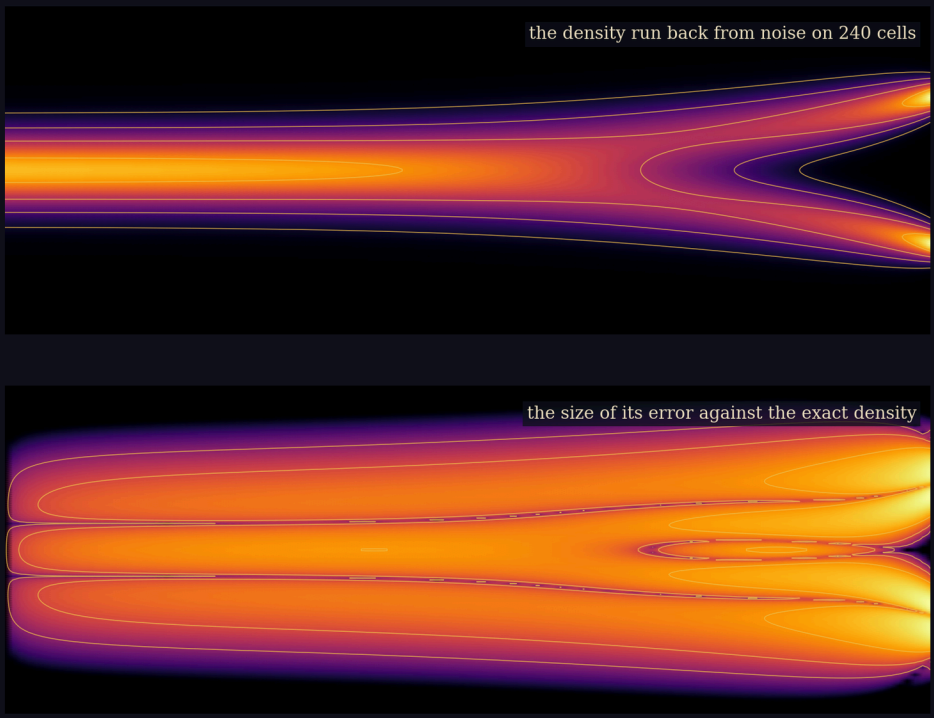


A density run back from noise on a grid, and the size of its error.

Numerical Methods for Fokker–Planck Equations I

Finite Volumes, Time Steps and Phase Space

Sanne Kooistra

ODIN PRESS

Numerical Methods for Fokker–Planck Equations I

Copyright © 2026 by **Sanne Kooistra**.
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Paperback version, First edition: 2026.
ISBN: forthcoming (Bowker allocation pending)
Trim Size: 6×9 inch. Paper color: White.

Source-code copyright. All Python source code listings reproduced in this volume—including those in the Appendix, the chapter bodies, and any auxiliary file distributed alongside this book—are the copyrighted intellectual property of **Sanne Kooistra**. No reproduction, redistribution, adaptation, translation into another programming language, or incorporation into any derivative work, whether academic, commercial, or otherwise, is permitted without the prior written permission of the copyright holder. Single-machine execution by the individual purchaser of this volume for the sole purpose of reproducing a figure or table from this book is granted as a limited, non-transferable, non-sublicensable use and does not constitute redistribution.

Published by Odin Press.
`https://odinpress.org`
This book was typeset in LaTeX by Odin Press from a manuscript submitted by the author.

9 8 7 6 5 4 3 2 1 0

Preface

The same equation has followed me through three jobs. In a semiconductor device it moved electrons and holes. In a fusion plasma it moved the velocities of electrons. In a population of spiking neurons it moved membrane potentials towards a threshold. Each time it had another name, and each time the first scheme we wrote made the same mistakes: it lost mass, it went negative, or it drifted away from an equilibrium that the physics said it must hold. The equation was a Fokker–Planck equation every time. These two volumes are the book I wanted then.

Its argument is short. A Fokker–Planck equation carries structure: it keeps mass, it keeps a density positive, it has an equilibrium, often a Gibbs density with detailed balance, and a relative entropy that only decreases. A numerical method is good when its discrete version keeps that structure, and every method in the two volumes, from finite volumes to time steppers, spectral expansions, particles, tensor trains and neural networks, is judged by the same list: what it keeps, what it loses, and what it costs in accuracy. Accuracy and structure are different virtues. A scheme can be more accurate and still wrong, and the book shows the two side by side, measured, wherever they part.

A second thread runs underneath. A positive scheme that conserves mass is a continuous-time Markov chain on its grid. Its forward equation is then the transpose of a backward Kolmogorov scheme, and an expectation computed forward, by moving the density, and backward, by moving the observable, agrees to round-off. The third thread is modern: a diffusion model noises its data with an Ornstein–Uhlenbeck process and generates by running time backwards, and both directions are Fokker–Planck equations. With Gaussian-mixture data every density and every score is exact, so the generative models can be tested like any other scheme.

The book is written for numerical analysts, for physicists and engineers who solve drift–diffusion equations in devices, plasmas and neurons, for

quantitative modellers of rates and populations, for researchers who train diffusion models and want to know what their samplers keep, and for graduate students who want the proofs whole. It assumes Itô calculus and stochastic differential equations, the basic theory of Markov processes, linear algebra with eigenvalues and matrix exponentials, and a first course in the numerical solution of differential equations. Everything else is defined where it is first used, and every result is proved or cited.

The work comes in two volumes. This first volume lays the foundation and builds the grid. Chapter 1 derives the equation from its stochastic differential equation and introduces the running examples with their exact solutions: the Ornstein–Uhlenbeck and Cox–Ingersoll–Ross processes, the double-well potential, the integrate-and-fire neuron, the Wright–Fisher diffusion with its atoms at loss and fixation, and the noising of a Gaussian mixture. Chapter 2 states the structure a scheme must keep. Chapters 3 and 4 build finite volumes in one dimension, with the Scharfetter–Gummel and Chang–Cooper fluxes, and the time steps that keep positivity. Chapter 5 carries both into phase space and the kinetic equation. Volume II continues the numbering: spectral methods and escape rates in Chapter 6, schemes read as Markov chains in Chapter 7, nonlinear and nonlocal equations in Chapter 8, particle methods in Chapter 9, and tensor trains and neural solvers in Chapter 10, with a guide to choosing a method and a test suite of exact solutions in its Appendix B. This volume stands on its own; the second assumes it.

The algorithms are in the chapters, as boxes, one step to a line, beside the theorems that say why they work. The Python is in Appendix A: one script for each of this volume's five chapters, printed whole. Every figure in this volume is drawn from those scripts, and every number in its text comes from them, computed twice, by two different routes.

The order of each argument is the order in which I learned to check a code. Count the mass first. Then find the smallest value of the density. Then ask whether the scheme holds its equilibrium exactly. Only then look at the error, and only beside an exact solution.

Sanne Kooistra

Contents

Preface vii

List of Figures xiii

List of Tables xvii

1 Where Fokker–Planck Equations Come From 1

1.1 From a Particle to a Density: the Generator, Its Adjoint and the Flux 2

1.2 Two Equations for One Expectation: Backward Kolmogorov and Duality 8

1.3 Interest Rates: Ornstein–Uhlenbeck and Cox–Ingersoll–Ross in Closed Form 12

1.4 The Double Well and the Integrate-and-Fire Neuron . 19

1.5 Wright–Fisher Diffusion: Fixation, Heterozygosity and Kimura's Density 24

1.6 Diffusion Models Read as Fokker–Planck Equations . . 29

1.7 A First Failure, and the Map of the Book 35

2 The Structure a Scheme Must Keep 43

2.1 Conservation Form and What Happens at a Wall . . . 44

2.2 Positivity, the Maximum Principle and Contraction . . 49

2.3 Equilibrium, Detailed Balance and Circulating Flux . . 53

2.4 Relative Entropy, Fisher Information and the H-Theorem 57

2.5 Poincaré, Log-Sobolev and the Spectral Gap 60

2.6 The Fokker–Planck Equation as a Wasserstein Gradient Flow 66

2.7 Five Properties of a Structure-Preserving Scheme . . . 69

2.8 Learned Scores Break Their Fokker–Planck Equation . 73

3 Finite Volumes in One Dimension **77**
3.1 Cells, Faces and the Semi-Discrete Balance 78
3.2 Central and Upwind Fluxes and the Cell Péclet Number 81
3.3 The Scharfetter–Gummel Flux from a One-Cell Problem 87
3.4 Chang–Cooper and the Stolarsky Family of Equilibrium Fluxes 91
3.5 A Discrete H-Theorem and What It Does Not Promise 96
3.6 Walls: the CIR Degenerate Boundary and the Wright–Fisher Vertex Scheme 101
3.7 The Reverse-Time Equation on a Grid with the Exact Score 107

4 Time Steps That Keep Positivity **111**
4.1 The Semi-Discrete System and the Explicit Step 112
4.2 Stiffness and the Spectrum of the Grid Operator 115
4.3 Implicit Euler and the M-Matrix 117
4.4 θ-Methods, Crank–Nicolson and the Order Barrier . . 122
4.5 Patankar Weights and Modified Patankar–Runge–Kutta Schemes 126
4.6 Exponential Integrators and Krylov Projections 130
4.7 IMEX Steps, Splitting and the Lost Equilibrium 135
4.8 The Reverse-Time Equation Near $t = 0$ 140

5 Phase Space and the Kinetic Equation **145**
5.1 Finite Volumes in the Plane: Anisotropy and Mixed Derivatives 146
5.2 Divergence-Free Fluxes from a Stream Function 151
5.3 The Klein–Kramers Equation in Phase Space 156
5.4 A Chain That Keeps the Gibbs Density 160
5.5 Discrete Hypocoercivity 168
5.6 Kramers Escape and the Turnover on the Grid 173

A Source Code **181**
A.1 Chapter 1 — Where Fokker–Planck Equations Come From 181
A.2 Chapter 2 — The Structure a Scheme Must Keep . . . 203
A.3 Chapter 3 — Finite Volumes in One Dimension 229
A.4 Chapter 4 — Time Steps That Keep Positivity 259

A.5 Chapter 5 — Phase Space and the Kinetic Equation . 288

References **315**

Index **327**

About the Author **329**

Chapter 1

Where Fokker–Planck Equations Come From

A Fokker–Planck equation moves probability. Its unknown is a density $p(t, x)$, and a density carries four things that no computation may lose: its mass, which stays one; its sign, which stays nonnegative; an equilibrium, which the flow approaches; and a relative entropy to that equilibrium, which only decreases. This book is about numerical methods for these equations, and its test of a method is short. Does the discrete solution keep the mass to round-off? Does it stay nonnegative at every step? Does it hold its equilibrium exactly, and does a discrete entropy decrease? And, as a separate question, how large is its error against an exact solution? Accuracy and structure are different virtues. A scheme can have a small error and still be wrong. Finite volumes, time steppers, spectral expansions, Markov chains, particles, tensor trains and neural networks are all judged by this one list in the chapters that follow.

Chapter 1 has two jobs. The first is to derive the equation: from a stochastic differential equation, by Itô's formula and one integration by parts, to the Fokker–Planck equation, its probability flux J, and its backward partner, the Kolmogorov equation for expectations. The second is to fix the test problems, each with the exact answer the book checks against: two short-rate models from finance, a particle in a double well from physics, a population of integrate-and-fire neurons from neuroscience, the Wright–Fisher diffusion from population genetics, and a fifth from machine learning, the noising of data in a diffusion model, where the score $\nabla \log p$ is known in closed form. The chapter ends with a first failure, proved rather than ob-

served: central differences on a strong-drift Ornstein–Uhlenbeck equation keep the mass exactly and lose the sign.

The seven sections follow that order: the derivation (§1.1), the backward equation and duality (§1.2), the interest-rate models (§1.3), the double well and the neuron (§1.4), Wright–Fisher (§1.5), diffusion models (§1.6), and the first failure together with a map of the methods against the structure each keeps (§1.7).

1.1 From a Particle to a Density: the Generator, Its Adjoint and the Flux

Let W be a Brownian motion in $\mathbb{R}^{d'}$ and let $X_t \in \mathbb{R}^d$ solve the stochastic differential equation

$$dX_t = b(X_t)\,dt + \sigma(X_t)\,dW_t,$$

with drift b and dispersion σ. A single path of X is random; the law of X_t is not. The question of this section is which equation that law obeys. The answer comes in two steps: Itô's formula gives an equation for expectations, and an integration by parts turns it into an equation for the density.

Assumption 1.1 (Standing assumptions on the diffusion)**.** The initial value X_0 is independent of W with $\mathbb{E}|X_0|^2 < \infty$. The drift $b : \mathbb{R}^d \to \mathbb{R}^d$ and the dispersion $\sigma : \mathbb{R}^d \to \mathbb{R}^{d\times d'}$ are locally Lipschitz with linear growth, $|b(x)| + \|\sigma(x)\| \le C(1+|x|)$, so that the equation has a unique strong solution with $\mathbb{E}\sup_{s\le t}|X_s|^2 < \infty$ for every t; we write $a := \sigma\sigma^{\mathsf{T}}$ for the diffusion matrix. Wherever a density is differentiated, b and a are C^∞ with bounded derivatives and a is uniformly elliptic: $\xi^{\mathsf{T}} a(x)\xi \ge c_0|\xi|^2$ for some $c_0 > 0$ and all $x, \xi \in \mathbb{R}^d$.

Several examples break these hypotheses on purpose: two diffusions degenerate at a wall (§1.3, §1.5), one diffuses in velocity only and one is reset at a threshold (§1.4). The degeneracies are where schemes go wrong, which is why they are in the book.

For $f \in C^2(\mathbb{R}^d)$, Itô's formula reads

$$f(X_t) = f(X_0) + \int_0^t \mathscr{L} f(X_s)\,ds + \int_0^t \nabla f(X_s)^{\mathsf{T}} \sigma(X_s)\,dW_s,$$

where the operator $\mathscr{L}$ collects the drift and the second-order Itô correction. It has a name, and it is the first object of the book.

Definition 1.2 (Generator)**.** For $f \in C_c^2(\mathbb{R}^d)$ the generator of the diffusion is

$$\mathscr{L} f := b \cdot \nabla f + \tfrac{1}{2}\, a : \nabla^2 f = \sum_i b_i\, \partial_i f + \tfrac{1}{2} \sum_{i,j} a_{ij}\, \partial_{ij} f.$$

By Itô's formula the process $f(X_t) - f(X_0) - \int_0^t \mathscr{L} f(X_s)\, ds = \int_0^t \nabla f(X_s)^{\mathsf{T}} \sigma(X_s)\, dW_s$ is a martingale, and $\mathscr{L} f(x) = \lim_{t \downarrow 0} (\mathbb{E}[f(X_t) \mid X_0 = x] - f(x))/t$.

The martingale property holds because $\nabla f^{\mathsf{T}} \sigma$ is continuous with compact support, hence bounded; taking expectations, dividing by t and letting $t \downarrow 0$ gives the limit. The generator acts on observables. The density needs its adjoint.

Definition 1.3 (Adjoint operator)**.** For $p \in C^2(\mathbb{R}^d)$ the adjoint operator of $\mathscr{L}$ is

$$\mathscr{L}^* p := -\nabla \cdot (b\, p) + \tfrac{1}{2} \nabla \cdot \nabla \cdot (a\, p) = -\sum_i \partial_i (b_i p) + \tfrac{1}{2} \sum_{i,j} \partial_{ij} (a_{ij} p),$$

characterised by $\int_{\mathbb{R}^d} (\mathscr{L} f)\, p\, dx = \int_{\mathbb{R}^d} f\, (\mathscr{L}^* p)\, dx$ for all $f, p \in C_c^\infty(\mathbb{R}^d)$.

A density that evolves needs a kernel that moves it, and the kernel is the law of the process started at a point.

Definition 1.4 (Transition density)**.** If there is a measurable $p(t, x \mid y) \geq 0$ with $\mathbb{P}(X_t \in A \mid X_0 = y) = \int_A p(t, x \mid y)\, dx$ for every $t > 0$, every $y \in \mathbb{R}^d$ and every Borel set A, then $p(t, x \mid y)$ is the transition density of X. By the Markov property it satisfies the Chapman–Kolmogorov equation $p(t+s, x \mid y) = \int p(t, x \mid z)\, p(s, z \mid y)\, dz$ [1], and a process started from a density p_0 has the density $p(t, x) = \int p(t, x \mid y)\, p_0(y)\, dy$.

The theorem below is the forward equation of Kolmogorov [2], known in physics as the Fokker–Planck equation [3]; we follow the modern treatment of [1]. Its hypotheses ask for a smooth density. They do not ask for decay at infinity, because the test functions of the proof have compact support.

Theorem 1.5 (The Fokker–Planck equation)**.** Let the standing assumptions hold, let X_0 have density p_0, and suppose the law of X_t has a density

Chapter 3

Finite Volumes in One Dimension

A drift–diffusion solver that takes the logarithm of a carrier density cannot survive one negative cell, however small the negative value. The remedy is not a finer grid. It is a better flux, chosen so that the density cannot change sign, and this chapter is about choosing it.

Chapter 2 listed what a scheme must keep: mass, positivity, the equilibrium, a relative entropy that only decreases, and detailed balance where the equation has it. It did not say how to build a scheme that has them. The construction here is the finite volume method: the unknowns are the cell masses m_i, and a scheme is its choice of numerical flux $J_{i+1/2}$ through each face. Mass is then kept by construction, and every other property is decided face by face.

The chapter delivers a tested set of one-dimensional fluxes (central, upwind, Scharfetter–Gummel, Chang–Cooper and the Stolarsky family) with the cell Péclet number as the positivity criterion, a discrete H-theorem, exact wall treatments for the degenerate CIR boundary and for Wright–Fisher fixation, and a positive grid solver for the reverse-time equation of a diffusion model, every claim checked against an exact solution.

Section 3.1 sets up the semi-discrete balance and §3.2 reads the central and upwind fluxes through the cell Péclet number. Section 3.3 derives the Scharfetter–Gummel flux from a one-cell problem, and §3.4 finds it again as the Chang–Cooper flux and in the Stolarsky family. Section 3.5 proves a discrete H-theorem and shows what it does not promise, §3.6 puts mass

at the walls, and §3.7 runs the noising of a diffusion model backwards on a grid.

3.1 Cells, Faces and the Semi-Discrete Balance

Chapter 2 wrote the Fokker–Planck equation in conservation form, $\partial_t p + \partial_x J = 0$, and asked a scheme to keep mass, positivity, the equilibrium and a decreasing entropy. In one dimension, with drift $b(x)$ and a constant diffusion coefficient $D > 0$, the probability flux is $J = bp - D\,\partial_x p$, and a finite volume scheme follows the mass $m_i(t)$ of each cell.

Definition 3.1 (Finite volume method and cell average)**.** Split $[x_L, x_R]$ into N cells $K_i = [x_{i-1/2}, x_{i+1/2}]$ of width $h = (x_R - x_L)/N$ with centres x_i. For a density solving the conservation form $\partial_t p + \partial_x J = 0$ with probability flux $J = bp - D\,\partial_x p$, the cell mass is $m_i(t) = \int_{K_i} p(t,x)\,dx$ and the cell average is $p_i = m_i/h$. A finite volume method replaces the exact face flux $J(t, x_{i+1/2})$ by a numerical flux $J_{i+1/2}$ computed from cell averages and evolves $dm_i/dt = -(J_{i+1/2} - J_{i-1/2})$ for $i = 1, \dots, N$, with $J_{1/2} = J_{N+1/2} = 0$ at zero-flux walls.

Integrating the conservation form over K_i gives, with no approximation,

$$\frac{d}{dt}\int_{K_i} p(t,x)\,dx = -\big(J(t,x_{i+1/2}) - J(t,x_{i-1/2})\big).$$

The balance of a cell is exact; the only approximation is the face flux, and the sum over cells telescopes to the wall fluxes whatever numerical flux is used. The unknowns are cell masses, not point values, and the two differ:

$$p_i = \frac{1}{h}\int_{-h/2}^{h/2} p(x_i + s)\,ds = p(x_i) + \frac{h^2}{24}\,\partial_x^2 p(x_i) + O(h^4),$$

since the odd powers of s integrate to zero and the average of s^2 over the cell is $h^2/12$. Every error in this chapter is therefore measured against exact cell masses from a distribution function (normal for the Ornstein–Uhlenbeck process, noncentral chi-square for the CIR process), never against point values, so that the offset above does not enter a measured order.

Definition 3.2 (Two-point numerical flux, consistent and monotone)**.** A two-point numerical flux is $J_{i+1/2} = \alpha_{i+1/2}\,p_i - \beta_{i+1/2}\,p_{i+1}$ with coefficients

depending on h, D and the drift near $x_{i+1/2}$. It is consistent of order q if $\alpha_{i+1/2}\,p(x_i) - \beta_{i+1/2}\,p(x_{i+1}) = J(x_{i+1/2}) + O(h^q)$ for every smooth p, and monotone if $\alpha_{i+1/2} \geq 0$ and $\beta_{i+1/2} \geq 0$ at every interior face.

In masses the flux reads $J_{i+1/2} = (\alpha_{i+1/2} m_i - \beta_{i+1/2} m_{i+1})/h$: the face carries mass from cell i to cell $i+1$ at rate $\alpha_{i+1/2}/h$ and back at rate $\beta_{i+1/2}/h$. The balance becomes $dm/dt = Am$ with a tridiagonal matrix, $A_{i+1,i} = \alpha_{i+1/2}/h$, $A_{i,i+1} = \beta_{i+1/2}/h$ and $A_{ii} = -(\alpha_{i+1/2} + \beta_{i-1/2})/h$, where the wall faces contribute nothing ($\beta_{1/2} = \alpha_{N+1/2} = 0$).

Proposition 3.3 (Exact conservation of mass)**.** For every numerical flux and zero-flux walls, $\frac{d}{dt}\sum_{i=1}^{N} m_i = 0$. Equivalently the semi-discrete system $dm/dt = Am$ has zero column sums; when the flux is monotone the off-diagonal entries $A_{i+1,i} = \alpha_{i+1/2}/h$ and $A_{i,i+1} = \beta_{i+1/2}/h$ are nonnegative and $A = Q^{\mathsf{T}}$ for a rate matrix Q (rows summing to zero, off-diagonal entries nonnegative).

Proof. Summing the balance over the cells, $\frac{d}{dt}\sum_i m_i = -\sum_{i=1}^{N}(J_{i+1/2} - J_{i-1/2}) = -(J_{N+1/2} - J_{1/2}) = 0$, because every interior face flux appears once with each sign and the wall fluxes vanish. Column j of A holds $A_{j-1,j} = \beta_{j-1/2}/h$, $A_{jj} = -(\alpha_{j+1/2} + \beta_{j-1/2})/h$ and $A_{j+1,j} = \alpha_{j+1/2}/h$, with the wall conventions at $j = 1$ and $j = N$; they add up to zero. If the flux is monotone the off-diagonal entries are nonnegative, and $Q := A^{\mathsf{T}}$ has nonnegative off-diagonal entries and row sums equal to the column sums of A, which vanish. □

Mass needs nothing from the flux; positivity needs everything from it. Chapter 2 proved the next lemma for a constant matrix. The reverse-time equation of §3.7 has coefficients that move with time, so we prove the general form.

Lemma 3.4 (Positivity is decided face by face)**.** Let $A(t)$ be a real $N \times N$ matrix, continuous in t. The solution of $dm/dt = A(t)\,m$ is entrywise nonnegative for $t \geq s$, for every starting time s and every entrywise nonnegative $m(s)$, if and only if $A_{ij}(t) \geq 0$ for all $i \neq j$ and all t. Consequently a finite volume method with zero-flux walls is positivity preserving if and only if its numerical flux is monotone at every interior face and every time.

Proof. Suppose $A_{ij}(t) \geq 0$ for all $i \neq j$ and all t. Fix $s < t_1$, let $c = \max_{s \leq r \leq t_1} \max_i |A_{ii}(r)|$, finite by continuity, and put $B(r) = A(r) + cI$, entrywise nonnegative on $[s, t_1]$. The vector $\tilde{m}(r) = e^{c(r-s)} m(r)$ solves

Electrons in a device, particles in a double well, a population of firing neurons, gene frequencies drifting to fixation, and the noising of data in a diffusion model all move probability by the same law. A Fokker–Planck equation keeps mass, keeps its density positive, holds an equilibrium and lets a relative entropy only decrease. A scheme that loses one of these is wrong, however small its error.

This first volume derives the equation from its stochastic differential equation, sets out five running examples with their exact solutions, and states the structure a scheme must keep. It then builds the schemes that keep it: finite volumes with Scharfetter–Gummel and Chang–Cooper fluxes, time steps that stay positive at any size, and the kinetic equation in phase space with its Gibbs density held exactly. Where the structure-keeping scheme is the less accurate one, the book measures by how much.

Every main result is proved. Every figure and every number comes from the Python printed whole in the appendix, checked against an exact solution and computed two ways.

Sanne Kooistra is a numerical analyst from Delft who has simulated electrons in semiconductors and fusion plasmas and the firing of neuron populations, and has counted the mass at every step.

ISBN 13 978-873-95-1015-1
Paperback
9 788739 510151