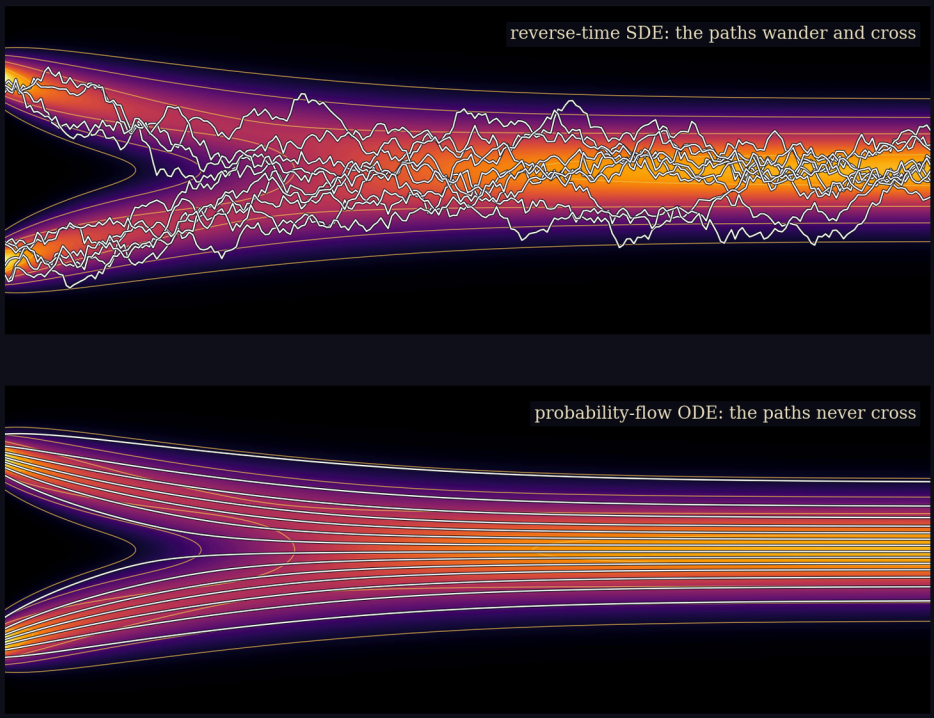


From noise back to data along two samplers of one density.

Numerical Methods for Fokker–Planck Equations II

Spectral Methods, Markov Chains, Particles and Neural Solvers

Sanne Kooistra

ODIN PRESS

Numerical Methods for Fokker–Planck Equations II

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

Volume I ended with a scheme that holds its equilibrium exactly in phase space. It took five chapters to get there: the equation derived from its stochastic differential equation, the structure a scheme must keep, finite volumes in one dimension, time steps that keep positivity, and the kinetic equation on a grid. This second volume asks what lies beyond that grid, and it judges every method by the same list as the first: mass, positivity, the equilibrium, and a relative entropy that only decreases. Accuracy is still a number and structure is still a property. They still part ways, and where they do, the book says so with the measurement beside it.

Chapter 6 turns to spectral methods, Hermite expansions on the line and Gegenbauer polynomials on the Wright–Fisher interval, and uses them to compute what a grid computes slowly: the spectral gap, and through it the rate at which a particle escapes a double well. Chapter 7 reads every scheme of the book as a continuous-time Markov chain. Its forward equation is the transpose of a backward Kolmogorov scheme, so an expectation computed forward and backward agrees to round-off; committors and mean first-passage times come out of the same matrix, and the time reversal of a diffusion model is exact on the grid. Chapter 8 treats equations whose drift depends on the density itself: the Desai–Zwanzig model with its phase transition, the Keller–Segel equation, and populations of neurons that excite one another, where the self-consistent Scharfetter–Gummel scheme keeps the free energy decreasing and a Lagrangian scheme keeps the transient more accurately. Chapter 9 gives up the grid for particles, which keep mass and positivity for nothing and pay in statistical error, and it reads the sampler of a diffusion model as a particle method. Chapter 10 goes to four, five and six dimensions, where a full grid becomes expensive, and measures tensor trains and a score network against the grid on problems with exact answers. Appendix B closes the work with a guide from

the features of a problem to a method, and a test suite of exact solutions to hold any new method against.

The numbering continues from Volume I, so a reference to Chapter 3 or to Figure 3.2 means that chapter and that figure of the first volume. This volume assumes it, or an equivalent knowledge of finite volumes for the Fokker–Planck equation, together with Itô calculus, the basic theory of Markov processes and linear algebra with eigenvalues and matrix exponentials. Everything else is defined where it is first used, and every result is proved or cited.

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. One script trains a small network, on a graphics card, and the chapter says how long and on what.

The checks have not changed between the volumes. Count the mass. Find the smallest value of the density. Ask whether the method holds its equilibrium. Then look at the error, beside an exact solution, and ask what the method cost to get it.

Sanne Kooistra

Contents

Preface vii

List of Figures xiii

List of Tables xvii

6 Spectral Methods and Escape Rates 1

6.1 From Fokker–Planck to Schrödinger: the Ground-State Transform . . . 2

6.2 Hermite Functions, the Mehler Formula and the Noising of Data . . . 5

6.3 Hermite–Galerkin for the Double Well . . . 8

6.4 Spectral Gaps, Metastability and Escape Rates . . . 14

6.5 Gegenbauer Polynomials and Kimura's Solution . . . 20

6.6 Fourier Modes for a Tilted Periodic Potential . . . 23

6.7 Gibbs Oscillations, Filters and Lost Positivity . . . 29

7 Schemes as Markov Chains 35

7.1 Positive Conservative Schemes Are Continuous-Time Markov Chains . . . 36

7.2 The Forward Scheme Is the Transpose of the Backward Scheme . . . 41

7.3 Detailed Balance, Reversible Chains and the Dirichlet Form . . . 45

7.4 Local Consistency and the Convergence of Expectations . . . 50

7.5 Committors, Absorption Times and the Moran Model . 55

7.6 Time Reversal on the Grid and the Discrete Diffusion Model . . . 63

8 Nonlinear and Nonlocal Equations **69**
8.1 McKean–Vlasov Drifts and Their Free Energy 70
8.2 Finite Volumes That Dissipate the Free Energy 73
8.3 The Desai–Zwanzig Phase Transition 80
8.4 The JKO Step in Mass Coordinates 85
8.5 Keller–Segel Collapse Above the Critical Mass 90
8.6 Neurons That Fire and Reset 95

9 Particle Methods **103**
9.1 Euler–Maruyama Ensembles and the Empirical Measure 104
9.2 Weak Error and Statistical Error 107
9.3 Kernel Density Estimates and the Bandwidth Ridge . . 110
9.4 Interacting Particles and Propagation of Chaos 115
9.5 Deterministic Particles and the Blob Method 120
9.6 Escape Times by Particles 123
9.7 Sampling a Diffusion Model Is a Particle Method . . . 128
9.8 Grids or Particles as the Dimension Grows 133

10 Tensor Trains, Neural Solvers and the Grid **139**
10.1 Low-Rank Densities and the Tensor-Train Format . . . 140
10.2 The Fokker–Planck Operator in Tensor-Train Form . . 145
10.3 Physics-Informed Networks and the Normalisation Constant 152
10.4 The Score Fokker–Planck Equation and Score-Based Solvers 155
10.5 Score Matching and Denoising Score Matching 159
10.6 A Benchmark in Four to Six Dimensions 162
10.7 What Each Method Keeps: the Decision Map 168

A Source Code **173**
A.1 Chapter 6 — Spectral Methods and Escape Rates . . . 173
A.2 Chapter 7 — Schemes as Markov Chains 201
A.3 Chapter 8 — Nonlinear and Nonlocal Equations 225
A.4 Chapter 9 — Particle Methods 253
A.5 Chapter 10 — Tensor Trains, Neural Solvers and the Grid 287

B Choosing a Method, and the Exact Solutions That Check It **325**

B.1 Choosing a Method 325
B.2 The Test Suite of Exact Solutions 332
B.3 The Gaussian-Mixture Benchmark 341

References **349**

Index **363**

About the Author **365**

Chapter 6

Spectral Methods and Escape Rates

A Hermite expansion of a neuron population density can converge faster than any power of the number of modes and still dip below zero next to the firing threshold, so that the firing rate read from it is negative for a moment. Accurate and right are different words, and this chapter is where they are furthest apart. Chapters 3 to 5 (Volume I) built finite volumes that keep the structure exactly at first or second order. Spectral methods turn that trade around: they expand the density in eigenfunctions of an operator close to the Fokker–Planck operator, their error falls geometrically in the number of modes, and they give up the exact structure to get there. This chapter opens Volume II; Volume I derived the equation and the structure a scheme must keep, and built the grid schemes that keep it in one dimension, in time and in phase space.

The chapter delivers a Hermite–Galerkin solver for the ground-state-transformed operator, with Gegenbauer and Fourier counterparts, that computes the double-well spectral gap and escape rate, Kimura's solution and the effective diffusion of a tilted periodic potential to round-off against exact checks, and a measured account of what each loses. Section 6.1 turns the Fokker–Planck operator into a Schrödinger operator $\mathscr{H}$; Section 6.2 solves the Ornstein–Uhlenbeck case in Hermite functions and reads its gap as the rate at which a diffusion model forgets its data; Sections 6.3 and 6.4 compute the double-well gap and read it as an escape rate; Sections 6.5 and 6.6 treat Kimura's solution and a tilted periodic potential; Section 6.7 counts what spectral methods lose.

6.1 From Fokker–Planck to Schrödinger: the Ground-State Transform

Chapters 3 to 5 discretised the flux J face by face; an expansion of p in eigenfunctions needs instead an operator $\mathscr{L}^*$ that is self-adjoint in one fixed inner product. For a gradient drift there is one.

Assumption 6.1 (Confining gradient diffusion)**.** Let $b = -D\nabla V$ with constant $D > 0$, $V \in C^\infty(\mathbb{R}^d)$, $Z = \int e^{-V}\,dx < \infty$ and $p_\infty = e^{-V}/Z$, so that $\partial_t p = \mathscr{L}^* p = D\nabla \cdot (\nabla p + p\nabla V)$ is the Fokker–Planck equation of the overdamped Langevin equation $dX_t = -D\nabla V(X_t)\,dt + \sqrt{2D}\,dW_t$. The Schrödinger potential $V_{\mathrm{S}} = \frac{1}{4}|\nabla V|^2 - \frac{1}{2}\Delta V$ is bounded below and $V_{\mathrm{S}}(x) \to \infty$ as $|x| \to \infty$.

The generator $\mathscr{L} f = D p_\infty^{-1} \nabla \cdot (p_\infty \nabla f)$ is symmetric in $L^2(p_\infty)$ and the Fokker–Planck operator in $L^2(p_\infty^{-1})$; the square root of the Gibbs density sits halfway.

Definition 6.2 (Ground-state transform)**.** Under the assumption, the ground-state transform is the map $\mathcal{U} : L^2(\mathbb{R}^d; p_\infty^{-1}dx) \to L^2(\mathbb{R}^d)$, $\mathcal{U}p = \psi = p/\sqrt{p_\infty}$, with inverse $p = \sqrt{p_\infty}\,\psi$. It is unitary, $\int \mathcal{U}p\,\mathcal{U}q\,dx = \int pq\,p_\infty^{-1}dx$, it sends the Gibbs density p_∞ to $\sqrt{p_\infty}$, and the mass of p is $\int \sqrt{p_\infty}\,\psi\,dx$.

Definition 6.3 (Schrödinger operator of a gradient diffusion)**.** The Schrödinger operator of the diffusion is $\mathscr{H}\psi = D(-\Delta\psi + V_{\mathrm{S}}\psi)$ on $L^2(\mathbb{R}^d)$, with Schrödinger potential $V_{\mathrm{S}} = \frac{1}{4}|\nabla V|^2 - \frac{1}{2}\Delta V$; it is the Witten Laplacian on functions, scaled by D [1].

With $g = \sqrt{p_\infty} = Z^{-1/2}e^{-V/2}$ we have $\nabla g = -\frac{1}{2}g\nabla V$ and $\Delta g = -\frac{1}{2}\nabla g \cdot \nabla V - \frac{1}{2}g\,\Delta V = g(\frac{1}{4}|\nabla V|^2 - \frac{1}{2}\Delta V)$, that is,

$$\Delta\sqrt{p_\infty} = V_{\mathrm{S}}\,\sqrt{p_\infty}.$$

Theorem 6.4 (The ground-state transform turns the Fokker–Planck operator into minus a Schrödinger operator)**.** Under the assumption, for smooth compactly supported ψ: $\mathscr{L}^*(\sqrt{p_\infty}\,\psi) = -\sqrt{p_\infty}\,\mathscr{H}\psi$ and $\mathscr{L}(\psi/\sqrt{p_\infty}) = -\mathscr{H}\psi/\sqrt{p_\infty}$; moreover $\langle \psi, \mathscr{H}\psi\rangle = D\int p_\infty\,|\nabla(\psi/\sqrt{p_\infty})|^2\,dx \geq 0$ and $\mathscr{H}\sqrt{p_\infty} = 0$. Hence $\mathscr{H}$ is symmetric and nonnegative, and $p = \sqrt{p_\infty}\,\psi$ solves the Fokker–Planck equation exactly when $\partial_t\psi = -\mathscr{H}\psi$.

Proof. Write $g = \sqrt{p_\infty}$ and $p = g\psi$. In gradient form $\mathscr{L}^* p = D\nabla \cdot (p_\infty \nabla(p/p_\infty))$, and $p_\infty \nabla(\psi/g) = g\nabla\psi - \psi\nabla g$. Hence

$$\mathscr{L}^*(g\psi) = D(\nabla g \cdot \nabla\psi + g\Delta\psi - \nabla\psi \cdot \nabla g - \psi\Delta g),$$

where the two first-order terms cancel and $\Delta g = V_{\mathrm{S}}\, g$ turns the rest into $Dg(\Delta\psi - V_{\mathrm{S}}\psi) = -g\mathscr{H}\psi$. For the generator the same vector field appears: $\mathscr{L}(\psi/g) = Dg^{-2}\nabla \cdot (g\nabla\psi - \psi\nabla g) = -g^{-1}\mathscr{H}\psi$. For the quadratic form, with $f = \psi/g$, $\langle \psi, \mathscr{H}\psi\rangle = -\int g f\, g\, \mathscr{L} f \, dx = -D\int f\, \nabla \cdot (p_\infty \nabla f)\, dx = D\int p_\infty |\nabla f|^2\, dx$ after one integration by parts. Polarisation gives the symmetric form $D\int p_\infty \nabla(\varphi/g)\cdot\nabla(\psi/g)\, dx$, and $\mathscr{H} g = 0$ is the display before the theorem. Last, $\partial_t(g\psi) = \mathscr{L}^*(g\psi)$ reads $g\,\partial_t\psi = -g\mathscr{H}\psi$, and $g > 0$. □

The generator, the Fokker–Planck operator and $-\mathscr{H}$ are one operator seen in $L^2(p_\infty)$, $L^2(p_\infty^{-1})$ and $L^2(dx)$, and confinement makes its spectrum discrete.

Theorem 6.5 (Discrete spectrum under confinement)**.** Under the assumption, $\mathscr{H}$ is self-adjoint on its natural domain with compact resolvent, so its spectrum is a sequence of real eigenvalues $0 = \lambda_0 < \lambda_1 \le \lambda_2 \le \ldots \to \infty$ with an orthonormal basis of real eigenfunctions ψ_k of $L^2(\mathbb{R}^d)$, and the eigenvalue 0 is simple with $\psi_0 = \sqrt{p_\infty}$; the generator and the Fokker–Planck operator are unitarily equivalent to $-\mathscr{H}$ [2].

Corollary 6.6 (Eigen-expansion and decay at the spectral gap)**.** Under the assumption, every solution with $p_0/\sqrt{p_\infty} \in L^2$ is $p(t) = \sqrt{p_\infty}\sum_{k\ge 0} e^{-\lambda_k t}\langle \psi_k, p_0/\sqrt{p_\infty}\rangle\, \psi_k$, the $k = 0$ term is p_∞ times the mass, and $\int (p(t) - p_\infty)^2 p_\infty^{-1} dx \le e^{-2\lambda_1 t}\int (p_0 - p_\infty)^2 p_\infty^{-1} dx$ for a probability density p_0; λ_1 is the spectral gap and $1/\lambda_1$ the best constant in the Poincaré inequality for p_∞ written with the Dirichlet form, $\mathrm{Var}_{p_\infty}(f) \le \lambda_1^{-1} D\int |\nabla f|^2 p_\infty\, dx$.

Proof. By the theorem $\psi = p/g$ solves $\partial_t\psi = -\mathscr{H}\psi$, so $\psi(t) = \sum_k e^{-\lambda_k t}\langle \psi_k, \psi(0)\rangle \psi_k$; multiply by g. With $\psi_0 = g$, $\langle \psi_0, p_0/g\rangle = \int p_0$ and the $k = 0$ term is $p_\infty \int p_0$. For a probability density, unitarity and Parseval give

$$\int \frac{(p(t) - p_\infty)^2}{p_\infty}\, dx = \sum_{k\ge 1} e^{-2\lambda_k t}\big|\langle \psi_k, \psi(0)\rangle\big|^2,$$

at most $e^{-2\lambda_1 t}\|\psi(0) - g\|^2$. Last, by min–max $\lambda_1 = \inf\{\langle \psi, \mathscr{H}\psi\rangle/\|\psi\|^2 : \psi \perp g\}$; with $\psi = gf$ the constraint is $\int f p_\infty = 0$, $\|\psi\|^2 = \mathrm{Var}_{p_\infty}(f)$ and $\langle \psi, \mathscr{H}\psi\rangle = D\int |\nabla f|^2 p_\infty$. This infimum is the spectral gap of Chapter

Chapter 8

Nonlinear and Nonlocal Equations

Every equation in the first seven chapters had a drift fixed before the density arrived. A population of neurons does not behave that way: each neuron that fires kicks the others, so the drift depends on how fast the population leaves through its own threshold, and a scheme that gets that rate wrong feeds its error back into the drift. This chapter treats Fokker–Planck equations whose coefficients read the solution: the McKean–Vlasov equation, whose drift contains the convolution of the density with an interaction potential $\mathcal{W}$, the aggregation–diffusion equations around it, and a neuron network whose nonlinearity is a boundary flux.

Chapters 3, 4 (Volume I) and 7 built schemes that keep mass, positivity, the Gibbs density and a decreasing relative entropy; all of them knew the equilibrium in advance. Here the equilibrium solves an equation of its own, it can have several solutions, and the quantity that decreases is a free energy F with an interaction term.

The chapter delivers four schemes, each checked against an exact answer: a self-consistent Scharfetter–Gummel scheme with exact discrete equilibria and a free energy that decreases at every step for attractive interactions, a Lagrangian JKO step in mass coordinates, a radial Keller–Segel scheme, and a chain with a reset jump for the neuron network. Section 8.1 states the equations, §8.2 builds the finite volumes, §8.3 computes the Desai–Zwanzig phase transition, §8.4 the Lagrangian step and the benchmark, §8.5 the Keller–Segel collapse, and §8.6 the neurons.

8.1 McKean–Vlasov Drifts and Their Free Energy

The chains of Chapter 7 had their rates fixed before the first step; here the drift reads the density it transports, through an interaction potential.

Definition 8.1 (McKean–Vlasov equation and interaction potential)**.** Let $U : \mathbb{R}^d \to \mathbb{R}$ be a confining potential, $T > 0$ a temperature and $\mathcal{W} : \mathbb{R}^d \to \mathbb{R}$ an interaction potential: an even function, $\mathcal{W}(-x) = \mathcal{W}(x)$, called attractive where it increases with $|x|$. The McKean–Vlasov SDE is $dX_t = -\nabla U(X_t)\,dt - (\nabla \mathcal{W} * p_t)(X_t)\,dt + \sqrt{2T}\,dW_t$, where p_t is the law of X_t. Its density solves the McKean–Vlasov equation

$$\partial_t p = \nabla \cdot (T \nabla p + p\,\nabla U + p\,\nabla(\mathcal{W} * p)),$$

a Fokker–Planck equation $\partial_t p = \mathscr{L}_p^{*}\, p$ whose drift $b(x,p) = -\nabla U(x) - (\nabla \mathcal{W} * p)(x)$ depends on the unknown density and whose diffusion is $D = T$.

The process is McKean's [35]. The equation follows from Itô's formula as in Chapter 1 (Volume I), with the drift evaluated on the law: for a smooth test function f with compact support,

$$\frac{d}{dt}\int f\, p_t\, dx = \int \left(b(x,p_t)\cdot \nabla f + T\Delta f\right) p_t\, dx,$$

the generator $\mathscr{L}_p$ with the density frozen; two integrations by parts give the equation. It is also the limit of N diffusions each pushed by the average of $-\nabla \mathcal{W}$ over the others [36], an interacting particle system that Chapter 9 simulates [37].

Definition 8.2 (Aggregation–diffusion and nonlinear Fokker–Planck equations)**.** Let $H : [0,\infty) \to \mathbb{R}$ be convex (an internal-energy density), U a confining potential and $\mathcal{W}$ an interaction potential. The aggregation–diffusion equation is $\partial_t p = \nabla \cdot (p\,\nabla \xi)$ with chemical potential $\xi = H'(p) + U + \mathcal{W} * p$, and its free energy is

$$F(p) = \int H(p)\, dx + \int U\, p\, dx + \frac{1}{2}\iint \mathcal{W}(x-y)\, p(x)\, p(y)\, dx\, dy;$$

the choice $H(p) = Tp\log p$ gives the McKean–Vlasov equation. More generally a nonlinear Fokker–Planck equation is $\partial_t p = -\nabla \cdot (b(x,p)\,p) + \frac{1}{2}\nabla \cdot$

$\nabla\cdot(a(x,p)\,p)$, in which the drift or the diffusion depends on the unknown density, locally through $p(t,x)$ or nonlocally through integrals of $p(t,\cdot)$ or a boundary flux.

For $H(p)=Tp\log p$ the two definitions agree. Carrillo, Chertock and Huang built finite volumes on this form [38], and Bailo, Carrillo and Hu made them fully discrete [39].

Theorem 8.3 (The free energy dissipates)**.** Let p be a smooth positive solution of the aggregation–diffusion equation with $\mathcal{W}$ even, on $\mathbb{R}^d$ with enough decay or on a bounded domain with the zero-flux boundary condition $p\,\partial_n\xi=0$. Then $\frac{d}{dt}F(p(t))=-\int p\,|\nabla\xi|^2\,dx\le 0$, with equality exactly when $\nabla\xi=0$ wherever $p>0$; for positive p on a connected domain, exactly when ξ is constant. For $H(p)=Tp\log p$ and $\mathcal{W}=0$ one has $F(p)-F(p_\infty)=T\,H(p\,\|\,p_\infty)$ with $p_\infty\propto e^{-U/T}$, and the identity becomes $\frac{d}{dt}H(p\,\|\,p_\infty)=-T\,I(p\,\|\,p_\infty)$.

Proof. The first variation of F is ξ: for a perturbation εq the internal and potential energies change by $\varepsilon\int(H'(p)+U)q+O(\varepsilon^2)$, and the interaction energy by $\frac{\varepsilon}{2}\iint\mathcal{W}(x-y)(q(x)p(y)+p(x)q(y))+O(\varepsilon^2)=\varepsilon\int(\mathcal{W}*p)\,q+O(\varepsilon^2)$, the two halves being equal because $\mathcal{W}$ is even. Along the flow $q=\partial_t p=\nabla\cdot(p\nabla\xi)$, and one integration by parts gives

$$\frac{d}{dt}F(p(t))=\int\xi\,\nabla\cdot(p\nabla\xi)\,dx=-\int p\,|\nabla\xi|^2\,dx;$$

the boundary term $\int_{\partial\Omega}\xi\,p\,\partial_n\xi\,dS$ vanishes by the zero-flux condition, or on $\mathbb{R}^d$ by the decay. The integrand is nonnegative and vanishes exactly where $p|\nabla\xi|^2=0$. For the second statement let $\mathcal{W}=0$, $H(p)=Tp\log p$ and $p_\infty=e^{-U/T}/Z$. Then $TH(p\,\|\,p_\infty)=T\int p\log p+\int Up+T\log Z=F(p)+T\log Z$ and $F(p_\infty)=-T\log Z$. Finally $\xi=T\log(p/p_\infty)+T-T\log Z$, so $\nabla\xi=T\nabla\log(p/p_\infty)$, the dissipation is $T^2I(p\,\|\,p_\infty)$, and dividing by T gives the last identity. □

With no interaction and $V=U/T$, this F is T times the free energy of Chapter 2 (Volume I), the same functional in units of energy, so Chapter 2's JKO step with $D=T$ minimises $\mathrm{W}_2(p,p_k)^2/(2\Delta t)+F(p)$, as does the static JKO step of §8.4 with the interaction included.

Proposition 8.4 (Stationary states are self-consistent Gibbs densities)**.** For $H(p)=Tp\log p$, a positive stationary solution with zero flux satisfies the self-consistency equation

A grid that keeps mass, positivity and its equilibrium is only the beginning. Escape from a well can take longer than any grid will run; a drift can depend on the density itself; in six dimensions a full grid no longer fits; and a diffusion model generates its samples by running a Fokker–Planck equation backwards.

This second volume takes the structure-keeping schemes of Volume I beyond the grid: spectral methods and the spectral gap that sets an escape rate, every scheme read as a Markov chain whose forward and backward equations agree to round-off, nonlinear and nonlocal equations from phase transitions to chemotaxis and coupled neurons, particle methods and the samplers of diffusion models, and tensor trains and score networks measured against the grid on problems with exact answers. A closing guide matches the features of a problem to a method, with a test suite of exact solutions.

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 13 978-957-82-3462-8
Paperback
9 789578 234628