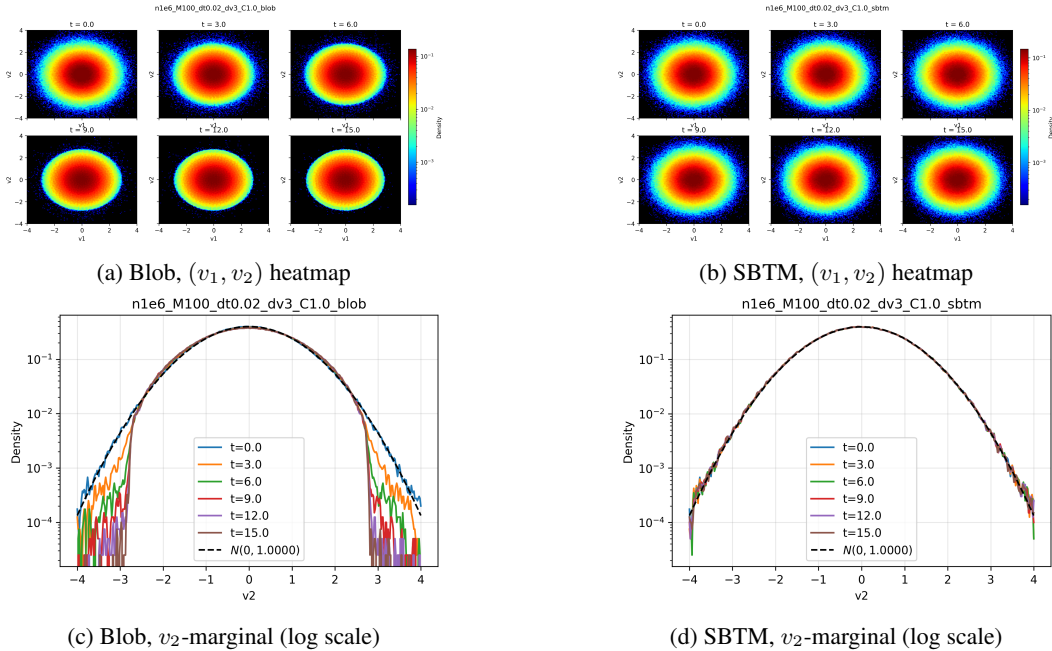


Figure 5: Landau damping: velocity-space distributions at  $\nu = 1.0$ ,  $n = 10^6$ . Top:  $(v_1, v_2)$  heatmaps. Bottom:  $v_2$ -marginal density on a log scale. The blob method’s  $v_2$ -marginal drops abruptly at  $|v_2| \approx 3$ , while SBTM maintains smooth Gaussian tails.

Table 3: Runtime and peak GPU memory comparison ( $M = 100$ ,  $d_v = 3$ ). All runs on L40S except \*blob at  $n = 3 \times 10^6$  (H200, 141 GB required).

Benchmark	Method	Runtime ( $n = 10^6$ )		Peak Memory (GB)	
		Time	Speedup	$n = 10^6$	$n = 3 \times 10^6$
Two-stream ( $T=50, \Delta t=0.05$ )	Blob	11h 27m	–	16.4	65.1*
	SBTM	7h 17m	$1.57\times$	4.4	17.2
Weibel ( $T=125, \Delta t=0.1$ )	Blob	3h 26m	–	8.6	65.1*
	SBTM	2h 37m	$1.32\times$	8.7	33.1
Landau damping ( $T=15, \Delta t=0.02$ )	Blob	44m	–	8.9	65.2*
	SBTM	35m	$1.27\times$	4.9	17.1

#### 4.5 COMPUTATIONAL EFFICIENCY

Table 3 compares runtime and peak GPU memory. SBTM is  $1.3 - 1.6\times$  faster at  $n = 10^6$ . At  $n = 3 \times 10^6$ , blob requires 65 GB ( $O(n^2)$  within each cell) vs. 17 – 33 GB for SBTM. All runs on L40S except blob at  $n = 3 \times 10^6$  (H200, 141 GB).

### 5 CONCLUSION

SBTM provides a superior alternative to the blob method for deterministic particle simulation of the Vlasov-Maxwell-Landau equations: correct thermalization to Maxwellian equilibrium, entropy production decaying to zero, stable results across particle counts, and  $1.3 - 1.6\times$  faster runtime with  $2 - 4\times$  lower memory. Since SBTM modifies only the collision operator (acting in velocity space), it extends to higher spatial dimensions without modification.

## REFERENCES

- Kamyar Azzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 2024.
- Rafael Bailo, José A. Carrillo, and Jingwei Hu. The collisional particle-in-cell method for the Vlasov-Maxwell-Landau equations. *arXiv preprint arXiv:2401.01689*, 2024.
- Graeme A Bird. *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*. Clarendon Press, Oxford, 1994.
- Charles K Birdsall and A Bruce Langdon. *Plasma Physics via Computer Simulation*. CRC Press, 2018.
- Nicholas M Boffi and Eric Vanden-Eijnden. Probability flow solution of the Fokker-Planck equation. *Machine Learning: Science and Technology*, 4(3):035012, 2023.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, et al. Jax: composable transformations of python+ numpy programs. 2018.
- José A Carrillo, Jingwei Hu, Li Wang, and Jeremy Wu. A particle method for the homogeneous Landau equation. *Journal of Computational Physics: X*, 7:100066, 2020.
- José A Carrillo, Matias G Delgadino, and Jeremy S.H. Wu. Convergence of a particle method for a regularized spatially homogeneous Landau equation. *Mathematical Models and Methods in Applied Sciences*, 33(05):971–1008, 2023.
- José A Carrillo, Matias G Delgadino, Laurent Desvillettes, and Jeremy Wu. The Landau equation as a gradient flow. *Analysis & PDE (to appear)*, 2024.
- José Antonio Carrillo, Katy Craig, and Francesco S Patacchini. A blob method for diffusion. *Calculus of Variations and Partial Differential Equations*, 58:1–53, 2019.
- José Antonio Carrillo, Shi Jin, and Yijia Tang. Random batch particle methods for the homogeneous Landau equation. *Communications in Computational Physics*, 31(4):997–1019, 2022.
- Carlo Cercignani. *The Boltzmann Equation and Its Applications*. Springer, 1988.
- Francis F Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer International Publishing, 2016.
- Roger W Hockney and James W Eastwood. *Computer Simulation Using Particles*. CRC Press, 1988.
- Yan Huang and Li Wang. JKO for Landau: a variational particle method for homogeneous Landau equation. *arXiv preprint arXiv:2409.12296*, 2024a.
- Yan Huang and Li Wang. A score-based particle method for homogeneous Landau equation. *arXiv preprint arXiv:2405.05187*, 2024b.
- Michael F Hutchinson. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics—Simulation and Computation*, 18(3):1059–1076, 1989.
- Aapo Hyvärinen. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005.
- Vasily Ilin. Stability of the spatially homogeneous landau equation in relative entropy and applications to score-based numerical methods. *arXiv preprint arXiv:2510.15089*, 2025.
- Vasily Ilin and Peter Sushko. Discoformer: Plug-in density and score estimation with transformers. *arXiv preprint arXiv:2511.05924*, 2026.

- Vasily Ilin, Jingwei Hu, and Zhenfu Wang. Transport based particle methods for the Fokker-Planck-Landau equation. *Communications in Mathematical Sciences*, 23(7):1763–1788, 2025a.
- Vasily Ilin, Peter Sushko, and Jingwei Hu. Score-based deterministic density sampling. *arXiv preprint arXiv:2504.18130*, 2025b.
- Lev D Landau. Die kinetische Gleichung für den Fall Coulombscher Wechselwirkung. *Physikalische Zeitschrift der Sowjetunion*, 10(2):154–164, 1936.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.
- Jianfeng Lu, Yue Wu, and Yang Xiang. Score-based transport modeling for mean-field Fokker-Planck equations. *Journal of Computational Physics*, 503:112859, 2024.
- Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, 2021.
- Wallace M Manheimer, Martin Lampe, and Glenn Joyce. Langevin representation of Coulomb collisions in PIC simulations. *Journal of Computational Physics*, 138(2):563–584, 1997.
- M. A. Miller, R. M. Churchill, A. Dener, C. S. Chang, T. Munson, and R. Hager. Encoder-decoder neural network for solving the nonlinear Fokker-Planck-Landau collision operator in XGC. *Journal of Plasma Physics*, 87(2), 2021.
- Zebang Shen and Zhenfu Wang. Entropy-dissipation informed neural network for McKean-Vlasov type PDEs. In *37th Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Zebang Shen, Zhenfu Wang, Satyen Kale, Alejandro Ribeiro, Amin Karbasi, and Hamed Hassani. Self-consistency of the Fokker-Planck equation. In *Conference on Learning Theory*, pp. 817–841. PMLR, 2022.
- Tomonor Takizuka and Hirotsada Abe. A binary collision model for plasma simulation with a particle code. *Journal of Computational Physics*, 25(3):205–219, 1977.
- Cédric Villani. On the spatially homogeneous Landau equation for Maxwellian molecules. *Mathematical Models and Methods in Applied Sciences*, 8(06):957–983, 1998.
- Cédric Villani. A review of mathematical topics in collisional kinetic theory. In S Friedlander and D Serre (eds.), *Handbook of Mathematical Fluid Mechanics*, volume I, pp. 71–305. North-Holland, 2002.
- Jiyoung Yoo, Jingwei Hu, and Lee F Ricketson. An explicit energy-conserving particle method for the vlasov-fokker-planck equation. *arXiv preprint arXiv:2510.03960*, 2025.

## A EQUILIBRIUM CHARACTERIZATION

We provide proofs of Theorem 1. The proof structure was developed with the assistance of Gemini DeepThink.

*Proof of Theorem 1(i) – VML case.* The proof proceeds in two stages: (1) the Landau entropy dissipation forces  $f$  to be a local Maxwellian in  $v$ , and (2) the Vlasov–Maxwell transport equations constrain the spatial dependence.

**Step 1: Entropy dissipation and the local Maxwellian.** At equilibrium,  $\partial_t f = \partial_t E = \partial_t B = 0$ , so equation 1a reduces to  $v \cdot \nabla_x f + (E + v \times B) \cdot \nabla_v f = \nu Q[f, f]$ . Multiply by  $\log f$  and integrate over  $\mathbb{T}^3 \times \mathbb{R}^3$ . The spatial transport term vanishes by the divergence theorem on the periodic domain. The Lorentz force term vanishes upon integration by parts since  $\nabla_v \cdot (E + v \times B) = 0$ . This leaves:

$$\int_{\mathbb{T}^3} \left( \int_{\mathbb{R}^3} Q[f, f] \log f \, dv \right) dx = 0. \quad (14)$$

By the H-theorem for the Landau operator,  $\int Q[f, f] \log f \, dv \leq 0$  pointwise in  $x$ . Since its spatial integral is zero, it must vanish identically. The null space of the Landau dissipation is spanned by  $\{1, v, |v|^2\}$  (Villani, 2002; Cercignani, 1988), so  $f$  is a local Maxwellian:

$$f(x, v) = \frac{n(x)}{(2\pi T(x))^{3/2}} \exp\left(-\frac{|v - u(x)|^2}{2T(x)}\right), \quad (15)$$

with density  $n(x) > 0$ , bulk velocity  $u(x)$ , and temperature  $T(x) > 0$ .

**Step 2: Polynomial matching.** Since  $Q[f, f] = 0$ , the Vlasov equation becomes  $v \cdot \nabla_x(\log f) + (E + v \times B) \cdot \nabla_v(\log f) = 0$ . Writing  $\log f = a(x) + b(x) \cdot v + c(x)|v|^2$  with  $c = -1/(2T)$  and  $b = u/T$ , and expanding yields a polynomial in  $v$  that must vanish identically.

(a)  $\mathcal{O}(|v|^3)$ :  $\nabla_x c = 0$ , so  $T(x) \equiv T_\infty$  is constant (Cercignani, 1988, Eq. 10.7).

(b)  $\mathcal{O}(|v|^2)$ : Killing's equation  $\partial_{x_i} b_j + \partial_{x_j} b_i = 0$ ; on  $\mathbb{T}^3$  this forces  $b$  constant, hence  $u(x) \equiv u_\infty$  (Cercignani, 1988, Eq. 10.6).

(c)  $\mathcal{O}(|v|^1)$ : Force balance:  $\nabla_x \log n = (E + u_\infty \times B)/T_\infty$ .

**Step 3: Ampère's law forces  $u_\infty = 0$ .** The current density is  $J = n(x)u_\infty$ . Ampère's law at steady state gives  $\nabla_x \times B = n(x)u_\infty$ . Taking the dot product with  $u_\infty$  and using  $\nabla_x \cdot (B \times u_\infty) = u_\infty \cdot (\nabla_x \times B)$  (since  $\nabla_x \times u_\infty = 0$ ), integration over  $\mathbb{T}^3$  gives  $0 = |u_\infty|^2 \int n(x) \, dx$ . Since  $n > 0$ , we must have  $u_\infty = 0$ .

**Step 4: Density and field uniformity.** With  $u_\infty = 0$ , the force balance gives  $\nabla_x \log n = E/T_\infty$ . Taking the divergence and using Gauss's law  $\nabla_x \cdot E = n - \rho_{\text{ion}}$ :  $T_\infty \Delta_x(\log n) = n - \rho_{\text{ion}}$ . By the maximum principle on the periodic domain: at a maximum of  $n$ ,  $\Delta_x \log n \leq 0$  forces  $n \leq \rho_{\text{ion}}$ ; at a minimum,  $n \geq \rho_{\text{ion}}$ . Hence  $n \equiv \rho_{\text{ion}}$  and  $E = 0$ .

**Step 5: Magnetic field.** With  $J = 0$ , Ampère's law gives  $\nabla_x \times B = 0$ . Combined with  $\nabla_x \cdot B = 0$ , we get  $\Delta_x B = 0$  on  $\mathbb{T}^3$ , so  $B \equiv B_\infty$ . Conservation of  $\int_{\mathbb{T}^3} B \, dx$  (from Faraday's law on a periodic domain) fixes  $B_\infty = \frac{1}{|\mathbb{T}^3|} \int B_{\text{init}} \, dx$ . The temperature  $T_\infty$  is then determined by total energy conservation.  $\square$

*Proof of Theorem 1(ii) – VPL case.* Steps 1–2 of the VML proof carry over with  $B = 0$ : the steady state is a local Maxwellian with constant  $T_\infty$  and uniform  $u_\infty$ , satisfying  $\nabla_x \log n = E/T_\infty$ .

**Step 3 breaks:** without Ampère's law,  $u_\infty = 0$  is not forced. However, the electrostatic balance (Step 4) still gives  $n \equiv \rho_{\text{ion}}$  and  $E = 0$  by the same maximum principle argument.

**Determination of  $u_\infty$ .** The total kinetic momentum  $P = \iint v f \, dv \, dx$  is conserved: the collision term preserves momentum locally, the transport term integrates to zero on  $\mathbb{T}^3$ , and the electric self-force  $\int nE \, dx$  vanishes because  $E = -\nabla_x \phi$  is a gradient (so  $\int \rho_{\text{ion}} E \, dx = 0$ ) and  $\int (n - \rho_{\text{ion}})E \, dx = \int (\nabla \cdot E)E \, dx = 0$  by the stress tensor identity  $(\nabla \cdot E)E = \nabla \cdot (E \otimes E - \frac{1}{2}|E|^2 I)$  for curl-free  $E$ . Thus  $u_\infty = P(0)/(\rho_{\text{ion}}|\mathbb{T}^3|)$ , and  $T_\infty$  follows from energy conservation.  $\square$

## B CONSERVATION PROPERTIES

We provide detailed proofs of Theorem 2. The approach follows Bailo et al. (2024), unified through the weak formulation equation 7.

**Notation.** Define the collision force on particle  $p$ :

$$F_p^{\text{coll}} = -\nu \sum_{q=1}^n w_q \psi_\eta(x_p - x_q) A(v_p - v_q) [s_p - s_q], \quad (16)$$

where  $A(z) = |z|^{2-d_v} \Pi(z)$  with  $\Pi(z) = I_{d_v} - zz^\top / |z|^2$  the orthogonal projection onto  $z^\perp$ ,  $s_p = s(x_p, v_p)$  is the score approximation at particle  $p$ , and  $A_{pq} = A(v_p - v_q)$ . The key symmetries are:  $A(-z) = A(z)$  (so  $A_{pq} = A_{qp}$ ),  $\psi_\eta(-z) = \psi_\eta(z)$ , and  $s_q - s_p = -(s_p - s_q)$ . The diagonal terms  $p = q$  vanish since  $s_p - s_p = 0$ .

The full particle ODE equation 6 gives, for any test function  $\psi(x, v)$ :

$$\begin{aligned} \frac{d}{dt} \sum_p w_p \psi(x_p, v_p) &= \underbrace{\sum_p w_p \nabla_x \psi(x_p, v_p) \cdot v_p}_{C_x} + \underbrace{\sum_p w_p \nabla_v \psi(x_p, v_p) \cdot (E(x_p) + v_p \times B(x_p))}_{C_v} \\ &+ \underbrace{\sum_p w_p \nabla_v \psi(x_p, v_p) \cdot F_p^{\text{coll}}}_{D}. \end{aligned} \quad (17)$$

**Discrete weak form of the collision term.** The collision contribution  $D$  can be symmetrized by swapping  $p \leftrightarrow q$  in the double sum and averaging:

$$D = -\frac{\nu}{2} \sum_{p,q} w_p w_q \psi_\eta(x_p - x_q) [\nabla_v \psi_p - \nabla_v \psi_q]^\top A_{pq} (s_p - s_q), \quad (18)$$

where we write  $\nabla_v \psi_p = \nabla_v \psi(x_p, v_p)$  for brevity. To verify: start from  $D = -\nu \sum_{p,q} w_p w_q \psi_\eta(x_p - x_q) (\nabla_v \psi_p)^\top A_{pq} (s_p - s_q)$ . Swap  $p \leftrightarrow q$  and use  $A_{qp} = A_{pq}$ ,  $\psi_\eta(x_q - x_p) = \psi_\eta(x_p - x_q)$ ,  $s_q - s_p = -(s_p - s_q)$  to get  $-\nu \sum_{p,q} w_p w_q \psi_\eta(x_p - x_q) (-\nabla_v \psi_q)^\top A_{pq} (s_p - s_q)$ . Averaging gives equation 18.

We now apply the weak formulation equation 17 with different choices of  $\psi$ .

*Proof of Theorem 2. (i) Mass* ( $\psi = 1$ ). All three terms in equation 17 vanish since  $\nabla_x \psi = 0$  and  $\nabla_v \psi = 0$ . Thus  $\frac{d}{dt} \sum_p w_p = 0$ .

**(ii) Momentum** ( $\psi = v_i$ , the  $i$ -th velocity component). We have  $\nabla_v \psi = e_i$  (the  $i$ -th standard basis vector) and  $\nabla_x \psi = 0$ .

$C_x = 0$  since  $\nabla_x \psi = 0$ .

$C_v = \sum_p w_p (E_i(x_p) + (v_p \times B(x_p))_i)$ , which is the Lorentz force contribution.

$D$ : In equation 18,  $\nabla_v \psi_p - \nabla_v \psi_q = e_i - e_i = 0$ . So  $D = 0$ .

Therefore  $\frac{d}{dt} \sum_p w_p v_{i,p} = \sum_p w_p (E_i(x_p) + (v_p \times B(x_p))_i)$ , i.e., the collision operator conserves momentum exactly for any score  $s$ , and the total momentum changes only due to the Lorentz force.

**(iii) Total energy.** We first show that the collision operator conserves kinetic energy. Set  $\psi = \frac{1}{2}|v|^2$ , so  $\nabla_v \psi = v$  and  $\nabla_x \psi = 0$ .

$C_x = 0$  since  $\nabla_x \psi = 0$ .

$C_v = \sum_p w_p v_p \cdot (E(x_p) + v_p \times B(x_p)) = \sum_p w_p v_p \cdot E(x_p)$ , since  $v_p \cdot (v_p \times B) = 0$  (the magnetic force does no work).

$D$ : In equation 18,  $\nabla_v \psi_p - \nabla_v \psi_q = v_p - v_q$ . So

$$D = -\frac{\nu}{2} \sum_{p,q} w_p w_q \psi_\eta(x_p - x_q) (v_p - v_q)^\top A(v_p - v_q) (s_p - s_q).$$

We use the key algebraic identity: for any  $z \in \mathbb{R}^{d_v}$ ,

$$z^\top A(z) = |z|^{2-d_v} z^\top \Pi(z) = |z|^{2-d_v} z^\top \left( I_{d_v} - \frac{zz^\top}{|z|^2} \right) = |z|^{2-d_v} (z^\top - z^\top) = 0^\top. \quad (19)$$

Applying this with  $z = v_p - v_q$ , every term vanishes:  $(v_p - v_q)^\top A(v_p - v_q) = 0^\top$ . Hence  $D = 0$  for *any* score approximation  $s$ .

Therefore  $\frac{d}{dt} \sum_p \frac{1}{2} w_p |v_p|^2 = \sum_p w_p v_p \cdot E(x_p)$ , with the collision operator contributing nothing. The kinetic energy changes at rate  $\sum_p w_p v_p \cdot E(x_p) \approx \int J \cdot E dx$ . From Maxwell's equations  $\partial_t E = \nabla \times B - J$  and  $\partial_t B = -\nabla \times E$ , we compute the electromagnetic energy rate:

$$\begin{aligned} \frac{d}{dt} \frac{1}{2} \int (|E|^2 + |B|^2) dx &= \int (E \cdot \partial_t E + B \cdot \partial_t B) dx \\ &= \int (E \cdot (\nabla \times B - J) + B \cdot (-\nabla \times E)) dx. \end{aligned}$$

Using the vector identity  $\nabla \cdot (E \times B) = B \cdot (\nabla \times E) - E \cdot (\nabla \times B)$  and periodic boundary conditions (which eliminate the divergence term), this simplifies to  $-\int J \cdot E dx$ . The two contributions cancel:

$$\frac{d}{dt} \left[ \sum_p \frac{1}{2} w_p |v_p|^2 + \frac{1}{2} \int (|E|^2 + |B|^2) dx \right] \approx \int J \cdot E dx - \int J \cdot E dx = 0.$$

The approximation arises from replacing  $\sum_p w_p v_p \cdot E(x_p)$  by  $\int J \cdot E dx$ ; at the semi-discrete level with continuous fields, this is exact. Under forward Euler time discretization, the cancellation is approximate.

**(iv) Estimated entropy dissipation.** We motivate the definition of the estimated entropy production rate. For any test function  $\psi(x, v)$ , the weak formulation equation 17 gives

$$\frac{d}{dt} \sum_p w_p \psi(x_p, v_p) = C_x + C_v + D,$$

where  $D = \sum_p w_p \nabla_v \psi_p \cdot F_p^{\text{coll}}$  is the collision contribution. When  $\psi = \log f$  (so  $\nabla_v \psi = \nabla_v \log f$ ), the terms  $C_x$  and  $C_v$  vanish:  $C_v = 0$  because the Lorentz force is divergence-free in  $v$  ( $\nabla_v \cdot (E + v \times B) = 0$ ), so  $\int \nabla_v \log f \cdot (E + v \times B) f dv = 0$  by integration by parts; similarly  $C_x = 0$  by integration by parts in  $x$ . Thus, only the collision term  $D$  contributes to entropy dissipation.

We define the *estimated entropy production rate* by replacing  $\nabla_v \log f$  with the score approximation  $s$  in  $D$ :

$$\text{estimated entropy production rate} = \sum_p w_p s_p \cdot F_p^{\text{coll}}.$$

This is exactly the collision contribution  $D$  evaluated with  $\nabla_v \psi_p = s_p$ . Using equation 18 with  $\nabla_v \psi_p - \nabla_v \psi_q = s_p - s_q$ :

$$\sum_p w_p s_p \cdot F_p^{\text{coll}} = -\frac{\nu}{2} \sum_{p,q} w_p w_q \psi_\eta(x_p - x_q) (s_p - s_q)^\top A_{pq} (s_p - s_q) \leq 0,$$

since  $A_{pq}$  is positive semidefinite (as  $\Pi(z)$  is an orthogonal projection and  $|z|^{2-d_v} \geq 0$ ). This holds for *any* score approximation  $s$  – the collision operator always dissipates the estimated entropy. This is the same argument as in Ilin et al. (2025a, Proposition 3.1).  $\square$

**Homogeneous case.** In the spatially homogeneous setting (no  $x$ -dependence, no electromagnetic fields), the transport and Lorentz terms  $C_x$  and  $C_v$  are absent, and the estimated entropy dissipation  $D \leq 0$  holds exactly for any score approximation, including under forward Euler time discretization. This is stated precisely in Ilin et al. (2025a, Proposition 3.1).

## C LANDAU DAMPING

**Electric field decay.** Figure 6 shows the  $L^2$  norm of the electric field over time for blob and SBTM at  $n = 10^6$ , with  $\alpha = 0.1$ ,  $d_v = 3$ ,  $\Delta t = 0.02$ ,  $\nu = 0.4$ ,  $T = 15$ . Each plot shows the computed decay rate alongside the linear theory damping rate (plotted for reference; at  $\nu = 0.4$  the collision frequency is high enough that linear theory is not expected to be accurate). The SBTM rate remains constant as  $n$  varies from  $5 \times 10^5$  to  $3 \times 10^6$ , while the blob method’s rate changes with  $n$ , approaching the SBTM rate at higher  $n$ .

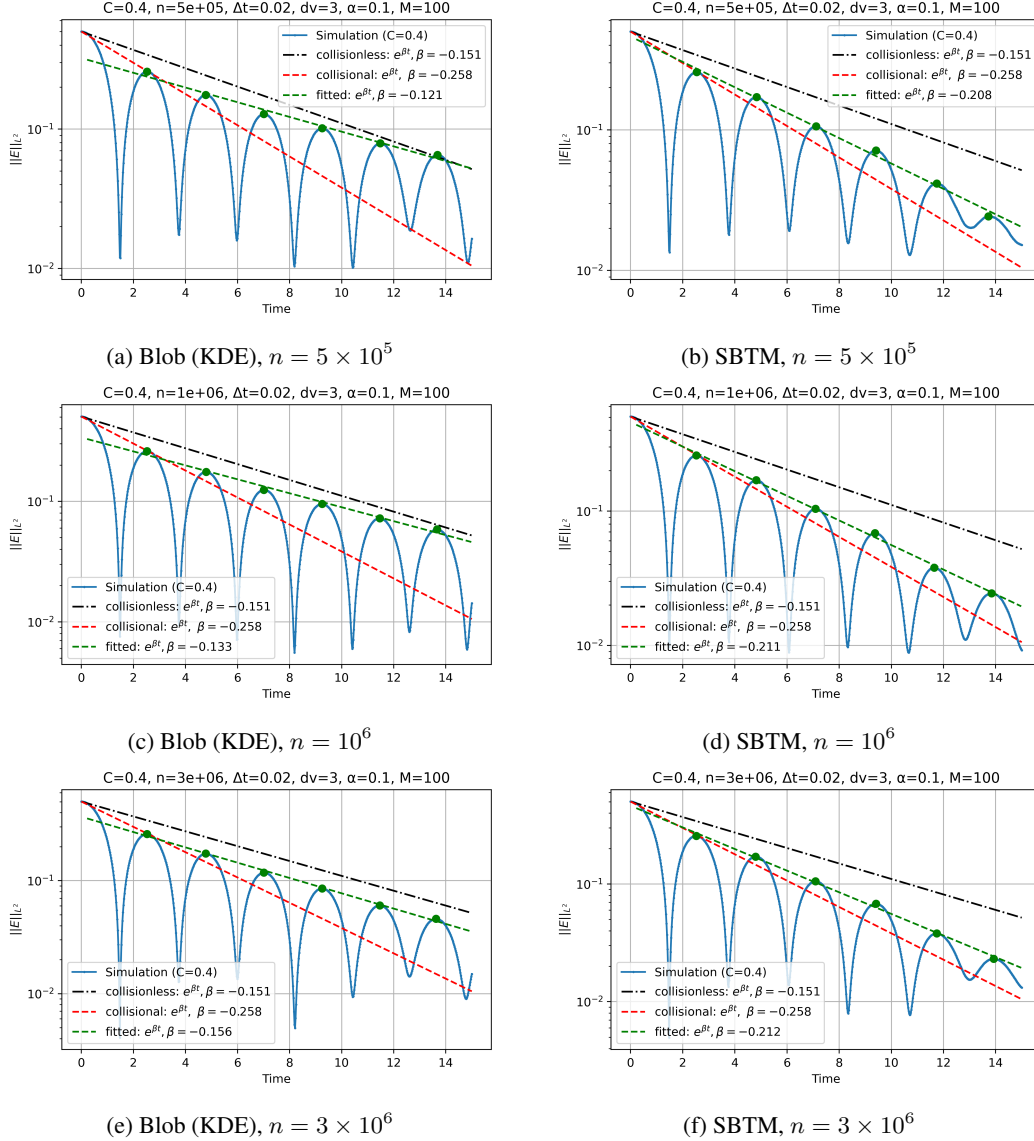


Figure 6: Landau damping:  $L^2$  norm of the electric field over time at  $\nu = 0.4$ . The linear theory damping rate is shown for reference (at  $\nu = 0.4$  it is not expected to be accurate). The blob method’s damping rate varies with  $n$  (compare (a), (c), and (e)), converging toward the SBTM rate. The SBTM rate is consistent across all particle counts (b), (d), (f).

**Estimated entropy production and total energy.** Figure 7 shows estimated entropy production and total energy conservation for the Landau damping benchmark. SBTM produces consistent estimated entropy production curves across all particle counts, while the blob method overestimates, especially

at lower  $n$ . SBTM also exhibits better total energy conservation, with the blob method showing larger drift.

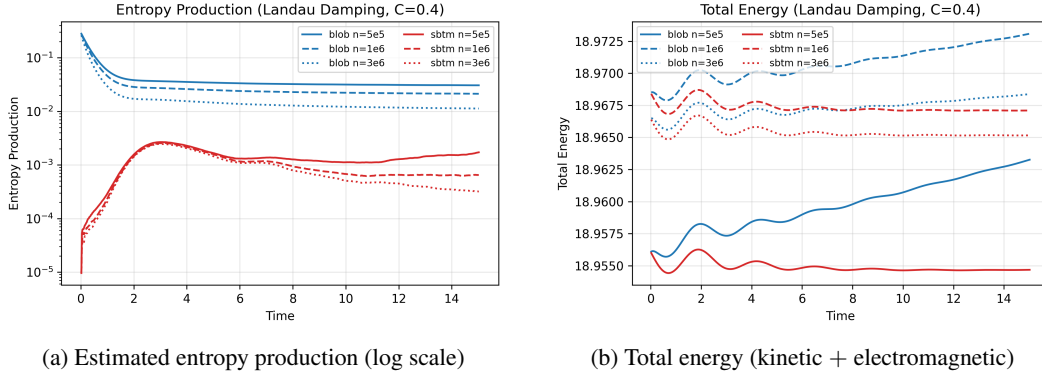


Figure 7: Landau damping at  $\nu = 0.4$ : estimated entropy production and total energy across particle counts. SBTM (red) dissipates estimated entropy consistently and maintains near-constant total energy. The blob method (blue) overestimates estimated entropy production and exhibits energy drift.

**Velocity-space phase space.** Figures 8–10 show  $(v_1, v_2)$  phase space heatmaps and  $v_2$ -marginal densities at  $n = 10^6$  for  $\nu \in \{0.1, 0.4, 1.0\}$ . SBTM produces smooth Gaussian tails across all collision frequencies, while the blob method exhibits sharp cutoffs in low-density regions.

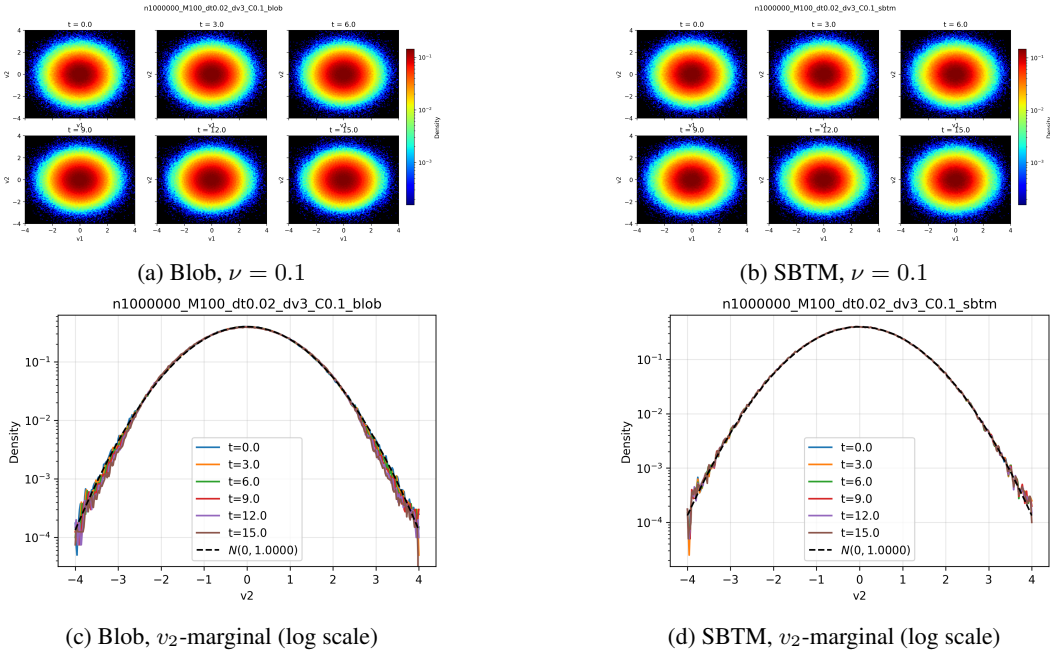


Figure 8: Landau damping:  $(v_1, v_2)$  phase space and  $v_2$ -marginal density at  $\nu = 0.1$ ,  $n = 10^6$ . At this low collision frequency, collisional effects are weak and both methods produce similar heatmaps. However, the  $v_2$ -marginal (log scale) reveals that SBTM maintains smooth Gaussian tails while the blob method already exhibits cutoff artifacts.

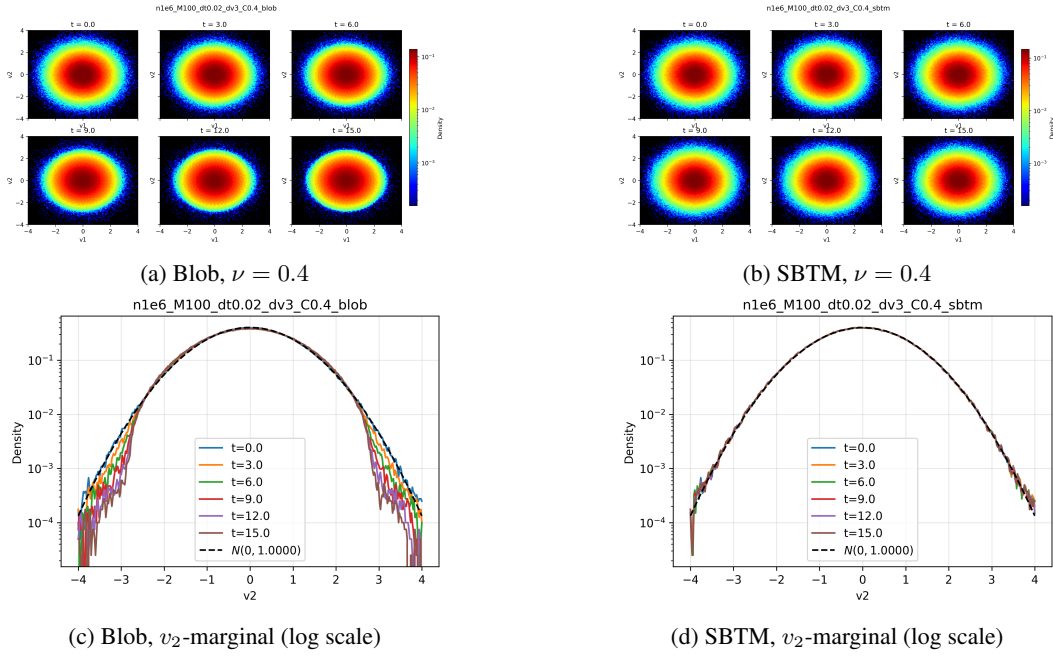


Figure 9: Landau damping:  $(v_1, v_2)$  phase space and  $v_2$ -marginal density at  $\nu = 0.4$ ,  $n = 10^6$ . SBTM maintains smooth Gaussian tails, while the blob method shows sharp cutoffs in low-density tail regions.

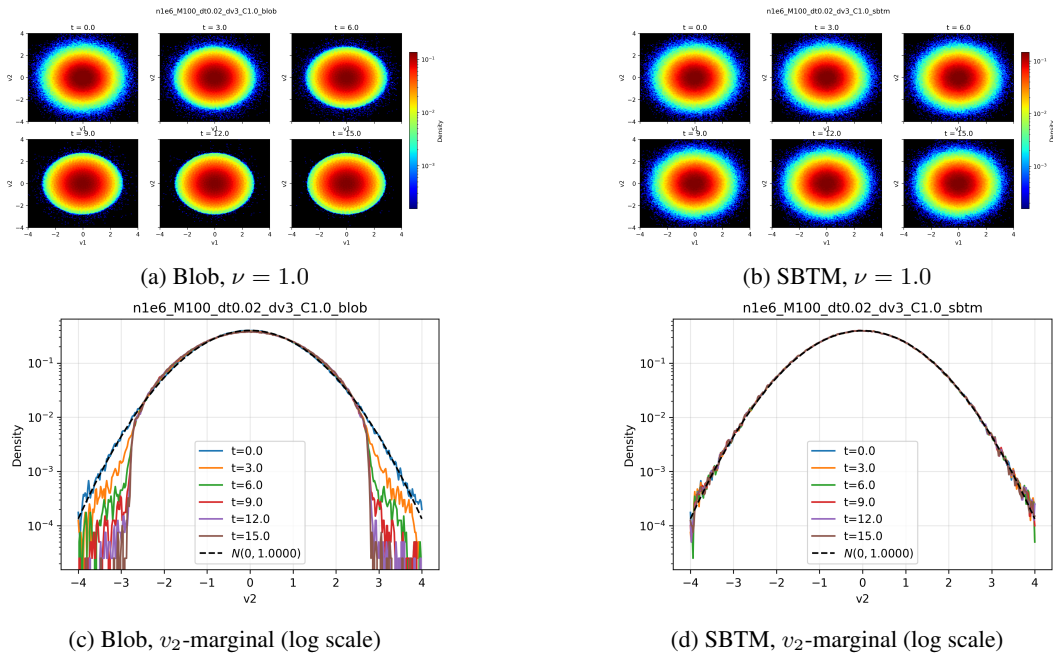


Figure 10: Landau damping:  $(v_1, v_2)$  phase space and  $v_2$ -marginal density at  $\nu = 1.0$ ,  $n = 10^6$ . At this high collision frequency, collisions dominate and both methods rapidly thermalize. The blob method's sharp tail cutoff is most pronounced here: the  $v_2$ -marginal drops abruptly at  $|v_2| \approx 3$ , while SBTM maintains smooth Gaussian tails throughout.

## D ADDITIONAL TWO-STREAM INSTABILITY RESULTS

**Phase space ( $\nu = 0.32$ ).** Figure 11 shows phase space at  $\nu = 0.32$  across three particle counts.

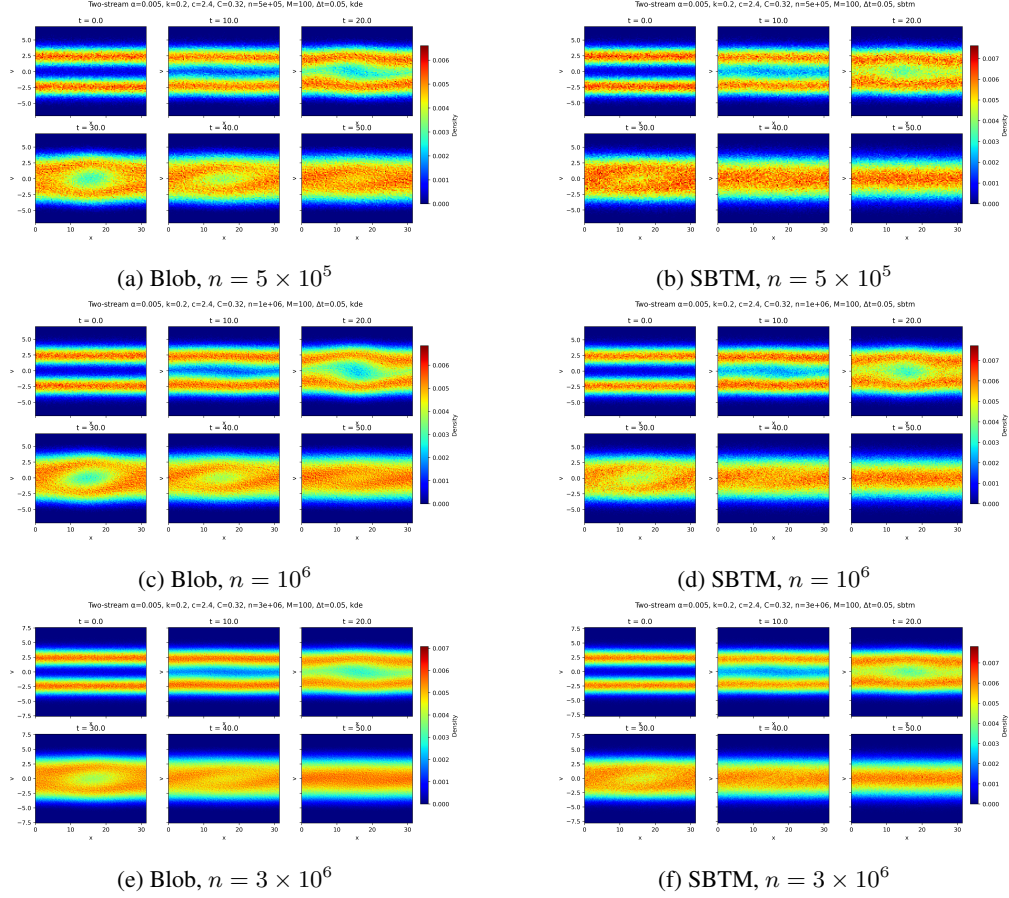


Figure 11: Two-stream instability:  $(x, v_1)$  phase space at  $\nu = 0.32$  across particle counts. At this higher collision frequency, SBTM achieves full vortex dissipation at all particle counts, while the blob method requires  $n = 3 \times 10^6$  to approach the same level of dissipation. Comparing with  $\nu = 0.24$  (Figure 1), stronger collisions accelerate thermalization for both methods, but SBTM remains more accurate at lower  $n$ .

**Energy and estimated entropy diagnostics ( $\nu = 0.32$ ).** Figure 12 shows electric field energy, kinetic energy, total energy, and estimated entropy production at  $\nu = 0.32$ . SBTM produces consistent curves across all  $n$ , while the blob method shows significant variation with  $n$ . The estimated entropy plot confirms the same pattern as  $\nu = 0.24$  (Figure 2): SBTM dissipates estimated entropy consistently while the blob method overestimates.

**Energy diagnostics ( $\nu = 0.24$ ).** Figure 13 shows electric field energy and kinetic energy at  $\nu = 0.24$ , complementing the estimated entropy production and total energy shown in the main text (Figure 2).

**Score quiver plots at additional particle counts ( $\nu = 0.24$ ).** Figure 14 shows score estimation quality across particle counts.

**Collision flow quiver plots ( $\nu = 0.24$ ).** Figures 15 and 16 show the collision velocity field at six times for blob and SBTM at  $n = 5 \times 10^5$ .

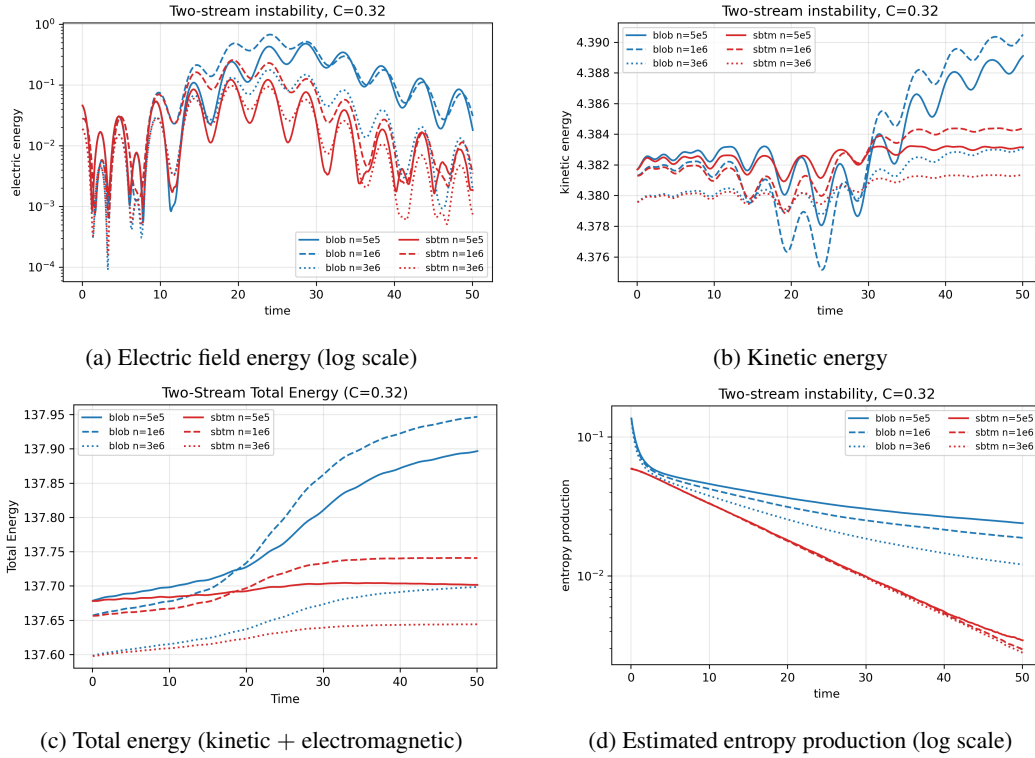


Figure 12: Two-stream instability ( $\nu = 0.32$ ): energy and estimated entropy diagnostics comparing blob (blue) and SBTM (red) at  $n = 5 \times 10^5$  (solid),  $10^6$  (dashed), and  $3 \times 10^6$  (dotted). SBTM produces consistent results across all particle counts, while the blob method overestimates estimated entropy production and exhibits energy drift.

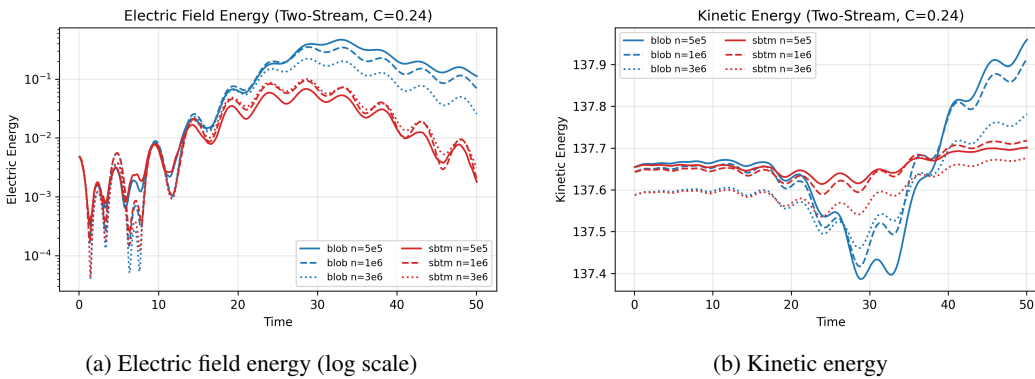


Figure 13: Two-stream instability ( $\nu = 0.24$ ): electric field energy and kinetic energy comparing blob (blue) and SBTM (red) across particle counts. SBTM produces consistent curves while the blob method varies with  $n$ .

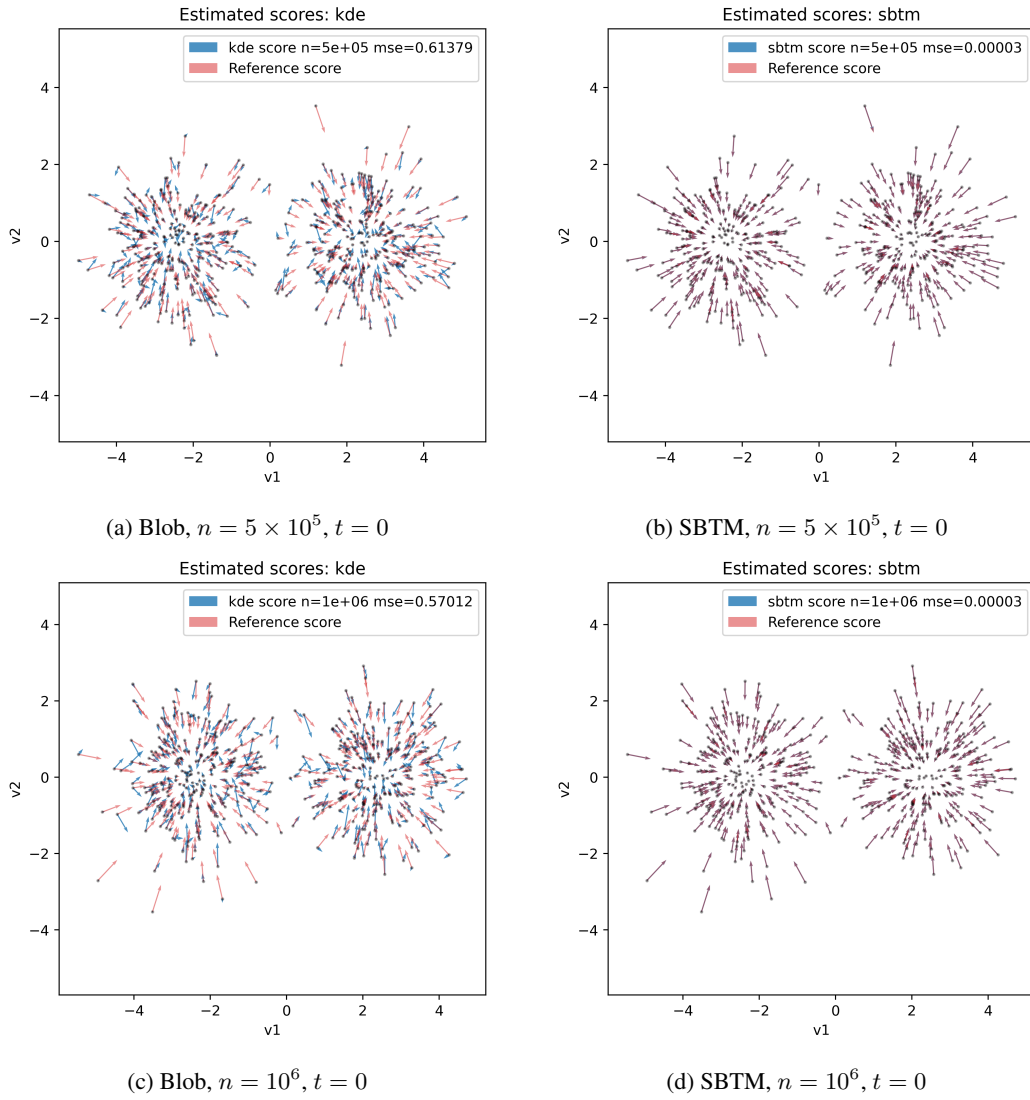


Figure 14: Two-stream instability: score quiver plots at  $\nu = 0.24$ , initial time, for additional particle counts. SBTM consistently achieves lower MSE than the blob method.

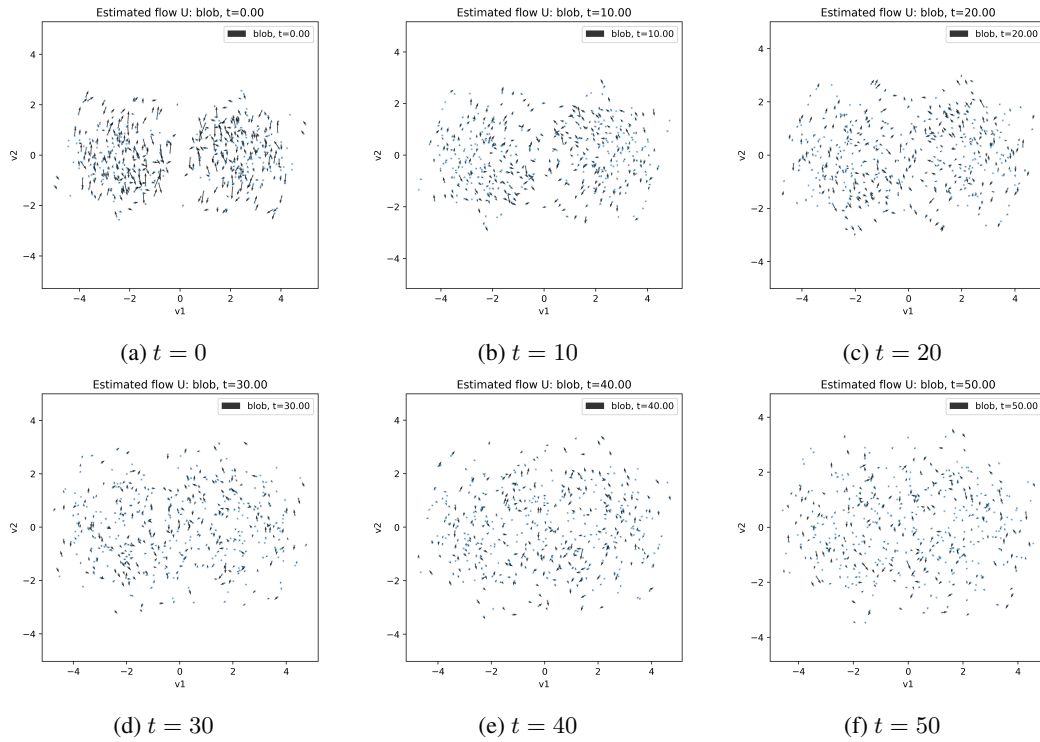


Figure 15: Two-stream instability: collision velocity field (blob) at  $\nu = 0.24$ ,  $n = 5 \times 10^5$ . The blob method produces noisy, incoherent collision fields that worsen as fine structure develops. Compare with SBTM (Figure 16).

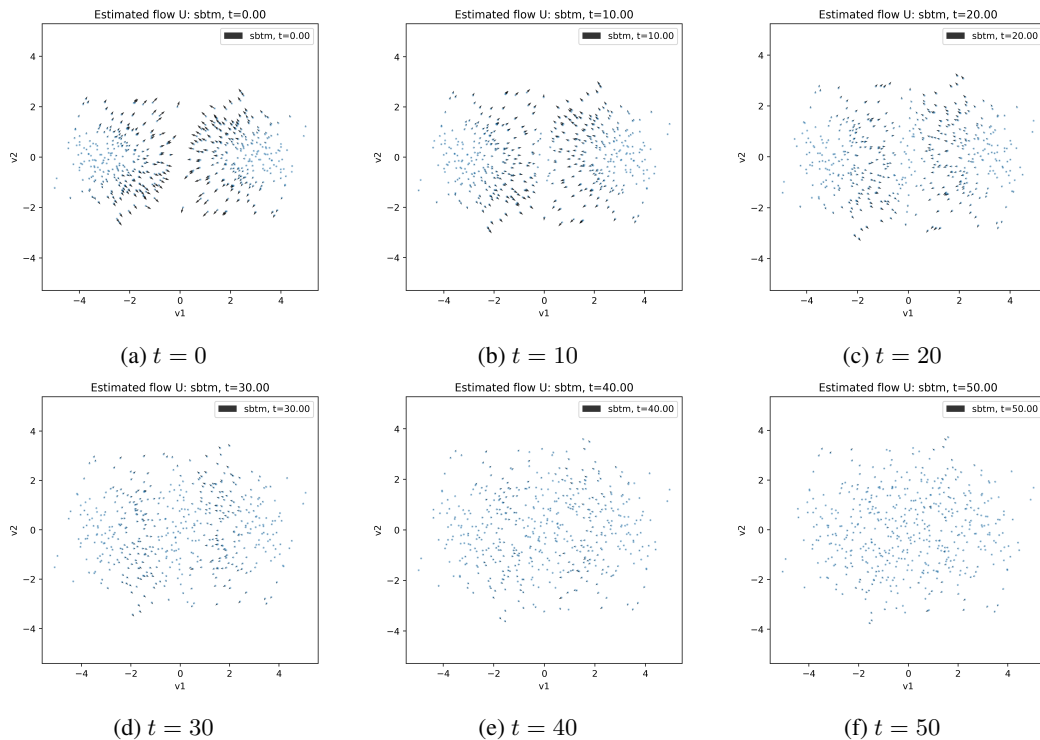


Figure 16: Two-stream instability: collision velocity field (SBTM) at  $\nu = 0.24$ ,  $n = 5 \times 10^5$ . SBTM produces smooth, coherent collision fields at all times, unlike the blob method (Figure 15).

## E ADDITIONAL WEIBEL INSTABILITY RESULTS

**Velocity distribution (linear scale).** Figure 17 shows the  $v_2$ -marginal density on a linear scale at  $\nu = 8 \times 10^{-4}$ .

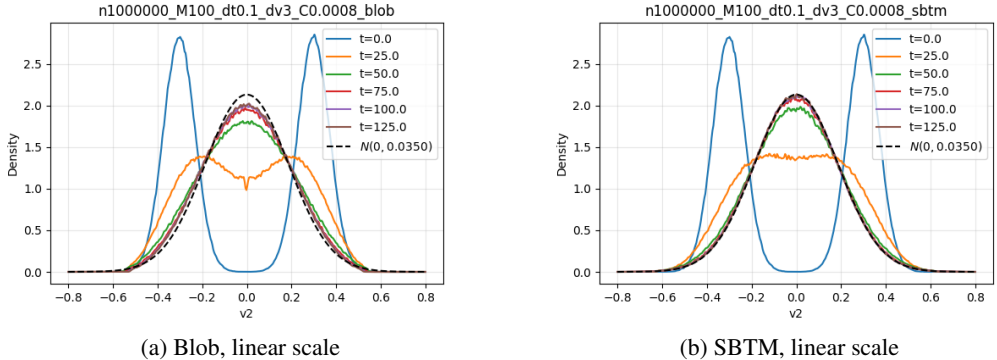


Figure 17: Weibel instability:  $v_2$ -marginal density evolution (linear scale) at  $\nu = 8 \times 10^{-4}$ . On a linear scale both methods appear similar, but the blob method produces a slightly broader distribution at late times. The difference is more visible in the log-scale version (Figure 4), where the blob method’s sharp tail cutoff is evident.

**Estimated entropy production.** Figure 18 shows estimated entropy production over time. SBTM correctly decays toward zero at all collision frequencies, with higher  $\nu$  producing faster decay. The blob method plateaus at nonzero values. At fixed  $\nu = 8 \times 10^{-4}$ , the blob plateau decreases with  $n$ , converging toward SBTM.

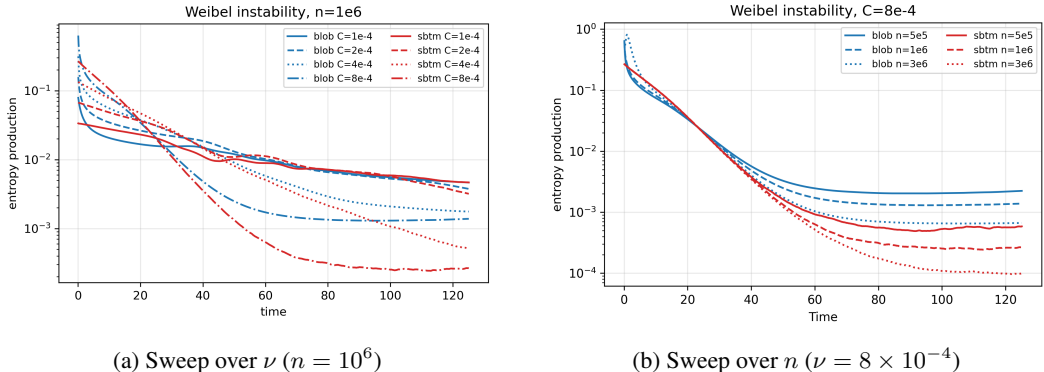


Figure 18: Weibel instability: estimated entropy production for blob (blue) and SBTM (red). Left: across collision frequencies; SBTM correctly decays to zero while the blob method plateaus. Right: across particle counts at  $\nu = 8 \times 10^{-4}$ ; the blob plateau decreases with  $n$ , converging toward SBTM.

**Energy diagnostics.** Figure 19 shows the energy diagnostics across collision frequencies. Electric and magnetic energies show the characteristic Weibel instability: exponential growth of the magnetic field, with higher  $\nu$  suppressing the instability more strongly. Blob and SBTM produce similar energy trajectories within each  $\nu$  value, confirming that the energy dynamics are primarily determined by the PIC framework rather than the score estimation method.

**Phase space.** Figure 20 shows the  $(v_1, v_2)$ -marginal density at several times for two collision frequencies. At low  $\nu$ , both methods are similar; at higher  $\nu$ , the blob method produces a non-Gaussian density while SBTM thermalizes correctly.

**Score estimation quality.** Figure 21 shows quiver plots of the estimated score in  $(v_1, v_2)$  at the initial and final times. At  $t = 0$ , SBTM achieves low MSE against the known analytic score; the

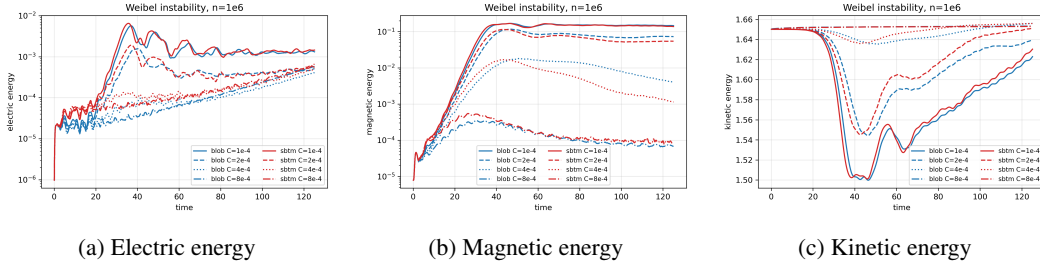


Figure 19: Weibel instability: energy diagnostics for blob (solid) and SBTM (dashed) across collision frequencies  $\nu \in \{10^{-4}, 2 \times 10^{-4}, 4 \times 10^{-4}, 8 \times 10^{-4}\}$ . Both methods produce similar energy trajectories within each  $\nu$ ; the score estimation method primarily affects the collision step.

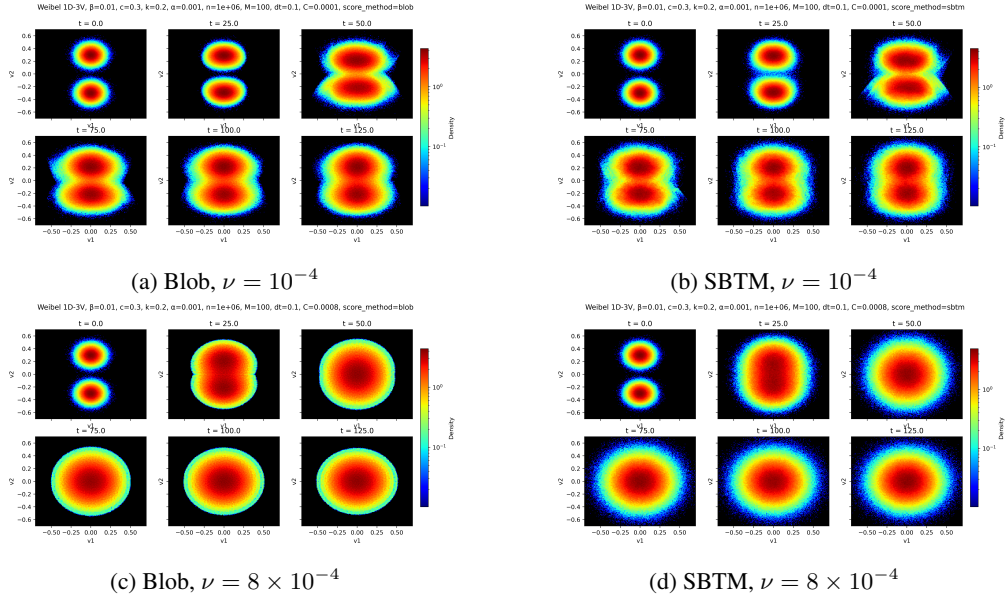


Figure 20: Weibel instability:  $(v_1, v_2)$ -marginal density. Top:  $\nu = 10^{-4}$  (weak collisions); bottom:  $\nu = 8 \times 10^{-4}$  (stronger collisions). At low  $\nu$ , both methods are similar. At higher  $\nu$ , the blob method produces a non-Gaussian density at late times, while SBTM thermalizes correctly.

blob method has substantially higher MSE. At  $t = T$ , SBTM scores are consistent with the expected linear score of a Gaussian ( $s(x, v) \propto -v$ ), while the blob method produces noisy scores in the tails.

**Velocity-space slices.** Figure 22 shows  $(v_1, v_2)$  density slices at a fixed spatial location  $x = 0$ . At  $\nu = 0$  (no collisions), both methods produce identical results, confirming that the PIC framework is the same. At  $\nu = 10^{-4}$ , the blob method already over-smooths the distribution, while SBTM retains the fine structures observed in the collisionless case.

#### PHASE SPACE SWEEP OVER $\nu$ AND $d_v$

Figures 23 and 24 show  $(v_1, v_2)$ -marginal density snapshots for the Weibel instability across collision frequencies  $\nu \in \{0, 10^{-4}, 2 \times 10^{-4}, 4 \times 10^{-4}, 8 \times 10^{-4}\}$  in  $d_v = 2$  and  $d_v = 3$  dimensions, respectively. At  $\nu = 0$ , both methods are identical. As  $\nu$  increases, the blob method increasingly produces non-Gaussian densities, while SBTM thermalizes correctly. The effect is more pronounced in  $d_v = 3$ , indicating better dimension scaling of SBTM.

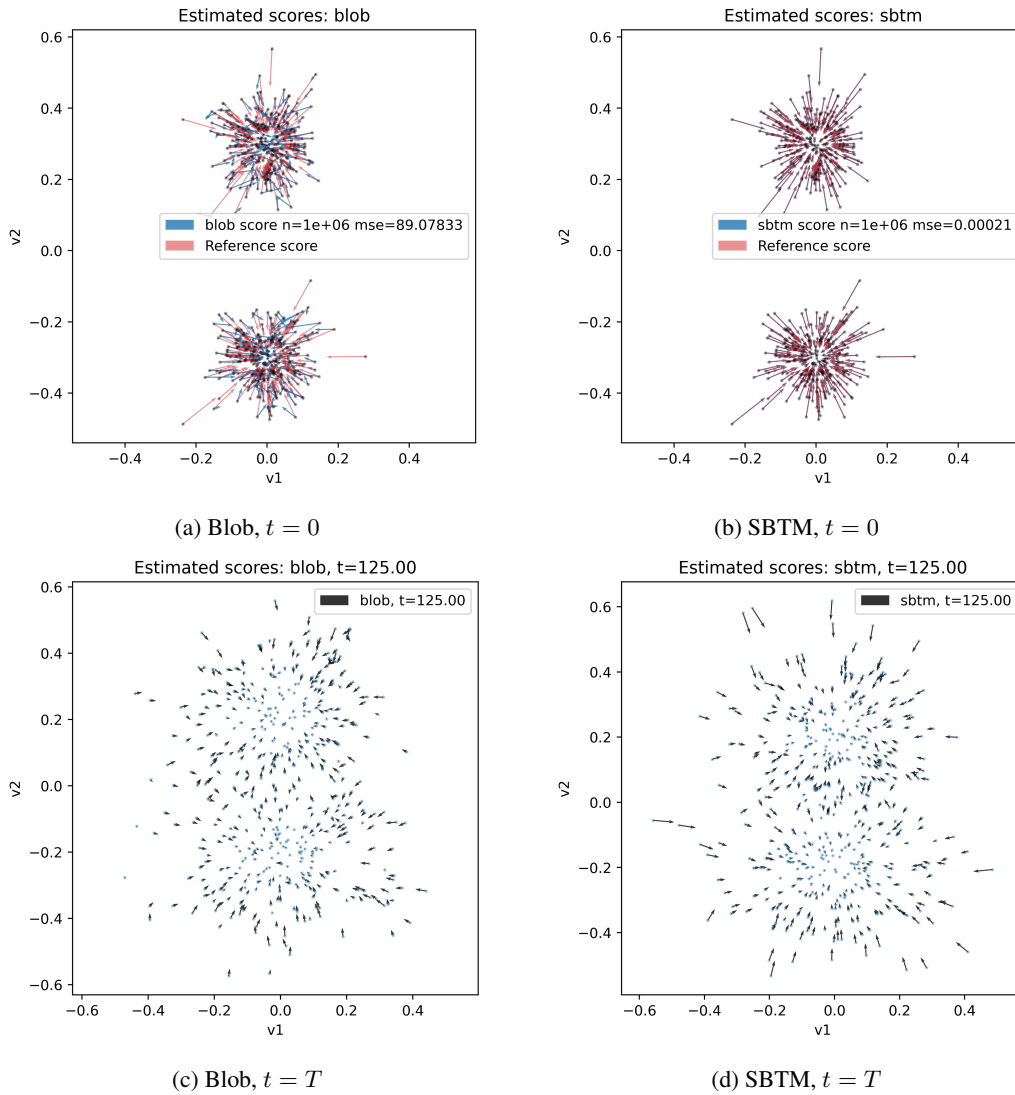


Figure 21: Weibel instability: score quiver plots at  $\nu = 10^{-4}$ . Top: initial time (true score in red; MSE in legend). Bottom: final time (true score unavailable). SBTM produces smooth, physically consistent scores; the blob method has high MSE at  $t = 0$  and noisy scores at  $t = T$ .

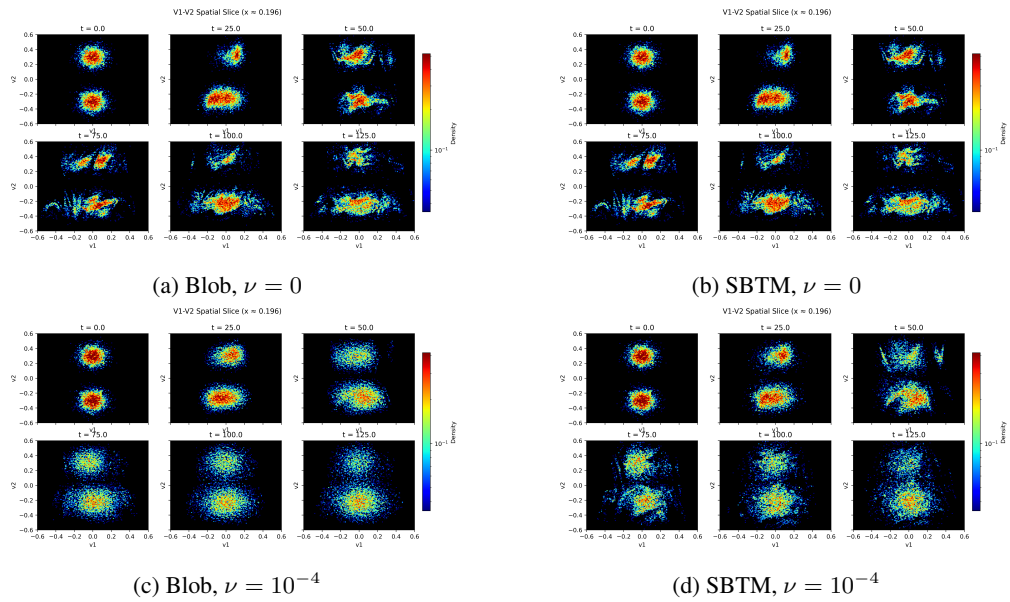


Figure 22: Weibel instability:  $(v_1, v_2)$  density slices at  $x = 0$ . At  $\nu = 0$  (top), both methods are identical. At  $\nu = 10^{-4}$  (bottom), the blob method over-smooths while SBTM preserves fine structures.

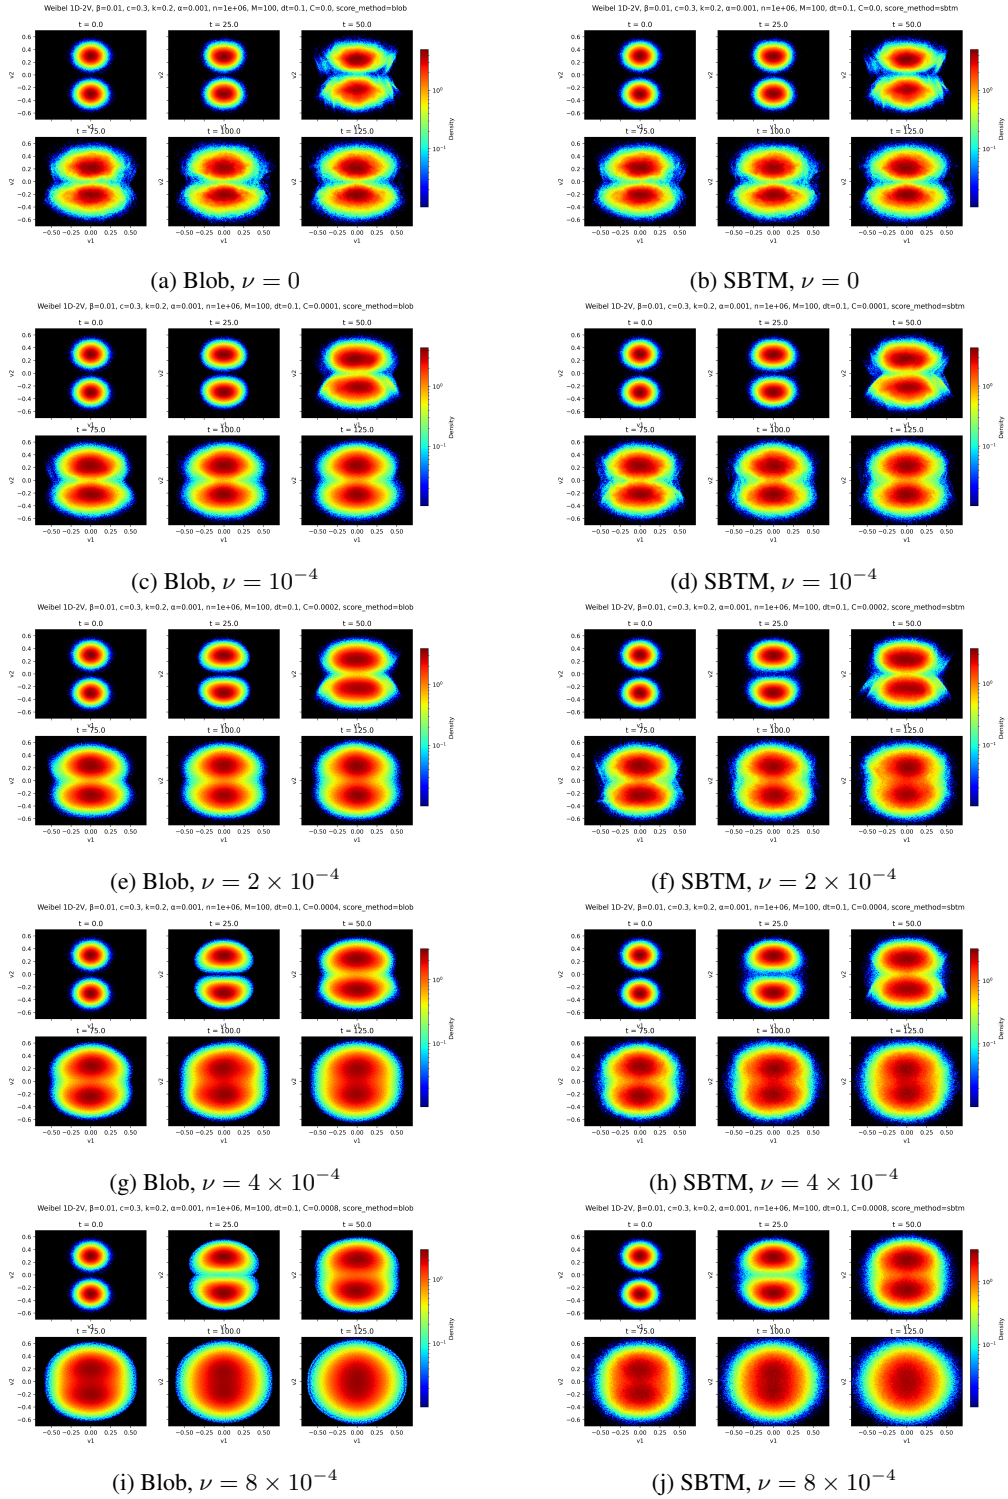


Figure 23: Weibel instability,  $d_v = 2$ :  $(v_1, v_2)$ -marginal density across collision frequencies. At  $\nu = 0$  both methods are identical. As  $\nu$  increases, the blob method produces increasingly non-Gaussian densities, while SBTM thermalizes correctly.

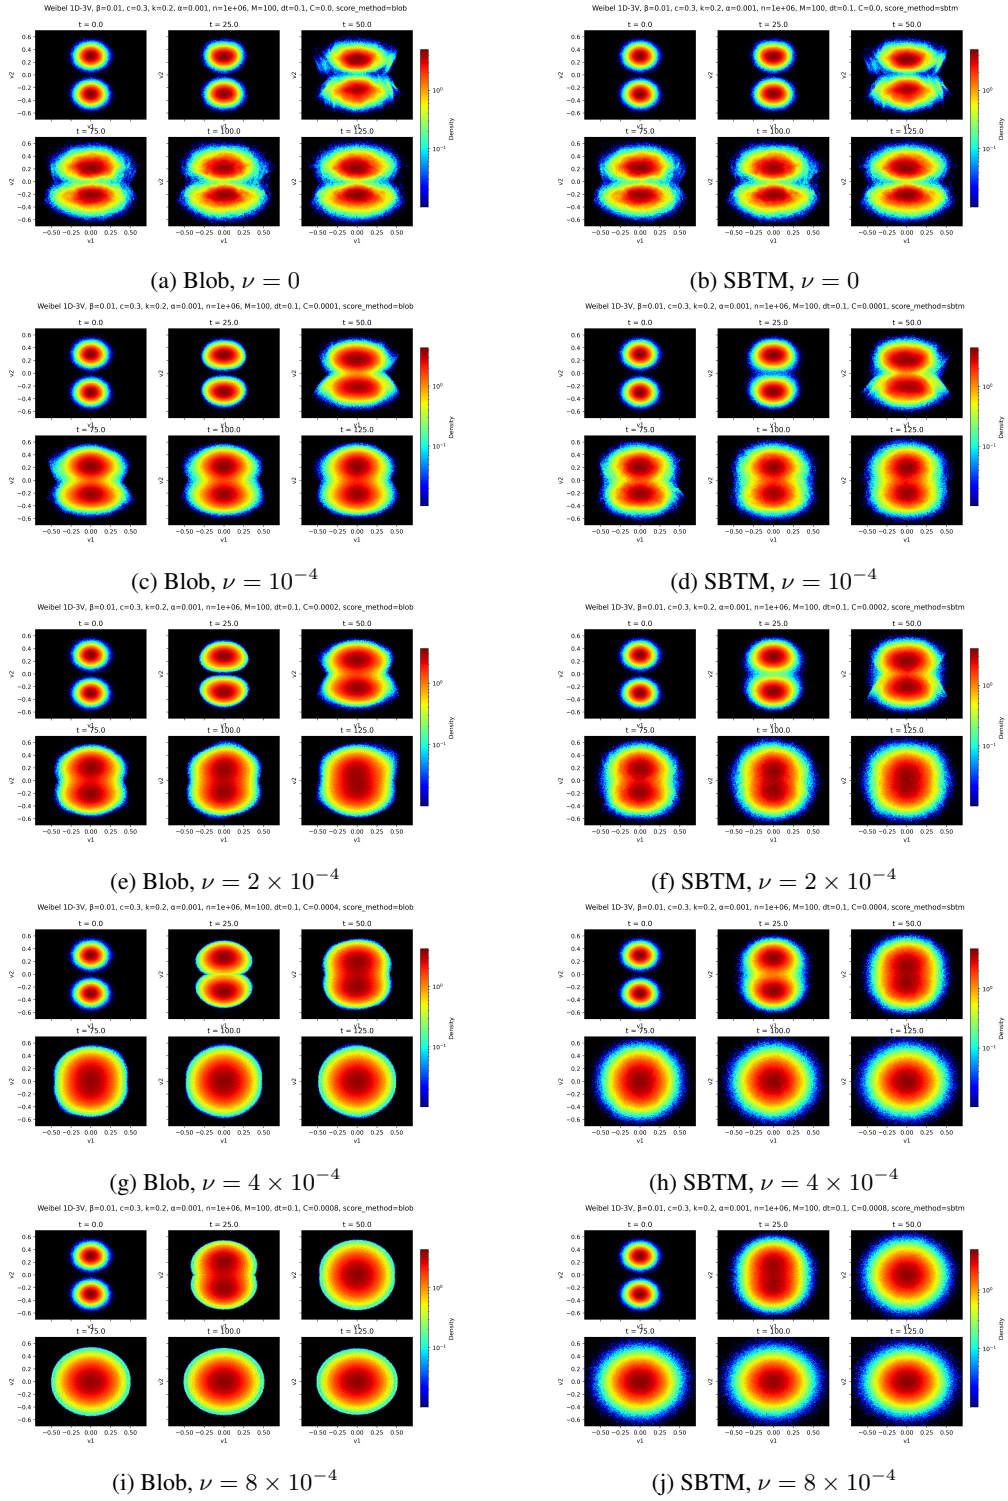


Figure 24: Weibel instability,  $d_v = 3$ :  $(v_1, v_2)$ -marginal density across collision frequencies. The blob method’s thermalization failure is more pronounced than in  $d_v = 2$  (Figure 23), indicating that SBTM scales better to higher dimensions.