UNIVERSITY OF CAMBRIDGE

Kinetics of Brownian Transport

Experiment & Simulation

Jannes Gladrow

Cavendish Laboratory
University of Cambridge

Corpus Christi College
July 2019

This dissertation is submitted for the degree of
Doctor of Philosophy

Declaration

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the Preface and specified in the text. It is not substantially the same as any that I have submitted, or, is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. I further state that no substantial part of my thesis has already been submitted, or, is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution except as declared in the Preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

Jannes Gladrow
January 2020

Kinetics of Brownian Transport

Jannes Gladrow

The rate of progress of Brownian processes is not easily quantifiable. An important measure of the ”speed” of Brownian motion is the mean first-passage time (FPT) to a given distance. FPTs exist in various flavours including exit- and transition-path times, which, for instance, can be used to quantify the length of reaction paths in folding transitions in molecules such as DNA. Due to their inherently stochastic nature, measurements of any FPTs require repeated experiments under controlled conditions. In my thesis, I systematically explore FPTs in various contexts using a custom-built automated holographic optical tweezers (HOT) setup. More precisely, I investigate transition- and exit-path-time symmetries in equilibrium systems and demonstrate the breakdown of the symmetry in out-of-equilibrium systems. Experimental data from folding DNA-hairpins show that the principles established on the mesoscale extend well into the molecular regime.

In Kramers escape problem, the reciprocal of the escape rate corresponds to the time of first-passage to leave the initial state. A lower bound for the achievable FPT, e.g. of the reaction coordinate of a folding molecule, therefore corresponds to a speed-limit of the ensemble reaction rate. Using my setup, I show that certain barrier shapes can substantially lower the escape time across the barrier without changing the overall energy balance. This result has deep implications for reaction kinetics, e.g. in protein folding.

Furthermore, I investigate the role of entropic forces in Brownian transport, show that hydrodynamic drag plays a crucial role in Brownian motion in confined systems, and give an experimental realisation of Fick-Jacobs theory.

The thermodynamic applications of HOTs considered here necessitate the creation of fine-tuned optical landscapes, which requires precise phase-retrieval to compute the necessary holograms. In order to address this problem, I explore novel algorithms based on deep conditional generative models and test whether such models can assist in finding holograms for a given desired light distribution. I compare several different models, including conditional generative-adversarial networks and conditional variational autoencoders, which are trained on data sets sampled on the HOT setup. Furthermore, I propose a novel forward-loss-minimising architecture and demonstrate its excellent performance on both validation and artificially-created test data sets.

This thesis is dedicated to my family.

Acknowledgements

I am deeply indebted to my research group, which always maintained an engaging, welcoming, and supportive environment over the last few years. It has been an amazing journey, full of hard work, dedication, curiosity, repeated failure, and eventual success. Without the occasional technical hint or advice, I would have probably given up on most projects. And without the occasional evening in a pub or college, I would have missed out on many great experiences, the memory of which will surely stay with me for many years. I also need to say special thanks to Alex without whom my figures would have remained unintelligible planes of grey concrete.

Of course, I am particularly indebted and grateful to my supervisor, Prof. Keyser, who not only had to endure my skepticism, but in spite of it, offered me to join the nanoTrans network as an early-stage researcher. The academic programme of the network was of high quality and I always felt supported on my journey towards completing this thesis. The bonds of friendship and professional relations that I have established in this network, passively or actively, will surely outlast my time as doctoral student.

With certainty, I can say that I will miss more than the free-of-charge Maxwell-centre coffee, which has never failed to fuel my energy. I look back at this time with a mixture of nostalgia, but also optimism about the future, and am proud and grateful to have been part of our amazing research group.

The research leading to the results presented in this thesis has received funding from the European Unions Horizon 2020 research and innovation program under European Training Network (ETN) Grant No. 674979-NANOTRANS (J.G.). J.G. furthermore acknowledges the support of the Winton Programme for the Physics of Sustainability.

Publications

During the course of this PhD, I contributed to the following publications:

- Particle transport across a channel via an oscillating potential
  Y Tan, **J Gladrow**, UF Keyser, L Dagdug, S Pagliara
  Physical Review E 96(5):052401, 2017.

- Experimental evidence of symmetry breaking of transition-path times
  J Gladrow, M Ribezzi-Crivellari, F Ritort, UF Keyser
  Nature Communications 10(1):55, 2019.

- Optimizing Brownian escape rates by potential shaping
  M Chupeau*, **J Gladrow***, A Chepelianskii, UF Keyser, E Trizac
  Proc. Natl. Acad. Sci. U.S.A., 2019.

- Direct detection of molecular intermediates from first-passage times
  AL Thorneywork, Y Tan, **J Gladrow**, AB Kolomeisky, UF Keyser
  bioRxiv:10.1101/772830, 2019.

- Digital phase-only holography using deep conditional generative models
  J Gladrow arXiv:1911.00904, 2019.

Furthermore, outside of the immediate scope, but within the duration of my PhD I contributed to the following publications:

- Broken detailed balance of filament dynamics in active networks
  J Gladrow, N Fakhri, FC MacKintosh, CF Schmidt, CP Broedersz
  Physical Review Letters 116(24):248301, 2016.

- Nonequilibrium dynamics of probe filaments in actin-myosin networks
  J Gladrow, CP Broedersz, CF Schmidt
  Physical Review E 96(2):022408, 2017.

- Broken detailed balance and non-equilibrium dynamics in living systems: a review
  FS Gnesotto*, F Mura*, **J Gladrow***, CP Broedersz
  Reports on Progress in Physics 81(6):066601, 2018.

*: The authors contributed equally.

Table of contents

Chapter 1

Theoretical Foundations

" Nothing in life is certain except death, taxes and the second law of thermodynamics."
- Seth Lloyd [1]

Any comprehensive quantitative description of biological systems will heavily draw on ideas from thermodynamics.

As a matter of fact, thermodynamics has seen a revival in the past thirty years spurred on by the development of small-scale stochastic thermodynamics, molecular biophysics, and a general search for principles in non-equilibrium systems. A prominent example is the recent quest for a *"thermodynamic niche of life"* in out-of-equilibrium systems [2–5]. In this niche, so the thinking goes, the highly improbable initial formation of reproducing structures that would then go on to evolve into all life (*Abiogenesis*) becomes more probable, more plausible. Gradients in intensive thermodynamic variables, such as temperature [6–8] and proton-concentration [9, 8] have been touted as candidates for this niche or aspects of it. Extreme instances of such gradients can indeed be found in nature, for instance near underwater volcanoes.

Like many ideas in the history of science, the idea of a thermodynamic niche is not entirely novel but rather the implication or extension of previous research, notably that of Ilya Prigogine and his work on dissipative structures and the thermodynamic arrow of time [10, 11]. Ultimately, however, the core paradoxon in this debate, the apparent thermodynamic viability of spontaneous pattern-formation, can be traced back to the beginnings of thermodynamics; it is closely related to early discussions around Loschmidt's paradox and Boltzmann's *H-theorem* [12]. Remnants of this debate resurface here and there in modern thermodynamics and even throughout this thesis, whenever questions of reversibility of thermodynamic processes are considered.

In any case, what seems clear is that living beings are curiously well-ordered entities, which seemingly defy popular interpretations of the second law: Can "*disorder*" really never decrease? Decreases in (local) entropy are, in fact, locally achievable at the expense of energy. It is for this reason that all living matter needs to constantly feed on (chemical) energy in order to be able to orchestrate its internal flows of matter and information and thereby, to quote Schrödinger, evade the decay to the state of maximum entropy - equilibrium [13].

The machinery of life can only hope to maintain this dynamic stability, if all its crucial molecular components, in particular cellular proteins, perform their respective tasks relatively seamlessly and with some reliability. An important example is the process by which proteins are produced in a cell: The ribosome pieces together long chains of amino acids according to a plan encoded in the language of messenger RNA, which in turn is translated from DNA in the genome. Once printed, these chains then need to fold into functional proteins, which happens spontaneously if and only if their functional forms correspond to thermodynamically favourable states. This means that the entire *process* of folding is likely to commence and proceed under many different initial configurations of the amino acid chain and the molecules of the surrounding medium. Due to this fundamental importance of all things molecular in cells, I believe that a thermodynamic description of its constituent microscopic systems is an important contribution to a wider "theory of life" [14].

Classical thermodynamics was originally developed for macroscopic systems that are close enough to a thermodynamic limit, such that microscopic fluctuations could be neglected. In its modern form, however, thermodynamic concepts are successfully applied to systems on the meso- and molecular-scale where fluctuations play a crucial role and have become a subject of interest in its own right [15, 16].

The purely statistical nature of the second law becomes apparent on smaller scales and has been conceptualised in a series of *fluctuation theorems* beginning in the early 1990s and 2000s [17–22]. Generally, fluctuation theorems relate the probability of forward (entropy producing) to backward (entropy consuming) processes. Among other things, this means that any forward process will see its backward counterpart be realised from time to time.

But if even the entropy (production) necessarily fluctuates on small scales, how can for instance molecular motors such as kinesin, myosin or dynein hope to perform any meaningful degree of work? And how do biological entities on all scales deal with the inherent uncertainty in sensing, copying DNA, or any molecular process?

While the "*how*" of molecular motors from a thermodynamics point-of-view is usually answered by pointing to power-stroke or ratchet-like mechanisms [23–25, 21], the question of "*how much*" work such motors can carry out with some certainty is subject of recent research activity. Based on large-deviation theory, a number of fundamental uncertainty relations have been derived for non-equilibrium processes [26–28], that seek to answer this question.

Notwithstanding the fluxes of material and energy, every process that involves information, ranging from cell-to-cell communication to error-correction of DNA-to-RNA or RNA-to-protein coding will involve a degree of thermodynamics and its extension in information theory [29–33]. A famous and early example is John Hopfields influential paper on kinetic proofreading [34], where thermodynamics of information lead to genuine biological insights.

Due to these deep connections into biophysics and the theory of information, stochastic thermodynamics seems very much alive as a research discipline.

This thesis is concerned with some of the thermodynamic concepts underlying first-passage times. First-passage times are an important measure of the "speed" of Brownian motion. For instance, one could ask how long it takes for a neurotransmitter to cross the synaptic cleft separating pre- and post-synaptic neurons. Since Brownian motion is a stochastic process, this first-passage time will be of stochastic nature too. It is therefore important to utilise some statistical machinery: One could for instance measure the mean first-passage time of crossing or even the full distribution of times. In any case, repeated experimental trials are necessary to collect enough data on this variable. This, in turn, necessitates an experimental setup, which is able to repeat in sequence or parallel an experiment 1,000 or even 1,000,000 times under the same conditions.

In my thesis, I construct a microscope equipped with an automated holographic optical tweezers setup, which enabled me to study first-passage times of Brownian particles in great detail with a high degree of control over relevant conditions.

The system is endowed with an autofocus routine, online particle localization routines, and an efficient holographic engine, which allows for active feedback control or the human operator to shape the distribution of light intensity in the microscope almost in realtime.

Even though I worked with different techniques, the bulk of this thesis has been carried out on the holographic optical tweezers setup, so that this document is firmly focused on holography and the physical theory used to describe Brownian motion - stochastic thermodynamics.

1.1 Stochastic thermodynamics

The idea of stochastic dynamics has its roots in the study of Brownian motion, named after its discoverer, Robert Brown. In 1828, Brown, a botanicist, studied the erratic movements of tiny objects (pollen) immersed in a fluid under a microscope and gave a phenomenological description of Brownian motion. It fell to a mathematical economist named Louis Bachelier to provide a comprehensive mathematical description of the time-evolution of the probability density of Brownian random walks in 1900 [35]. Incidentally, Bachelier also laid the groundwork for modern theories of financial markets, which (still) seem to be making heavy use of stochastic processes to manage risk and predict returns [36, 37]. I find it noteworthy that Bachelier also derived an early version of what would later become known as the Chapman-Kolmogorov equation.

Five years later, Albert Einstein derived an equation, which again, governed the time-evolution of the density but this time explicitly in the context of diffusion [38]. Einstein is also credited with being the first to provide an explicit equation of the diffusion constant of a particle suspended in a liquid $D = k_B T/\gamma$, with k_B denoting Boltzmann's constant and γ denoting the friction coefficient of the particle.

The polish physicist Marian Smoluchowski had independently arrived at the same equation of D and published his results in 1906, such that the equation became known as Einstein-Smoluchowski relation, in what only superficially appears to be an example of fair credit assignment. In fact, a physicist named William Sutherland had arrived at this equation already in 1904 [39].

A few years later, in 1908, Paul Langevin published a paper that sought to describe Brownian motion using a radically new concept - a stochastic differential equation. This equation was not designed to be solvable in the classical sense, but it can nevertheless lead to insights into the statistics of the underlying process and dependent variables. Even though, perhaps the full extent of the discovery was not clear at the time, Langevin's concept would inspire a host of mathematical, economical and physical research activity even more than 100 years later.

Generally speaking there are two complementary and equivalent descriptions in stochastic dynamics: (1) Langevin's microscopic description based on equations of motions that seek to describe the statistical properties of trajectories, and (2) Einstein's probabilistic description of the evolution of an *ensemble* of the system in question.

1.1.1 Langevin equation

The key idea in the Langevin picture of stochastic motion is to replace the set of equations of motion of a very high-dimensional, chaotic, yet deterministic system by a few or even a single stochastic equation of motion [40]. Of course, this is particularly useful whenever only a subset of variables can be tracked.

For instance, instead of tracking the momenta $\vec{p}_i$ and positions $\vec{x}_i$ of all water molecules in a bath, one replaces the collective interaction with a subsystem of interest, such as a colloid, by a stochastic force term $\vec{\xi}(t)$. It is this *ignorance* towards the neglected degrees of freedom that is a major source of entropy. The relation between ignorance and entropic effects will become relevant in Sec. 3.3.1.

In any case, the statistical properties of this stochastic force $\vec{\xi}(t)$ need to be specified, since they decide if the equations describe the entire system in a statistically correct way. Many physical systems that are studied in this fashion fulfil the condition of *ergodicity*, that is, averages can either be taken over an infinite number of different initial conditions (ensemble) or over a single, but infinitely long trajectory. In such cases, ensemble $\langle\cdot\rangle_{\mathrm{ens}}$ and temporal averages $\langle\cdot\rangle_{\mathrm{T}}$ are interchangeable. All systems considered in this thesis are ergodic and hence I will not specify the averaging procedure.

Let the subsystem of interest have a mass m and the surrounding fluid a temperature T. Linear response theory provides a framework in which many of the quantities of interest in this thesis can be derived. The central assumption here is the existence of a response function $\chi(t)$ that fulfils

$$\langle x_i(t)\rangle = \int_{-\infty}^{t} \mathrm{d}t'\, \chi_{i,j}\left(t'-t\right) f_j(t') \tag{1.1}$$

where $f_j(t)$ represents a force that perturbs the system at time t.

The statement in Eq. (1.1) is often written in the Fourier domain using the susceptibility $\hat{\chi}$, $\hat{x}_i(\omega) = \hat{\chi}_{i,j}(\omega)\hat{f}_j(\omega)$.

A cautionary note in the beginning: I assume here the subsystem to be far away from any obstacles or boundaries, a condition that is blatantly violated in microfluidic channels, that are investigated in this thesis. However, under these simplifying assumptions, effects related to the hydrodynamics of the surrounding fluid can be conveniently summarised in a single constant, the friction coefficient. Much research activity has been devoted to investigating effects specifically arising from the breakdown of this assumption [41–45].

For the sake of generality, I consider a general, possibly non-Newtonian fluid with a shear modulus $\hat{G}(\omega)$, such as starch solution or cellular cytoplasm [46]. From a dynamics

point-of-view, the crucial property of non-Newtonian fluids is the path-dependence of movements and velocities of objects. This path dependence is summarised in a viscoelastic kernel $\alpha(t)$, which is convolved in time with the momentum variable of the degree of freedom under consideration (see e.g. [47]).

A generalised Stokes formula relates the Fourier-transformed shear modulus $\hat{G}$ to the Fourier-transformed viscoelastic kernel $\hat{\alpha}$ of the medium, [48]

$$\hat{\alpha}(\omega) = \mathscr{A}\frac{\hat{G}(\omega)}{-i\omega}, \tag{1.2}$$

where $\mathscr{A}$ denotes a geometric factor. The same factor appears in the drag coefficient $\gamma = \mathscr{A}\eta$ of the subsystem. In the case of a sphere of radius R in a purely Newtonian medium, Eq. (1.2) simplifies to Stokes' formula $\gamma = 6\pi R\eta$ with η representing the viscosity of the medium. This follows from the fact, that in a purely Newtonian medium, the shear modulus and hence the kernel is purely imaginary $\hat{\alpha}(\omega) = i\mathscr{A}\omega\eta$ since a Newtonian liquid lacks any elastic response.

Within the boundaries of the linear response framework, the stochastic version of Newton's equations of motion of the subsystem read [49, 50]

$$\frac{\mathrm{d}\vec{p}}{\mathrm{d}t}(t) = -\frac{1}{m}\int_{-\infty}^{t}\mathrm{d}t'\,\alpha(t'-t)\vec{p}(t') + \vec{f}_{\mathrm{ext}}(\vec{x}(t),t) + \vec{\xi}(t) \tag{1.3}$$

$$m\frac{\mathrm{d}\vec{x}}{\mathrm{d}t}(t) = \vec{p}(t) \tag{1.4}$$

where $\vec{x}$ and $\vec{p}$ denote the position and momentum of the centre of mass of the subsystem. The stochastic term $\vec{\xi}(t)$ represents the collective effect of thermally-induced collisions. In the temporal Fourier domain, the same set of equations read

$$i\omega\,\hat{\vec{p}}(\omega) = -\frac{1}{m}\hat{\alpha}(\omega)\hat{\vec{p}}(\omega) + \hat{\vec{f}}_{\mathrm{ext}}(\vec{x},\omega) + \hat{\vec{\xi}}(\omega) \tag{1.5}$$

$$mi\omega\,\hat{\vec{x}}(t) = \hat{\vec{p}}(\omega). \tag{1.6}$$

The susceptibility can then be found by combining Eqs. (1.5) and (1.6) and reads $\hat{\chi}(\omega) = \left(m\omega^2 + i\hat{\alpha}(\omega)\omega\right)^{-1}$, which I use further down. In later chapters, the external force term that appears in Eq. (1.3), takes up the role of the force exerted by the optical tweezers.

Since the stochastic term $\vec{\xi}(t)$ seeks to describe the collective effect of a great number of collision events, its distribution must be close to a normal distribution on all timescales that are relevant in this thesis (see discussion in Sec. 1.1.1). Thus, once the mean $\langle\vec{\xi}\rangle = \vec{\mu}$

and covariance matrix $\langle \vec{\xi}\vec{\xi}^T \rangle = \mathbf{S}$ of this distribution are found, the statistical properties of $\vec{\xi}(t) \sim \mathcal{N}(\vec{\mu}, \mathbf{S})$ would be completely defined.

Interestingly, the scale of the fluctuations $\vec{\xi}$ is determined by the dissipative properties of the surrounding medium. The *Fluctuation-Dissipation Theorem* (FDT) provides an elegant and simple relation between the medium response and the correlator of the Fourier transformation of the stochastic force, $\hat{\vec{\xi}}(t)$[51, 49]

$$\langle \hat{\vec{\xi}}_i(\omega) \hat{\vec{\xi}}_j^{\dagger}(\omega) \rangle = \frac{2k_B T}{\omega} \delta_{ij} \mathrm{Im}\left[\hat{\alpha}(\omega)\right]. \tag{1.7}$$

where the † symbol marks the complex conjugate. Such a relation always holds in a steady state and within the realm of applicability of linear response theory. While FDT's certainly hold in equilibrium, they may also hold in non-equilibrium steady states, albeit, potentially only for a different choice of variables [52].

Equation (1.7) implies another FDT, which determines the scale of fluctuations in the position variable of the subsystem in relation to its susceptibility (see [49]),

$$\langle \hat{\vec{x}}_i(\omega) \hat{\vec{x}}_j^{\dagger}(\omega) \rangle = \frac{2k_B T}{\omega} \delta_{ij} \mathrm{Im}\left[\hat{\chi}(\omega)\right] \tag{1.8}$$

Accordingly, for a purely viscous medium, the correlator of the noise reads

$$\langle \hat{\vec{\xi}}_i(\omega) \hat{\vec{\xi}}_j^{\dagger}(\omega) \rangle = 2k_B T \mathscr{A}\eta \delta_{ij} \tag{1.9}$$

which implies

$$\langle \vec{\xi}_i(t) \vec{\xi}_j(0) \rangle = 2k_B T \mathscr{A}\eta \delta_{ij} \delta(t). \tag{1.10}$$

As for the correlator of the position variables $\vec{x}$ in a Newtownian fluid, Eq. (1.8) implies

$$\langle \vec{x}_i(t) \vec{x}_j(0) \rangle = \frac{2k_B T}{\mathscr{A}^2\eta^2} \delta_{ij} \left(e^{-\frac{\mathscr{A}\eta}{m}t} + \mathscr{A}\eta t \right) \overset{t \gg 1}{\to} \frac{2k_B T}{\mathscr{A}\eta} t \delta_{ij} \tag{1.11}$$

A comparison of the later term with the phenomenological definition of diffusion $\langle \vec{x}(t)\vec{x}(0) \rangle = 6Dt$ yields the Einstein-Smoluchowski relation discussed earlier,

$$D = \frac{k_B T}{\gamma} \tag{1.12}$$

with $\gamma = \mathscr{A}\eta$ as introduced before.

The phenomenological definition of the diffusion constant used above derives from a more general description that involves the velocity autocorrelation function (VACF) (see e.g. [53, 54])

$$\mathbf{D} \equiv \int_0^\infty \mathrm{d}t' \langle \vec{v}(t') \vec{v}^T(0) \rangle. \tag{1.13}$$

Crucially, the VACF fulfils a time-reversal symmetry (see e.g. [55]), meaning that

$$\langle \vec{v}(t') \vec{v}^T(t'') \rangle = \langle \vec{v}(t'') \vec{v}^T(t') \rangle. \tag{1.14}$$

This enables one to write

$$\begin{aligned} \langle x^2(t) \rangle &= \int_0^t \mathrm{d}t' \int_0^t \mathrm{d}t'' \langle v_x(t') v_x(t'') \rangle \\ &= 2 \int_0^t \mathrm{d}t' \int_0^{t'} \mathrm{d}t'' \langle v_x(t') v_x(t'') \rangle \end{aligned}$$

which implies $\partial_t \langle x^2(t) \rangle = 2 \int_0^t \mathrm{d}t' \langle v_x(t') v_x(0) \rangle$. For δ-correlated velocities, I retrieve $\langle x^2(t) \rangle = 2tD_x$, which leads back to the result in Eq. (1.11). Such time-reversal symmetries can breakdown under certain conditions in out-of-equilibrium systems. I refer the interested reader to my recent paper [56].

The VACF is related to the time-correlation of the random noise, which leads to another form of the above FDT [57]

$$\mathbf{D} = \frac{1}{2\mathscr{A}\eta} \int_{-\infty}^\infty \mathrm{d}t' \langle \vec{\xi}(t') \vec{\xi}^T(t') \rangle. \tag{1.15}$$

And, as a final remark, the above equation can also be written in matrix form, which is helpful in systems, where each degree of freedom experiences, for instance, a different temperature. In such a case, it may be a good idea to absorb the deviation of $\vec{\xi}$ in Eq. (1.10) into a *forcing* matrix $\mathbf{F}$, such that $\vec{\xi}(t) = \mathbf{F}\vec{\xi}_0$ with $\vec{\xi}_0$ having unit variance. The diffusion matrix $\mathbf{D}$ can then be written as

$$\mathbf{D} = \frac{1}{2} \mathbf{F}\mathbf{F}^T. \tag{1.16}$$

1.1.2 Overdamped limit

The VACF lends itself as a toe-hold to another crucial concept: overdamped dynamics. In this thesis, I mostly deal with dynamics of colloidal particles, which are on the order of 100 nm to 2 µm. In the following, I want to briefly discuss the relevant timescales for colloidal dynamics in aqueous (and thus purely viscous) media.

While the collision rate of water molecules with the colloid is on the order of 1 THz [58, 59], the inertia of the colloid itself decays on a slower time scale, $\tau_p = m/\gamma$. Disregarding hydrodynamic details, such as hydrodynamic resonances, the VACF of a colloid in water decays as a single exponential over $\tau_p \sim 100$ ns [60].

However, the relevant time-scales studied in this thesis are seconds, if not minutes. The relatively fast decay of the VACF effectively means that the velocity $\vec{v}(t)$ is so erratic, that it is not relevant to the slow dynamics of $\vec{x}$ considered here. This can be shown in a more formal way by considering the solution to Eq. (1.3) for $\hat{\alpha}(\omega) = i\gamma\omega$

$$v_x(t) = v_0 e^{-\frac{t}{\tau_p}} + \frac{\sqrt{2D_x}}{\tau_p}\int_0^t \mathrm{d}t'\, e^{\frac{(t-t')}{\tau_p}}\xi_x(t') \tag{1.17}$$

The equation shows that the velocity behaves like an Ornstein-Uhlenbeck stochastic process, with an internal decorrelation time τ_p. In the limit of τ_p approaching zero, the VACF will therefore converge to a white-noise correlator [37], that is

$$\langle v_x(t)v_x(0)\rangle = \frac{D_x}{\tau_p}e^{-\frac{|t|}{\tau_p}} \overset{\tau_p\to 0}{\to} D\delta(t). \tag{1.18}$$

A δ-correlation indicates that the underlying variable is a white-noise process. The Langevin equation can therefore be rewritten as $\dot{x}(t) = \mathrm{d}/\mathrm{d}t\left(\int_0^t \mathrm{d}t'\, v(t')\right) \approx \sqrt{2D_x}\xi_x(t)$ with $\xi_x(t)$ denoting unit-variance Gaussian white-noise, eliminating the velocity.

Integrating out the fast variable has a topological consequence: The dynamics of the reduced system are one-dimensional and of first-order only. The system is therefore bereft of its ability to oscillate. Instead, its only remaining type of (average) dynamics is a creeping motion, a drift, driven by the force term f. Such a change in topology can have thermodynamical consequences that I explore in chapter 3.1.3.

I also refer the interested reader to my papers on broken detailed balance in polymer dynamics [61, 56], where the question of topology of phase-space plays an important role.

The overdamped regime is relevant in much of biophysics and describes even dynamics on the molecular scale rather well. In the following I give a brief overview of small sample of the relevant literature.

Stochastic dynamics in biophysics

Stochastic dynamics, in particular the overdamped regime, are prevalent in biophysics and of importance to the topics discussed in this thesis and hence deserve a short discussion.

For instance, single-molecule measurements of folding dynamics of proteins and DNA are typically characterized using overdamped stochastic dynamics [62–67]. Results from molecular dynamics (MD) simulations give an impression of the magnitude of VACF decay times in molecular systems: velocities of individual atoms in the Lysine side chain of *Ribonuclease A* (RNase A), a much studied protein [68], decorrelate over a range of 0.1 ps to $\sim$ 300 ps [69]. The timescale of decorrelation is influenced by local friction and, importantly, the timescale of escape from local potential minima. Any motion that proceeds on a (much) longer timescale than $\sim$ 100 ps is found to be dominated by timescales set by the free energy landscape.

On the experimental side, Förster Resonance Energy Transfer (FRET) has been extensively used to study barrier-crossing dynamics of the folding process [62, 64–66]. The technique allows for a quantification of distance and relative orientation of two fluorophores on the scale of roughly 1-10 nm. If one attaches these fluorophores onto parts of a protein (or DNA) that are far apart in the unfolded, but close in the folded state, the FRET signal becomes a one-dimensional reporter on the progress of folding.

With this technique, William A. Eaton and his group studied the folding rate and transition-path time of several proteins in detail. Overall, folding rates derived from overdamped dynamics seem to fit the data (Kramers rates, see Sec. 1.2.3). Time-resolved force spectroscopy based on optical tweezers has been successfully applied in studies of the folding pathways of proteins and DNA [63, 70–74]. Measurements of ribosomal stepping times along RNA have led to insights into the molecular mechanics of gene translation [75]. The same technique has also been used to show a transition-path time symmetry in equilibrium folding and unfolding transitions of DNA-hairpins [72]. Recently, local velocities along folding trajectories of DNA have been measured with high resolution of the folding coordinate, which shed light on the frequency of recrossing events in relation to all free-energy barrier crossings [76].

The focus on transition-path times and folding rates in these studies derives from a desire to study transition paths and the thermodynamics of folding. FRET and other techniques, however, do not capture the full phase space but only a projection along

some reaction coordinate. Crucially, this does not imply that the folding process proceeds along a single pathway in phase space. The question of pathway *multiplicity* in molecular folding has been recently debated in the literature [67, 77–81].

1.1.3 Fokker-Planck equation

The Fokker-Planck equation (FPE) describes the time evolution of the probability mass of an ensemble of points in phase space of the subsystem (here) under the influence of Gaussian white noise. Depending on whether fast-decorrelating variables of the subsystem, such as velocities, have been integrated out, the FPE is sometimes referred to as Smoluchowski equation. I adopt here a naming convention used in mathematics which does not make this distinction and always refers to probability-governing equations as FPE's.

The idea to describe the dynamics of a probability instead of individual trajectories did not originate in stochastic dynamics. Liouville's equation in classical mechanics provides a similar probabilistic description of the time evolution of an ensemble of initial conditions under given equations of motion. In fact, following Robert Zwanzig [82], the FPE can be directly obtained by adding a Langevin-style noise term to Liouville's equation and subsequent averaging.

The derivation is unusual in so far as it involves a probability distribution of a probability distribution. The first step is to write down Liouville's equation with an additional noise term. For the sake of simplicity, I will assume here white-noise, such that $\langle \vec{\xi}(t)\vec{\xi}(t')^T \rangle = 2k_B T\gamma\delta(t-t')\mathbf{1}$. Liouville's equation is essentially a conservation law for the probability mass $\psi(\vec{x})$ in the zero-noise limit. In this limit, the probabilistic aspect of the time evolution arises solely from an initial distribution of initial conditions in phase space. Liouville's equation reads

$$\frac{\partial \psi}{\partial t}(\vec{x},t) = -\vec{\nabla}\cdot\left(\vec{A}(\vec{x},t)\psi(\vec{x},t)\right). \tag{1.19}$$

with $\vec{A}$ denoting the phase-space velocity, e.g. for ballistic movement in some potential U, it reads $\vec{A} = \left(-\vec{\nabla}U - \gamma\vec{p}/m,\ \vec{p}/m\right)^T$. I should clarify that $\vec{x}$ in this subchapter denotes positions in phase-space.

Now, Zwanzig adds the phase-space noise term $\vec{\xi}_{\mathrm{PS}} = \left(\vec{\xi}(t), \vec{0}\right)$, which appears as a velocity

$$\frac{\partial \psi}{\partial t}(\vec{x},t) = -\vec{\nabla}\cdot\left(\left(\vec{A}(\vec{x},t) + \vec{\xi}_{\mathrm{PS}}(t)\right)\psi(\vec{x},t)\right) \tag{1.20}$$

$$= -\mathbf{L}\psi(\vec{x},t) - \vec{\nabla}\cdot\left(\vec{\xi}_{\mathrm{PS}}(t)\psi(\vec{x},t)\right). \tag{1.21}$$

where $\mathbf{L}$ is the linear Liouville operator $\mathbf{L}u = \vec{\nabla}\cdot\left(\vec{A}_u u\right)$ with $\vec{A}_u$ denoting the phase-space velocities associated with u. The latter form of the equation indicates how the equation can be solved, since it's form is reminiscent of inhomogeneous linear differential equations, $\frac{\mathrm{d}x}{\mathrm{d}t} = -ax(t) + b(t)$. Such equations are solved using a variation of constants, i.e. $x(t) = e^{at}x_0 + \int_0^t \mathrm{d}t'\, e^{a(t'-t)b(t')}$. Applied to Eq. (1.21), this leads to

$$\psi(\vec{x},t) = e^{-\mathbf{L}t}\psi(\vec{x},0) - \int_0^t \mathrm{d}t' \left[e^{\mathbf{L}(t-t')}\vec{\nabla}\cdot\left(\vec{\xi}_{\mathrm{PS}}(t')\psi(\vec{x},t')\right)\right]. \tag{1.22}$$

Zwanzig then reinserts this equation into Eq.(1.21), and, after some algebra, takes the average $\langle\cdot\rangle$ over the phase-space noise $\vec{\xi}_{\mathrm{PS}}$, which leads to

$$\frac{\partial\langle\psi\rangle}{\partial t}(\vec{x},t) = -\vec{\nabla}\cdot\left(\vec{A}(\vec{x},t)\langle\psi\rangle(\vec{x},t)\right) + \vec{\nabla}\cdot\mathbf{D}\vec{\nabla}\langle\psi\rangle(\vec{x},t). \tag{1.23}$$

The diffusion matrix here follows from Eq. (1.16) with $\langle\vec{\xi}_{\mathrm{PS}}(t)\vec{\xi}^{T}_{\mathrm{PS}}(t')\rangle = 2\mathbf{D}\delta(t-t')$. A simple redefinition of $\langle\psi\rangle \equiv \rho$ then leads to the Fokker-Planck equation,

$$\frac{\partial\rho}{\partial t}(\vec{x},t) = -\vec{\nabla}\cdot\left(\vec{A}(\vec{x},t)\rho(\vec{x},t) - \mathbf{D}\vec{\nabla}\rho(\vec{x},t)\right) \tag{1.24}$$

$$= -\vec{\nabla}\cdot\vec{j}(\vec{x},t). \tag{1.25}$$

with $\vec{j}$ denoting the probability current. The form of Eq. (1.25) shows that the FPE really is a conservation equation.

Importantly, for initial conditions $\vec{x}(t_0) = \vec{x}_0$, the solution to the FPE is a conditional probability density $\rho(\vec{x},t|\vec{x}_0,t_0)$ which describes transitions $(\vec{x}_0,t_0) \to (\vec{x},t)$ rather than states $\vec{x}$.

The FPE can be solved using various methods. For constant forces $\mathbf{A}(x,t) = \mathbf{A}$ and diffusion matrices $\mathbf{D}$, the solution is particularly simple and of use in later chapters. Under

such conditions and initial conditions $\rho(\vec{x}_1, t_0|\vec{x}_0, t_0) = \delta(\vec{x}_1 - \vec{x}_0)$ it reads [37]

$$\rho(\vec{x}_1, t_0 + \Delta t|\vec{x}_0, t_0) = \mathcal{N}^{-1}\mathrm{Exp}\left[-\frac{(\vec{x}_1 - \vec{x}_0 - \mathbf{A}\Delta t)^T \mathbf{D}^{-1}(\vec{x}_1 - \vec{x}_0 - \mathbf{A}\Delta t)}{2\Delta t}\right] \tag{1.26}$$

with a normalisation factor $\mathcal{N} = \sqrt{(4\pi\Delta t)^d \det[\mathbf{D}]}$.

Detailed balance and microscopic irreversibility

The probability density behaves as an incompressible fluid, which has important consequences. The probability current density $\vec{j}$ of an overdamped system in steady-state ($\partial_t \rho = 0$) must either be purely rotational, that is, $\mathrm{div}(\vec{j}) = \vec{\nabla} \cdot \vec{j} = 0$ or zero everywhere $\vec{j} = \vec{0}$. The former state is referred to as non-equilibrium steady state (NESS).

For the system to be truly in equilibrium, it is therefore not sufficient to demand $\partial_t \rho = 0$. The system is only truly in equilibrium in complete absence of any probability currents, $\vec{j} = \vec{0}$, a condition known as *detailed balance*.

It is worthwhile to consider the consequences of detailed balance or the absence thereof for observable dynamics of the subsystem (e.g. the position of a colloid). A flow of probability between regions of phase space means that the system will on average transition between these regions more frequently into the direction of the flow. A probability flow implies a drift motion, a tendency to move into the direction of the flow.

By contrast, in the absence of probability flow, the transition rate between any two states $\vec{x}_0$ and $\vec{x}_1$ should cancel out, the net transition rate must be zero. For two states with steady-state probabilities $p(\vec{x}_0)$ and $p(\vec{x}_1)$, the transition $(\vec{r}_0, t_0) \to (\vec{r}_1, t_1)$ is then equally likely as its time-reversed counterpart $(\vec{r}_1, t_0) \to (\vec{r}_0, t_1)$, written in forward time. I tacitly assume here, that all variables considered in the overdamped system behave as even variables in the sense that their signs do not flip under time reversal. In such circumstances, the detailed balance condition expressed in terms of transition probabilities reads[37]

$$p(\vec{x}_1, t_1|\vec{x}_0, t_0)\, p(\vec{x}_0) = p(\vec{x}_0, t_0|\vec{x}_1, t_1)\, p(\vec{x}_1). \tag{1.27}$$

In systems, where the underlying equations of motion are time-symmetric (Newton's laws are time-symmetric), a property known as *microscopic reversibility* guarantees that every forward transition is associated with a backward transition. In equilibrium, the rates of these two processes must be equal (detailed balance).

In such systems it may be difficult for observers to infer the direction of time from microscopic trajectories of observed subsystems. For instance, if one forgot to save the time-direction of microscope movies of colloids undergoing Brownian motion, it would be

impossible, in equilibrium, to find that direction. If detailed balance is violated, one must gather sufficient information about the probabilities of forward and backward processes in order to decide which is which. In fact, for steady-state thermodynamic processes, the "amount of violation" of detailed balance [20, 83, 84] and, interestingly, the amount of data it takes to decide the direction of time [85] are related to the amount of entropy produced.

Example - Overdamped harmonic oscillator driven by coloured noise

In order to give an example of the theory laid out above, I solve the FPE for the one-dimensional, overdamped harmonic oscillator driven by coloured noise $\nu(t)$. For the sake of simplicity, I set $k_B = \gamma = 1$ and work in arbitrary units.

The coloured noise is here implemented as an Ornstein-Uhlenbeck process with a relaxation time τ. Later in this thesis, I will use a different coloured-noise process known as telegraph process. However, the two processes share the form of their steady-state correlator $\langle \nu(t)\nu(0)\rangle \propto e^{-|t|/\tau}$, such that the effects of Ornstein-Uhlenbeck and telegraph processes as driving forces are often phenomenologically similar.

The equations of motion corresponding to the scenario sketched in Fig. 1.1**a** read

$$\begin{pmatrix} \dot{x}(t) \\ \dot{\nu}(t) \end{pmatrix} = \begin{pmatrix} -k & \varepsilon \\ 0 & -\frac{1}{\tau} \end{pmatrix} \begin{pmatrix} x(t) \\ \nu(t) \end{pmatrix} + \begin{pmatrix} \sqrt{2T} & 0 \\ 0 & \sqrt{2T_\nu/\tau} \end{pmatrix} \begin{pmatrix} \xi_x(t) \\ \xi_\nu(t) \end{pmatrix} \tag{1.28}$$

where the first matrix is now abbreviated as **A** and the second as **F**. In order to ensure that the variance of the noise process ν does not depend on τ, I rescale the temperature of the noise, T_ν, by τ^{-1}. The diffusion matrix in this example can be obtained using Eq. (1.16).

I am here only interested in the steady-state. In this case, the phase-space density at some point $\vec{r} = (x, \nu)^T$ is time-constant and reads

$$\rho(\vec{r}) = \frac{1}{2\pi\sqrt{\det[\mathbf{C}]}} e^{-\frac{1}{2}\vec{r}^T\mathbf{C}^{-1}\vec{r}}. \tag{1.29}$$

The correlation matrix **C** in Eq. (1.29) can be obtained by solving the following Lyapunov-type differential equation [86]

$$\frac{\mathrm{d}\mathbf{C}}{\mathrm{d}t}(t) = \mathbf{A}\mathbf{C}(t) + \mathbf{C}(t)\mathbf{A}^T + 2\mathbf{D}. \tag{1.30}$$

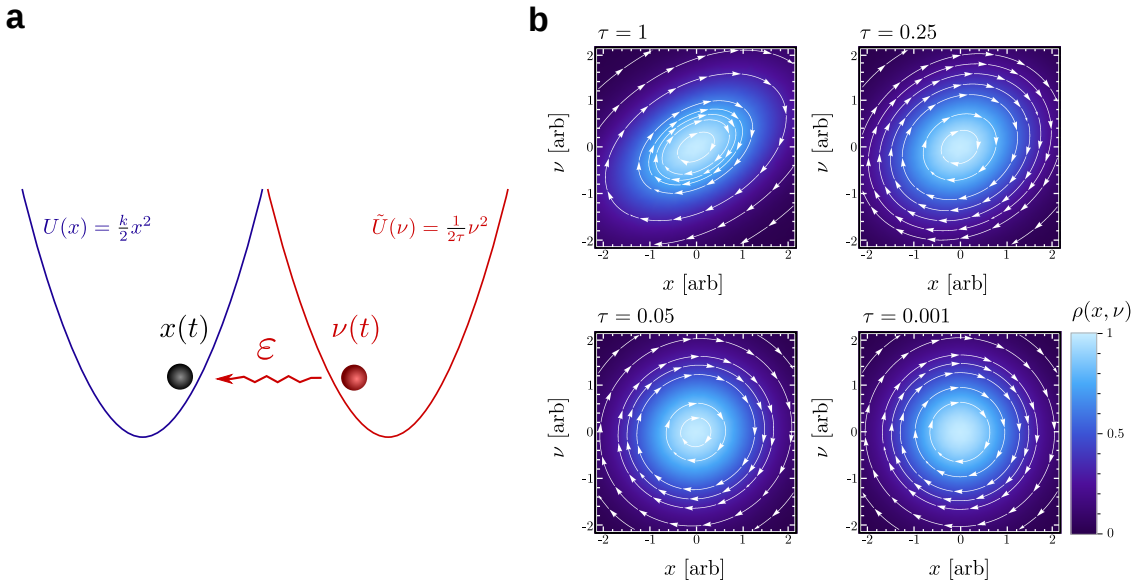

Fig. 1.1 Harmonic oscillator with a one-way coupling to an Ornstein-Uhlenbeck process. **a** Sketch of the potentials and variables of the system $x(t)$ and the noise $\nu(t)$. **b** Steady-state phase-space density $\rho(x,\nu)$ and current $\vec{j}(x,\nu)$ for different values τ. The vector plot is normalised, the magnitude of $\vec{j}$ is not shown. In the plots, I set $k = T = T_\nu = 1$.

For $t \to \infty$, the equation simplifies to $\mathbf{AC} + \mathbf{CA}^T = -2\mathbf{D}$, which, when solved yields the steady-state correlation matrix $\mathbf{C}$,

$$\mathbf{C} = \begin{pmatrix} \frac{T+kT\tau+T_\nu\varepsilon^2\tau}{k+k^2\tau} & \frac{T_\nu\varepsilon\tau}{1+k\tau} \\ \frac{T_\nu\varepsilon\tau}{1+k\tau} & T_\nu \end{pmatrix} \xrightarrow{\tau\to 0} \begin{pmatrix} \frac{T}{k} & 0 \\ 0 & T_\nu \end{pmatrix}. \tag{1.31}$$

In panel **b**, I show the effect of sending τ to zero, reminiscent of the overdamped limit considered in Sec. 1.1.2. As the limit in Eq. (1.31) shows, the offdiagonal elements of the correlation matrix approach zero for vanishing τ. This means that the phase-space density looses its tilt as shown in Fig. 1.1**b**: The dynamics of the system state $x(t)$ becomes independent of the internal noise state $\nu(t)$ even though it is still driven by it.

The asymmetric coupling of the noise to the system results in a breakdown of detailed balance and a finite, purely rotational probability current $\vec{j}(\vec{r}) = (j_x(\vec{r}), j_\nu(\vec{r}))^T$. The current can be obtained from $\vec{j}(\vec{r}) = \mathbf{U}\vec{r}\rho(\vec{r})$ using the frequency matrix $\mathbf{U}$, which is given by $\mathbf{U} = \mathbf{A} + \mathbf{DC}^{-1}$ [86]. The imaginary part of the positive eigenvalue of this matrix,

$$\omega = \frac{\varepsilon\sqrt{kT_\nu}}{\sqrt{T + 2kT\tau + T_\nu\varepsilon^2\tau + k^2T\tau^2}} \xrightarrow{\tau\to 0} \varepsilon\sqrt{\frac{kT_\nu}{T}}, \tag{1.32}$$

corresponds to the average cycling frequencies of a point $\vec{r}$ in phase space in the Langevin picture [86, 61]. Here, ω serves as a measure of strength of the current shown in Fig. 1.1**b**. The cycling frequency increases monotonically with decreasing τ.

1.2 First-passage processes

How long does it take for a stochastic process to (directly) move from a point $\vec{x}_a$ to a point $\vec{x}_b$? What are the odds of the process reaching $\vec{x}_b$ before touching some other point $\vec{x}_c$? First-passage processes offer a systematic and sometimes mathematically tractable approach to such questions. First-passage-related phenomena have received increased theoretical attention in recent years, for instance in the context of molecular folding [87–89, 74], nanopore translocations [90–92], and stepping behaviour of molecular motors [23, 93, 25].

In the following, I want to briefly introduce a few first-passage concepts that provide valuable context to the experiments described in this thesis.

If a stochastic process visits a particular point $\vec{x}$ at time t for the first time since the beginning of observation at some previous time t_0, this time t is known as the first-passage time (FPT) of the process at $\vec{x}$.

More generally, if one is interested in the probability of a future first-passage event at time T and place $\vec{x}$, one can use the backward Fokker-Planck equation (bFPE) to develop the probability density of the system in forward time towards this event. A solution to the bFPE is a measure of how likely this future first-passage event is given the current state at time $t < T$. From a mathematical point-of-view, the bFPE-operator is the adjoint-operator to the (forward) FPE and has a similar form [37],

$$\frac{\partial \rho}{\partial t}\left(\vec{x}, T|\vec{x}', t\right) = -\mathbf{A}\left(\vec{x}'\right)\vec{\nabla}_{\vec{x}'}\rho\left(\vec{x}, T|\vec{x}', t\right) - \left(\mathbf{D}\cdot\vec{\nabla}_{\vec{x}'}\right)\cdot\vec{\nabla}_{\vec{x}'}\rho\left(\vec{x}, T|\vec{x}', t\right). \tag{1.33}$$

Importantly, the gradient and Laplace operator act on $\vec{x}'$ instead of $\vec{x}$. The bFPE is central to many derivations of first-passage-type solutions.

FPTs exist in many different flavours. One of the simpler versions of a FPT is the exit-path time. The exit-path time τ is defined as the time it takes for a stochastic process to leave a certain region S. The appropriate boundary conditions for this problem are absorbing, that is, ∂S: $\rho(\vec{x}', t|\vec{x}, 0) = 0 \,\forall \vec{x} \in \partial S$ [37].

Unfortunately, the language in this area of physics appears to be somewhat loose. I have adopted the convention of referring to exit times as "exit-path times" and to direct

transition times as "transition-path times" (introduced later) in order to emphasise that these are durations of a process which is not instantaneous [94].

Properties of exit-path times are typically derived from the *survival probability* $G(\vec{x},t)$, which is defined as the aggregate probability mass within S, i.e. $G(\vec{x},t)=\int_S \mathrm{d}\vec{x}'\,\rho(\vec{x}',t|\vec{x},0)$, assuming the particle is placed within S at $t=0$. If T denotes the time the processes crosses ∂S and leaves S, the survival probability can be written as $\mathrm{Prob}\,(t\leqslant T)=G(\vec{x},t)$.

Since G is a linear superposition of ρ, it must also solve the bFPE. Using the fact that $p(\vec{x}',t|\vec{x},0)=p(\vec{x}',0|\vec{x},-t)$ [37], the negative sign in Eq. (1.33) disappears and one obtains

$$\frac{\partial G}{\partial t}(\vec{x},t)=\mathbf{A}(\vec{x})\,\vec{\nabla}G(\vec{x},t)+\left(\mathbf{D}\cdot\vec{\nabla}\right)\cdot\vec{\nabla}G(\vec{x},t)\,, \tag{1.34}$$

where I also exchanged the variable names $x'\leftrightarrow x$ in order to simplify the notation. Moreover, the absorbing boundary condition on ∂S requires G to vanish on the boundary, i.e. $G(\vec{x},t)=0\,\forall\vec{x}\in\partial S$.

$G(\vec{x},t)$ can be interpreted as the probability that $T\geqslant t$ for a given initial $\vec{x}$. In other words, it describes the probability that the particle has not left the interval at "time" t. The cumulative probability of T for a given initial $\vec{x}$, is thus given by $1-G(\vec{x},t)$. Therefore, $-\partial_t G(\vec{x},t)$ must be the corresponding probability density of the exit time T.

If the mean exit-path time is written as $\langle\tau(\vec{x}_0)\rangle$ where $\vec{x}_0$ denotes the initial position of the process, $\langle\tau(\vec{x}_0)\rangle$ needs to fulfil [37]

$$\langle\tau(\vec{x}_0)\rangle=-\int_0^\infty \mathrm{d}t\;t\partial_t G(\vec{x}_0,t)=\int_0^\infty \mathrm{d}t\,G(\vec{x}_0,t) \tag{1.35}$$

where I used integration by parts in the last step. Since initially all probability mass is contained within S, $G(\vec{x}_0,0)=1$. However, for $t\to\infty$, the mass will have leaked into the absorbing boundary conditions and thus $G(\vec{x}_0,\infty)=0$.

This implies $\int_0^\infty \mathrm{d}t'\,\partial_{t'}G(\vec{x}_0,t')=-1$.

Since $\langle\tau(\vec{x}_0)\rangle$ is a linear superposition of G, it must solve the bFPE and, using the integral result from above, I obtain a differential equation for mean exit-path times

$$-1=\mathbf{A}(\vec{x}_0)\,\vec{\nabla}\langle\tau\rangle(\vec{x}_0)+\left(\mathbf{D}\cdot\vec{\nabla}\right)\cdot\vec{\nabla}\langle\tau\rangle(\vec{x}_0)\,. \tag{1.36}$$

This well-known equation [37] can be used to derive explicit equations for the mean exit-path time. For instance, for a one-dimensional process with a position-dependent

diffusion coefficient $D(x)$ and symmetrically-placed exits at $\pm L/2$, one obtains [37],

$$\langle \tau(x_0)\rangle = \frac{\int_{-L/2}^{x_0} \mathrm{d}x \int_{x_0}^{L/2} \mathrm{d}x' \int_{-L/2}^{x'} \mathrm{d}x'' \, e^{\frac{U(x)}{k_BT}} e^{\frac{U(x')}{k_BT}} e^{\frac{-U(x'')}{k_BT}} / D(x'')}{\int_{-L/2}^{L/2} \mathrm{d}x \, e^{\frac{U(x)}{k_BT}}} - \frac{\int_{x_0}^{L/2} \mathrm{d}x \int_{-L/2}^{x_0} \mathrm{d}x' \int_{-L/2}^{x'} \mathrm{d}x'' \, e^{\frac{U(x)}{k_BT}} e^{\frac{U(x')}{k_BT}} e^{-\frac{U(x'')}{k_BT}} / D(x'')}{\int_{-L/2}^{L/2} \mathrm{d}x \, e^{\frac{U(x)}{k_BT}}}. \tag{1.37}$$

The above equation simplifies for constant forces f and a constant diffusion profile $D(x) = D$. Under such circumstances, I obtain

$$\langle \tau(x_0)\rangle = \frac{\gamma}{2f} \frac{L - 2Le^{-\frac{f(L-2x_0)}{2k_BT}} + e^{\frac{fx_0}{k_BT}}(L-2x_0) + 2x_0}{e^{\frac{fx_0}{k_BT}} - 1}, \tag{1.38}$$

which for $x_0 = 0$ collapses to

$$\langle \tau \rangle = \frac{L\gamma}{2f} \tanh\left(\frac{fL}{4k_BT}\right). \tag{1.39}$$

1.2.1 Conditional exit-path times

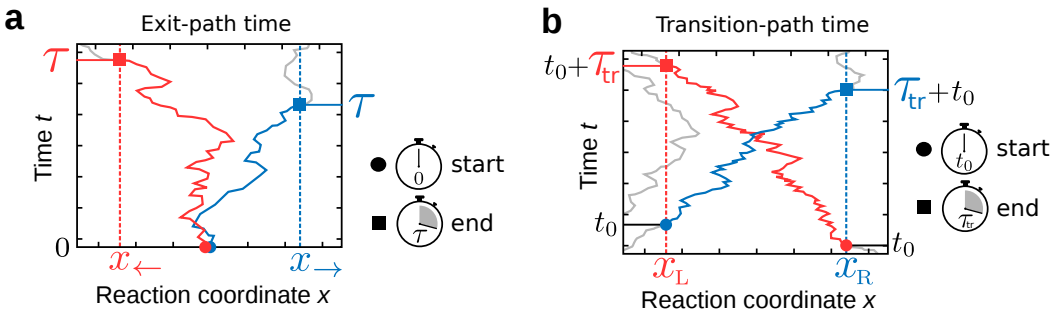

Fig. 1.2 Conditional exit- and transition-path time. **a** Sketch of two instances of a conditional exit-path time. The two traces do not occur simultaneously in real time, but are two instances of an exit-path. **b** Sketch of two transition paths and corresponding path times. The grey traces show paths that are not valid transition paths.

A slightly more complex question that one could ask is: what is the mean exit-path time from a region S through a particular exit $E \subset \partial S$ (if there are several)?

This question can be answered within the Fokker-Planck framework. If the system is initialised with certainty within S, the initial probability mass $\rho_0 = \rho(\vec{x}, 0)$ will be contained within S and then spread into all possible directions, according to the FPE time evolution. At some point it will reach and cross the exit of interest, E. The partial probability to ever exit through E is thus

$$\text{Prob}\left(\text{exit } S \text{ through } E\right) = \int_0^\infty \mathrm{d}t' \int_{\vec{x}\in E} \mathrm{d}\vec{n}\left(\vec{x}\right) \cdot \vec{j}\left(\vec{x}, t'\right) \tag{1.40}$$

where $\vec{n}$ denotes a vector that is perpendicular to ∂S and outward-pointing at the point $\vec{x} \in E$. The double integral in Eq. (1.40) simply measures the total amount of probability flowing through E. It is worth noting, that in a one-dimensional system, the inner integral will conveniently disappear. In this case, all that is left to do is to solve the FPE with the appropriate boundary conditions in order to compute the current density $j(x,t)$.

Consequently, the probability to exit through E at a particular point in time τ reads

$$\rho_E(\tau) = \frac{\int_{\vec{x}\in E} \mathrm{d}\vec{n}\left(\vec{x}\right) \cdot \vec{j}\left(\vec{x}, \tau\right)}{\text{Prob}\left(\text{exit } S \text{ through } E\right)}. \tag{1.41}$$

And the mean exit-path time conditioned on exit E reads

$$\left\langle \tau_E \right\rangle = \frac{\int_0^\infty \mathrm{d}\tau\, \tau \int_{\vec{x}\in E} \mathrm{d}\vec{n}\left(\vec{x}\right) \cdot \vec{j}\left(\vec{x}, \tau\right)}{\text{Prob}\left(\text{exit } S \text{ through } E\right)}. \tag{1.42}$$

Conditional exit times of Brownian processes exhibit rich and surprising behaviour. For instance, a process $\dot{x}(t) = \xi(t)$ that is initialised in the centre $x_0 = (a+b)/2$ of an one-dimensional interval $[a,b]$ will, on average, exit the interval to either side at the same time, i.e. $\langle \tau_a \rangle = \langle \tau_b \rangle$. Crucially, this remains so even under the influence of an external drift force $f(x) = c$. The process then reads $\dot{x}(t) = c + \xi(t)$.

To my knowledge, this central result in stochastic dynamics was first derived by Frederick Stern in 1977 [95]. In subsequent years, mathematicians derived the density of hitting times and places of spheres of a Brownian process with a drift initialised within and outside of spheres [96–98].

Conditional exit time from a channel

In spite of its apparent simplicity, it can be cumbersome to derive the conditional exit-time density for a one-dimensional process analytically from Eq. (1.41). The problem arises from the boundary conditions of the FPE, which require the density to vanish at the exits, that is, for a one-dimensional process at $x = \pm L/2$. In this case, one can construct the position density $\rho(x,t)$ using the method of images, which involves an infinite series. However, its derivative, the flux j, is not easily reformulated in an insightful way such that I omit the solution here. Instead, I resort to solving most these distributions numerically by solving the corresponding FPE.

I begin here by stating the total exit probability though a particular exit, which can be directly obtained from Eq. (1.40). These probabilities are also known as splitting probabilities and read [37]

$$P_{\rightarrow}(x_0) = \frac{\int_{-L/2}^{x_0} \mathrm{d}x\, e^{\frac{U(x)}{k_B T}}}{\int_{-L/2}^{L/2} \mathrm{d}x\, e^{\frac{U(x)}{k_B T}}} \tag{1.43}$$

$$= \frac{1}{1 + e^{\frac{fL}{2k_B T}}} \qquad \text{for} \qquad x_0 = 0,\ U(x) = -fx \tag{1.44}$$

$$P_{\leftarrow}(x_0) = 1 - P_{\rightarrow}(x_0) = \frac{e^{\frac{fL}{2k_B T}}}{1 + e^{\frac{fL}{2k_B T}}} \qquad \text{for} \qquad x_0 = 0,\ U(x) = -fx \tag{1.45}$$

Interestingly, the two probabilities, $P_{\rightarrow}(x_0 = 0, f)$ and $P_{\leftarrow}(x_0 = 0, f)$, fulfil what is known as a *detailed* fluctuation theorem [21], i.e.

$$\frac{P_{\rightarrow}(f)}{P_{\leftarrow}(f)} = e^{-\frac{fL}{2k_B T}} = e^{-\frac{\Delta s}{k_B}} \tag{1.46}$$

where Δs denotes the difference in produced (or absorbed) entropy. This underlines the fact that uphill-exiting trajectories that defy the external force *consume* entropy (heat) to do so. In order to get a sense of the forces required to create a measurable difference between the splitting probabilities $P_{\rightarrow}$ and $P_{\leftarrow}$, I plot the expression in Eqs. (1.45) and (1.44) in the inset of Fig. 1.3.

The time dependent density ρ and the current can be solved numerically. In Fig. 1.3, I plot the time-evolution of the density with the correct boundary conditions. Once the density is obtained, the exit-path time density can be calculated using Eq. (1.41). The results for the uphill and downhill directions are plotted in Fig. 1.4**a**. In the experimental

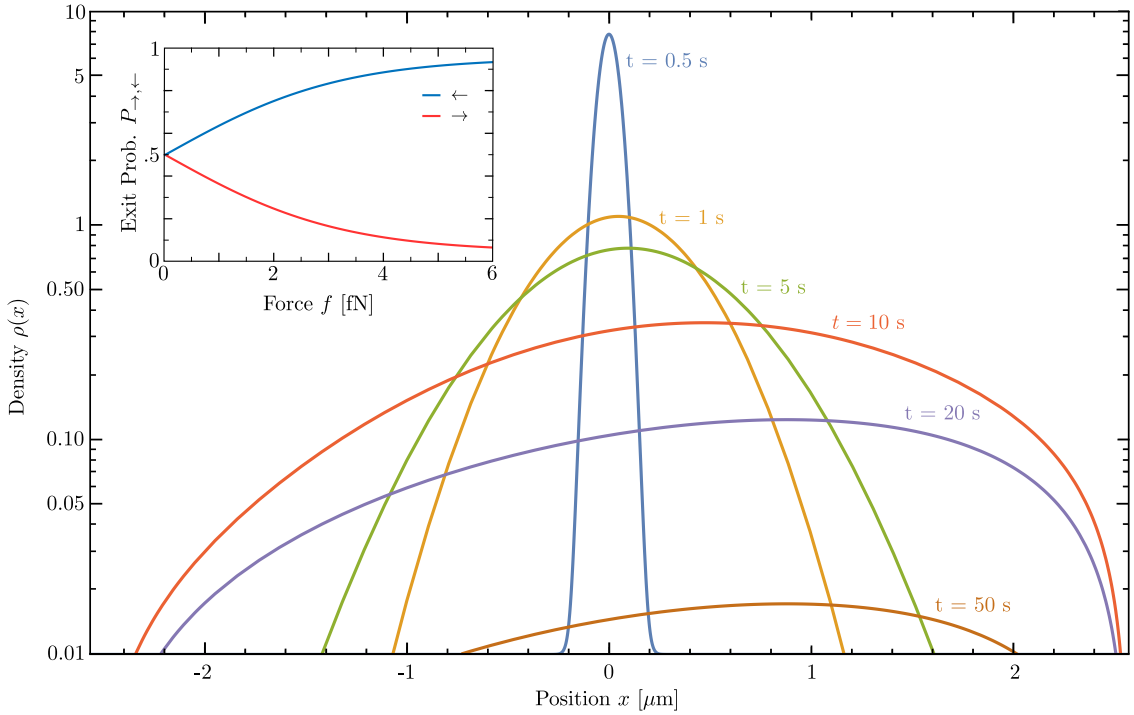

Fig. 1.3 Log-plot of the time-evolution of initial density (black line at $x_0 = 0$) with a constant drift force $f = +3$ fN over an interval of $L = 5.14\,\mu$m. The diffusion constant is $D = 0.13\,\mu\text{m}^2/s$. In the inset, exit probabilities (Eqs. (1.44) and (1.45)) are plotted against (red) and with (blue) the force for the same interval length and force. The plot is supposed to give a sense of the magnitude of forces required to create measurably different splitting probabilities.

sections, I will mostly compare cumulative distributions, since they can be compared independently of the choice of bin sizes. For this reason, I also plot the cumulative distributions in panel **b**.

As Fig. 1.4 clearly shows, the uphill and downhill exit-path time distributions agree. In contrast to their distributions, the mean conditional exit-path time can be readily calculated analytically [95]. Due to the observed exit-path symmetry, both expressions are exactly equal to the one given in Eqs. (1.38) and (1.39).

1.2.2 Transition-path times

Instead of asking how long it takes for a process to leave a certain region S, one might as well ask how long it takes for a process to *cross* a certain region S, e.g. the interval in Fig. 1.5**a**.

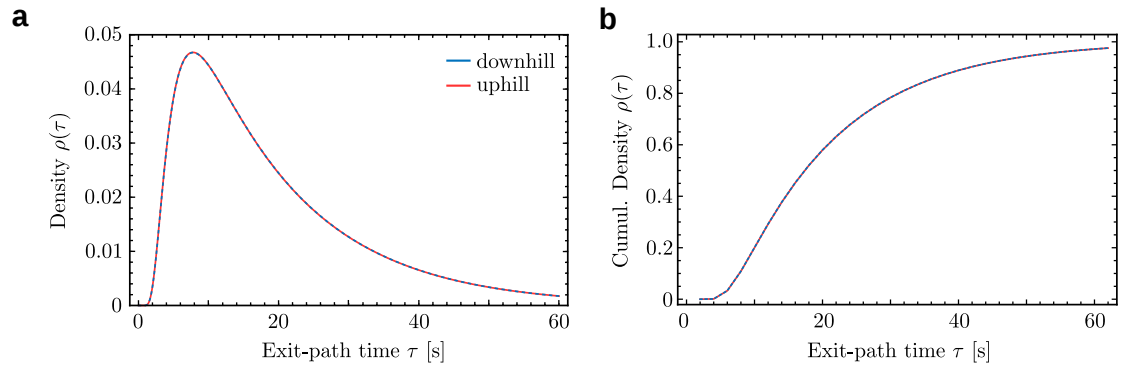

Fig. 1.4 **a** Exit-path time density $\rho(\tau)$ in the uphill and downhill direction (red-dashed, black). **b** The same distributions but plotted as cumulative distributions. In experimental situations, it is often better to compare cumulative distributions.

This question then leads to the notion of direct transition times, here referred to as transition-path times.

Transition-path times begin on one boundary or one part of the boundary of a connected region and end in another part of the boundary. This boundary is typically partially or completely absorbing, such that any theoretical treatment requires careful limiting procedures in order to arrive at sensible conclusions. From an experimental point-of-view, the procedure to measure such times is simple: One observes long, uninterrupted trajectories of the variables of interest and later on considers only those parts of the trajectories that cross the interval in one go. This procedure automatically implements purely absorbing boundaries at or close to the chosen boundaries. I will come back to this point in greater detail in Sec. 3.1.

Historically, A. Szabo is credited with having derived the first expression of a transition-path time [88, 43]. Comprehensive and modern papers on direct transition times in one-dimensional systems that are relevant here are Berezhkovskii *et al* [90–92], Kim *et al* [88], and Zhang *et al* [87]. The former derives analytical results for partially absorbing boundaries, which are readily simplified to the purely absorbing case. Kim and Zhang derive expressions for the distribution of transition-path times, assuming absorbing boundary conditions. In the following, I here draw on the results of all three authors.

Let $\rho(x,t|x_0,t_0=0)$ be the solution of the overdamped, one-dimensional forward FPE with an intial condition $\rho(x,0|x_0)=\delta(x_0-x)$. As sketched in Fig. 1.5**b**, I assume that there is a free-energy function $U(x)$ defined over the interval $x\in[x_l,x_r]$ of interest, portrayed in panel **a**. The partially-absorbing boundary conditions considered by Berezhkovskii that ρ has to fulfil are of the von-Neumann type and also known as "radition boundary

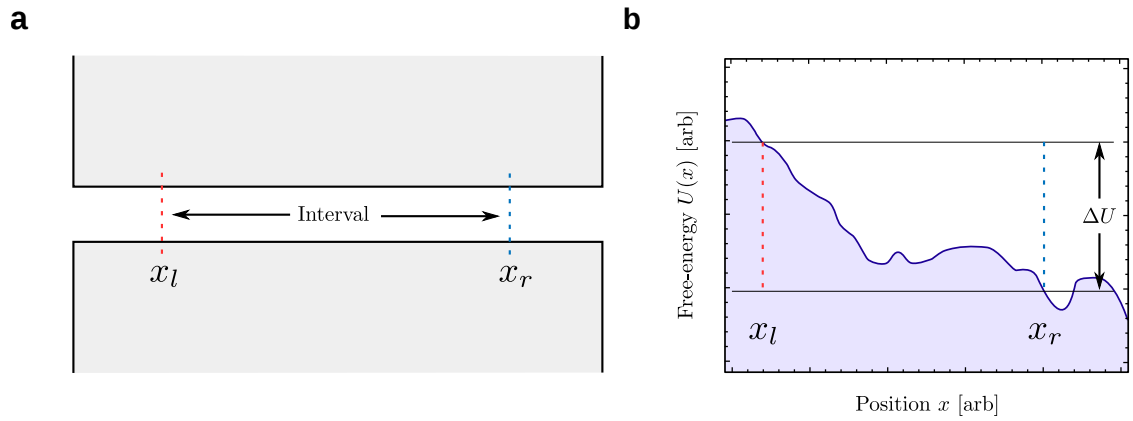

Fig. 1.5 **a** Sketch of the one-dimensional interval of interest (here realised in a confining channel). **b** Example of a free-energy trace over this interval.

conditions". They read

$$\partial_x \left(\rho(x,t|x_0) e^{\frac{U(x)}{k_B T}} \right) |_{x=x_l} = \frac{k_l}{D(x_l)} e^{\frac{U(x_l)}{k_B T}} \rho(x_l,t|x_0) \tag{1.47}$$

$$\partial_x \left(\rho(x,t|x_0) e^{\frac{U(x)}{k_B T}} \right) |_{x=x_r} = -\frac{k_r}{D(x_r)} e^{\frac{U(x_r)}{k_B T}} \rho(x_r,t|x_0) \tag{1.48}$$

with k_l and k_r denoting the rate of radiation at each boundary. The flux through the channel exits therefore has to fulfil $j(x_l,t|x_0) = -k_l \rho(x_l,t|x_0)$ and $j(x_r,t|x_0) = k_r \rho(x_r,t|x_0)$.

As before in Eq. (1.40), the total probability to exit through a particular exit is given by the time-integral of the respective flux, $P_l = \int_0^\infty \mathrm{d}t - j(x_l,t|x_0)$ and $P_r = \int_0^\infty \mathrm{d}t\, j(x_r,t|x_0)$.

The average conditional exit-path times (referred to as lifetimes by Berezhkovskii [90]) for the radiative boundaries take the form

$$\langle \tau_l(x_0) \rangle = \frac{1}{P_l(x_0)} \int_0^\infty \mathrm{d}t \, (-t j(x_l,t|x_0)) = \frac{k_l}{P_l(x_0)} \int_0^\infty \mathrm{d}t \, t \rho(x_l,t|x_0) \tag{1.49}$$

$$\langle \tau_r(x_0) \rangle = \frac{1}{P_r(x_0)} \int_0^\infty \mathrm{d}t \, t j(x_r,t|x_0) = \frac{k_r}{P_r(x_0)} \int_0^\infty \mathrm{d}t \, t \rho(x_r,t|x_0) \tag{1.50}$$

in accord with Eq. (1.42).

Finally, the average transition-path time, or translocation time can be written as $\langle \tau_{\mathrm{tr}}^{\rightarrow} \rangle = \langle \tau_r(x_l) \rangle$ and vice-versa $\langle \tau_{\mathrm{tr}}^{\leftarrow} \rangle = \langle \tau_l(x_r) \rangle$. Berezhkovskii then proceeds to solve the

system of equations, Eqs. (1.47)-(1.50), which yields

$$\langle \tau_{\mathrm{tr}} \rangle = \frac{\int_{x_l}^{x_r} dx\, e^{\frac{-U(x)}{k_B T}} \left(1 + k_l \int_{x_l}^{x} \mathrm{d}x'\, \frac{e^{\frac{U(x')}{k_B T}}}{D(x')}\right)\left(1 + k_r e^{\frac{-\Delta U}{k_B T}} \int_{x}^{x_r} \mathrm{d}x'\, \frac{e^{\frac{U(x')}{k_B T}}}{D(x')}\right)}{k_l + k_r e^{\frac{-\Delta U}{k_B T}} + k_l k_r e^{\frac{-\Delta U}{k_B T}} \int_{x_l}^{x_r} \mathrm{d}x\, \frac{e^{\frac{U(x)}{k_B T}}}{D(x)}}. \tag{1.51}$$

The expression above is independent of the direction of the transition, despite the possible difference in free-energy ΔU. Berezhkovskii acknowledges this fact and writes "*At the moment we have no simple qualitative explanation for the direction independence of the average translocation time.*". Testing and explaining this observation in experiments constitutes a major part of this thesis.

At any point in the above derivation, one can obtain Szabo's classical expression by sending both k's to infinity. This corresponds to an infinite rate at the respective boundary and thus imposes absorbing boundary conditions. For the sake of completeness, I state Szabo's result for the mean transition-path time for purely absorbing boundaries below

$$\langle \tau_{\mathrm{tr}} \rangle = \frac{\int_{x_l}^{x_r} dx\, e^{\frac{-U(x)}{k_B T}} \int_{x_l}^{x} \mathrm{d}x'\, \frac{e^{\frac{U(x')}{k_B T}}}{D(x')} \int_{x}^{x_r} \mathrm{d}x'\, \frac{e^{\frac{U(x')}{k_B T}}}{D(x')}}{\int_{x_l}^{x_r} \mathrm{d}x\, \frac{e^{\frac{U(x)}{k_B T}}}{D(x)}}. \tag{1.52}$$

The final aspect of transition-path times that I want to touch on here is their distribution $\rho_{\mathrm{tr},\leftarrow,\rightarrow}(\tau)$. Following Zhang *et al* [87], this density can be derived from Eq. (1.41) by sending the initial distribution $\rho_0(\vec{x}) = \delta(\vec{x} - \vec{x}_0)$ carefully close to the initial exit (i.e. the point of entry). Importantly, ρ_0 should be completely contained in S. For a one-dimensional example, this can for instance be achieved by parametrising the initial position x_0 by some $\varepsilon > 0$, $x_0 = x_{\mathrm{l}} + \varepsilon$ in this case of the left exit being designated as initial position. Again, as in the case of exit-path times, I consider the normalised flux through the target-exit surface as the distribution of transition-path times. For instance, for $l \to r$ transitions, this reads

$$\rho_{\tau_{\mathrm{tr}}\rightarrow}(\tau) = \lim_{\epsilon \searrow 0} \frac{j\left(x_{\mathrm{r}} | x_0(\epsilon)\right)}{P_{\rightarrow}\left(x_0(\epsilon)\right)}. \tag{1.53}$$

where $P_{\rightarrow}(x_0)$ is given in Eq. (1.43).

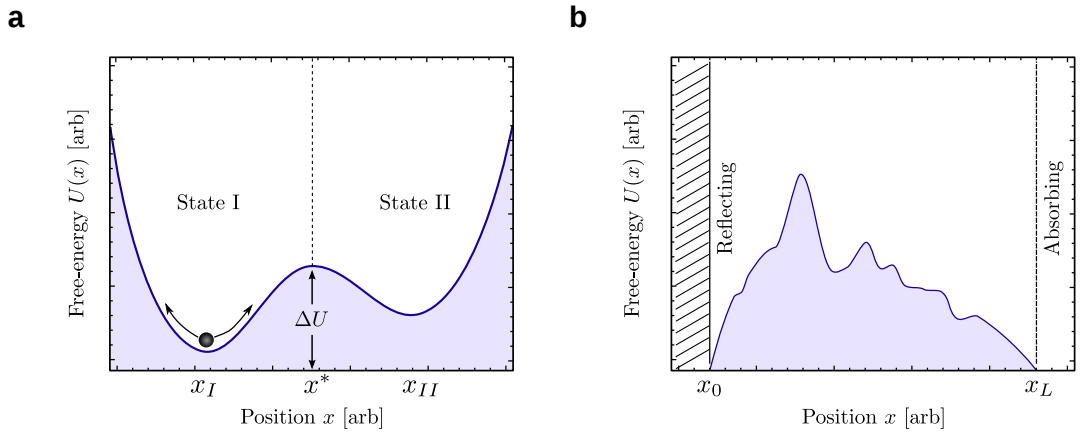

Fig. 1.6 **a** Sketch of a typical Kramer's reaction rate double-well scenario. **b** Sketch of a typical exit-path scenario. In this section, I calculate the exit-path time for one reflecting and one absorbing barrier.

1.2.3 Escape times & Kramers rate

The transition-path time discussed in the previous section quantifies the time it takes to cross some region, e.g. a free-energy barrier but it disregards unsuccessful attempts. Only the time actually spent on the barrier is taken into account. The rate of barrier crossings in an ensemble of two-state systems would thus not correspond to the inverse of the mean transition-path time. Instead, the rate must be obtained from the reciprocal of a different first-passage time, here referred to as Kramer's time.

In a one-dimensional system, with two states I, II that are connected by a single point, Kramers' time for the transition $I \to II$ equals the exit-path time from state I. Such a situation naturally occurs in bistable potential with each minimum being assigned a "state" as sketched in Fig. 1.6**a**.

As a first step towards deriving the rate of barrier transition in Fig. 1.6**a**, I begin with the case shown in Fig. 1.6**b**: the exit from an interval that is delimited to the left by a reflecting boundary at $x_{\text{refl.}}$ and to the right by an absorbing one at $x_{\text{abs.}}$.

The boundary conditions for this case are [37]

$$\partial_x G(x_{\text{refl.}}, t) = 0 \tag{1.54}$$

$$G(x_{\text{abs.}}, t) = 0. \tag{1.55}$$

If one solves the mean exit-path time equation, Eq. (1.42) observing these boundary conditions, one obtains

$$\left\langle \tau(x_0) \right\rangle = \int_{x_0}^{x_{\text{abs.}}} \mathrm{d}x\, e^{\frac{U(x)}{k_B T}} \int_{x_{\text{refl.}}}^{x} \mathrm{d}x'\, \frac{e^{-\frac{U(x')}{k_B T}}}{D(x')}. \tag{1.56}$$

In the case of $x_0 = x_{\text{refl.}} = 0$ and $x_{\text{abs.}} = L/2$, this result can also be obtained from Eq. (1.37) for a symmetric potential $U(x) = U(-x)$.

The aforementioned Kramers' time is readily obtained from Eq. (1.56) by exchanging x_0, $x_{\text{refl.}}$, and $x_{\text{abs.}}$ in the following way: $x_{\text{refl.}} \to -\infty$, $x_{\text{abs.}} \to x_{\text{II}}$, and $x_0 \to x_{\text{I}}$,

$$\left\langle \tau_{\text{I}\to\text{II}} \right\rangle = \int_{x_{\text{I}}}^{x_{\text{II}}} \mathrm{d}x\, e^{\frac{U(x)}{k_B T}} \int_{-\infty}^{x} \mathrm{d}x'\, \frac{e^{-\frac{U(x')}{k_B T}}}{D(x')}. \tag{1.57}$$

H.A. Kramers derived a simplified version of the above double integral in his influential paper about rates of barrier transition [99]. This approximation is derived from Eq. (1.57) under the assumption that the barrier height ΔU is much larger than the thermal energy $k_B T$. For simplicity, I also assume here a constant diffusion profile $D(x) = D$. In this case, one can, approximately, decouple the integrals and replace the lower limit of the inner integral by x_I, the initial position [37]. Furthermore, since the peak contributes the bulk of the value of the outer integral, the upper limit of the outer integral can be replaced by x^*, such that

$$\left\langle \tau_{\text{I}\to\text{II}} \right\rangle \approx \frac{1}{D} \left(\int_{x_{\text{I}}}^{x^*} \mathrm{d}x\, e^{\frac{U(x)}{k_B T}} \right) \left(\int_{x_{\text{I}}}^{x_{\text{II}}} \mathrm{d}x\, e^{-\frac{U(x)}{k_B T}} \right). \tag{1.58}$$

Both integrals can be further approximated: one develops $\frac{U(x)}{k_B T}$ into a series around the location of the minimum x_{I} for the left integral and around the peak x^* for the right integral, i.e. $U(x) \approx U(z) + \frac{1}{2}(x-z)^2 U''(z)$. The second derivative of the potential equals the local curvatures, which I abbreviate as $U''(x_{\text{I}}) \equiv \omega_1, U''(x^*) \equiv \omega_2$. Using the Gaussian integral, Eq. (1.58) simplifies to

$$\left\langle \tau_{\text{I}\to\text{II}} \right\rangle \approx (2\pi D)^{-1} \sqrt{\omega_1 \omega_2}\, e^{\frac{\Delta U}{k_B T}} \tag{1.59}$$

where I used the definition of the barrier height $\Delta U = U(x^*) - U(x_{\text{I}})$ in Fig. 1.56.

Kramers result in Eq. (1.59) is of paramount importance to chemistry. The crucial aspect of this equation is, as I mentioned before, that the reciprocal FPT here corresponds to the rate of barrier transitions. In the case of a reaction, such as protein folding, this microscopically-derived rate corresponds to the ensemble reaction rate, which is macroscopically observable. In this case, the scaling law in Eq. (1.59) is often also referred to as Arrhenius' law.
As I show in Sec. 3.2, Eq. (1.59) fundamentally breaks down for certain barrier shapes, with the implication that higher barriers may paradoxically accelerate reactions.

In this chapter, I introduce a few key ideas of stochastic thermodynamics and derive equations governing various first-passage processes. The next chapter summarises the setup and relevant experimental concepts used to test the theories described in this chapter.

Chapter 2

Experimental Foundations

The aim of this thesis is to combine the theory of first-passage times with clean and controlled Brownian dynamics experiment in microfluidic systems.

The tool of choice for such experiments are optical tweezers, due to their unique ability to exert and measure forces on the relevant scales (femto- to pico-Newton). By definition, first-passage time experiments require large sample sizes which in turn requires stable and repeatable experiments. My approach to this was to automate as many aspects of my experiments as possible such that the setup could independently carry out as many experiments as possible or required.

2.1 Introduction to Optical Tweezers

Optical tweezers are sensitive and precise devices, which can be used for manipulating particles on the micronscale and below. Their ability to trap particles in conjunction with their ability to measure particle displacements with high accuracy makes them an invaluable tool for biophysical research. Arguably, the history of optical trapping began with Arthur Ashkin [100] who discovered that dielectric particles experience a force when placed in a gradient of intensity of light. Ashkin provided the proof of principle of what would become known as *optical trap* and was awarded the Nobel prize in 2018 for his invention.

Optical tweezers are a particular form of an optical trap, in which a single laser beam is utilised to confine an object (Ashkin initially explored usage of two opposing lasers). The key insight that led to optical tweezers, is that gradient forces can not only overcome thermal forces, but also radiation pressure of the laser beam such that the particle can rest in a stable state near the focal point [101, 102]. Optical tweezers later became a popular tool in biophysics and thermodynamics [103]. While Ashkin in his 1970 paper still

referred to thermal forces as "*obscuring effects*", colloidal particles held by optical tweezers are nowadays an established toy model for studying thermal fluctuations. For instance, transient violations of the second law have been observed in optical traps [104] and its generalisations have been tested subsequently [105]. Moreover, Landauer's principle has been experimentally confirmed using optical trapping [106, 107].

2.1.1 Theory of optical trapping forces

The introduction to forces in optical tweezers in this subsection is an adaptation from my Master's thesis [108].

Introductions to the physics of optical tweezers typically begin by distinguishing the *Rayleigh* regime for small particles $d \ll \lambda$ with d denoting the diameter, from the *geometrical optics* regime applicable to larger ones $d \gg \lambda$ and then proceed to explain that typical colloidal handles are, unfortunately, in between, i.e. $d \sim \lambda$ [109, 110, 103]. It is most illuminating to discuss optical tweezers in the ray optics regime since one only has to consider the geometry of the effect of trapping on photon-momenta. By contrast, on the nanoscale, it is necessary to consider the interaction of induced dipoles with the laser light in order to explain the trapping effect. I therefore begin here with a more qualitative discussion of the first regime. In both cases, I assume the laser beam to adhere to Gaussian beam theory with a $1/e$ waist w_0 and a Gaussian intensity profile (TEM_{00} mode), as is the case in my experiments.

Ray optics regime - A plane electromagnetic wave is associated with a momentum flux $\vec{p}$ (momentum per volume) in the direction of the *Poynting vector* $\vec{S}$, $\vec{p} = 1/c^2 \vec{S} = \varepsilon_0 \left(\vec{E} \times \vec{B}\right)$ with a magnitude that is proportional to the intensity [111]. Bending or reflecting a laser beam will therefore result in a *reactio*-force proportional to beam power P_0, $F = P_0/c$.

I now consider a dielectric particle suspended in a viscous medium with respective indices of refraction obeying $n_p > n_m$. Only with this configuration of refraction, trapping is possible in my setup. Particles with a lower refractive index, such as air bubbles are pushed out of the trap. I assume the particle to be large compared to the optical wavelength, such that ray optics applies ($d > 20\lambda$, [109]). If such a particle inches closer to either side of the trap, a higher intensity will be refracted outward and the particle will hence experience an inward force (see Fig. 2.1**a**). In all cases, reflected or refracted rays impart a force directed downstream that must be somehow balanced for stable trapping. Using a high aperture objective, one can create highly con- and divergent rays at the focus. Beams refracted by the particle when above or below the focal point will be less divergent, as shown in Fig. 2.1**b**, resulting in a restoring force that may suffice to counteract radiation pressure (and thermal forces). Importantly, the axial scattering force pushing the bead

downstream is non-conservative. If integrated over a closed loop, these forces would result in a net force pointing downstream. The path dependence in optical scattering forces has been found to cause a breaking of detailed balance in fluctuations in the radial-axial coordinates of a particle held by optical tweezers [112, 113].

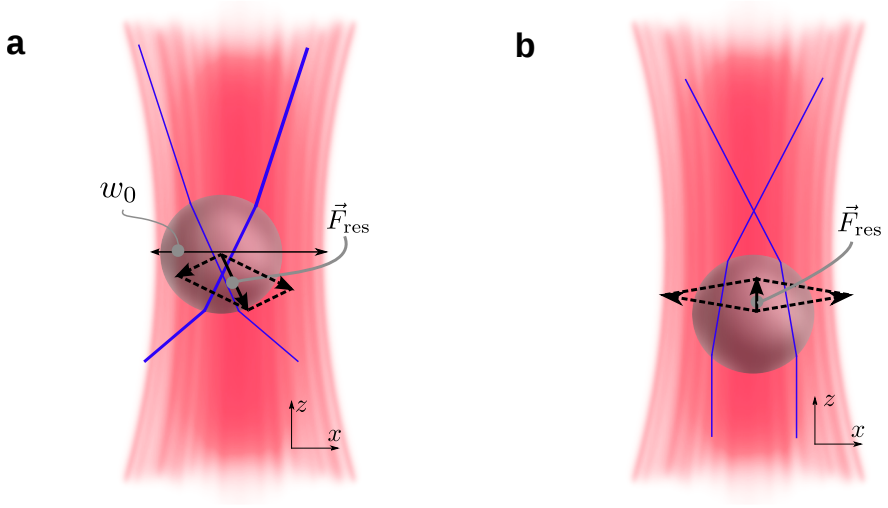

Fig. 2.1 Forces in the ray-optics regime in optical tweezers. This figure is a reprint of my figure in [108]. The thickness of the blue rays is a measure of their intensity. **a** Lateral restoring gradient force arising due to asymmetry of intensity experienced by a horizontally displaced particle. The minimal beam waist w_0 is indicated. **b** Vertical restoring force that counteracts radiation pressure in the axial direction. Angles and magnitudes are not drawn to scale in either panel.

Rayleigh regime - In the limit of small particle diameters, the theory of Rayleigh scattering becomes applicable. Electric fields $\vec{E}$ lead to charge separation and induce a dipole in dielectric materials. Inhomogeneous electric fields, such as the intensity gradient in the laser focus, exert Lorentz forces on the dipole. This force, averaged over time, reads for small dielectric spheres [103]

$$\vec{F}_{\mathrm{grad}}(\vec{r}) = \frac{\alpha}{2}\vec{\nabla}\langle\vec{E}^2(\vec{r})\rangle = n_m a^3 \left(\frac{(n_p/n_m)^2-1}{(n_p/n_m)^2+2}\right)\vec{\nabla} I(\vec{r}) \tag{2.1}$$

where α is referred to as *polarisability* and $I = cn\varepsilon_0/2\langle E^2\rangle$ denotes the intensity. This force is the small-particle-regime equivalent to the gradient force previously discussed. I note, that since $\vec{F}_{\mathrm{grad}}$ is derived from a gradient, namely, $\vec{\nabla} I$, it must be conservative.

By contrast, the scattering force that a particle in an optical trap is subject to points into the direction of the Poynting vector $\vec{S}$ and reads [103]

$$\vec{F}_{\text{scat}}(\vec{r}) = n_m \sigma \frac{\langle \vec{S}(\vec{r}) \rangle}{c} = \frac{8 n_m \pi (\kappa a)^4 a^2}{3c} \left(\frac{(n_p/n_m)^2 - 1}{(n_p/n_m)^2 + 2} \right)^2 \langle \vec{S}(\vec{r}) \rangle \tag{2.2}$$

with σ denoting the cross-section of a the particle with radius $a = d/2$. Gradient and scattering forces are only clearly separable in the Rayleigh regime [103]. A rigorous treatment of the intermediate regime, in which particle size and wavelength are comparable, requires a full solution of the Maxwell equations under geometrical constraints of the scenario (*Lorentz-Mie theory*, I refer for further reading to [109, 114]).

2.2 Holographic optical tweezers

In the previous section, I explain in some detail the forces that microscopic objects experience when placed into highly focused laser beams. The key lesson is that, on the micronscale, controlling light means controlling force.

Arguably, the most powerful optical system is a hologram. This is so, because holograms can be made to model almost any other optical element, such as lenses. But before I explain what holograms can do, I should explain what a hologram is. Modern reconstructing holograms were invented by the Hungarian physicist and Nobel-laureate Dennis Gabor in the late 1940s [115]. I use the qualifier "reconstructing", because there are nowadays different forms of holograms. Common to all forms is the fact that a hologram is essentially an image, but instead of absorbing part of the light spectrum, each "pixel", in addition, alters the phase of incoming light beams. It is even possible, to disregard the amplitude altogether and exclusively change the phase. In this case, a hologram is a spatial map of phase-delays or advancements. Occasionally, such phase-only optical devices are referred to as *kinoforms* [116]. Holography for optical tweezers is a recent development and was pioneered, among others, by David Grier's group [117, 116, 118–120].

2.2.1 A Fourier-optics perspective on holographic optical tweezers

The medium or rather physical device that implements the holographic transformation in my case is a spatial light modulator (SLM). I explain in greater detail the inner workings of SLMs in the next section. For now, it is sufficient to know that the SLM is essentially a reflective liquid crystal display with, in my case, 800×600 pixels, where each pixel can be assigned some value 8-bit value, i.e. $\text{SLM}_{ij} \in [0, 255]$. This value encodes a phase delay or

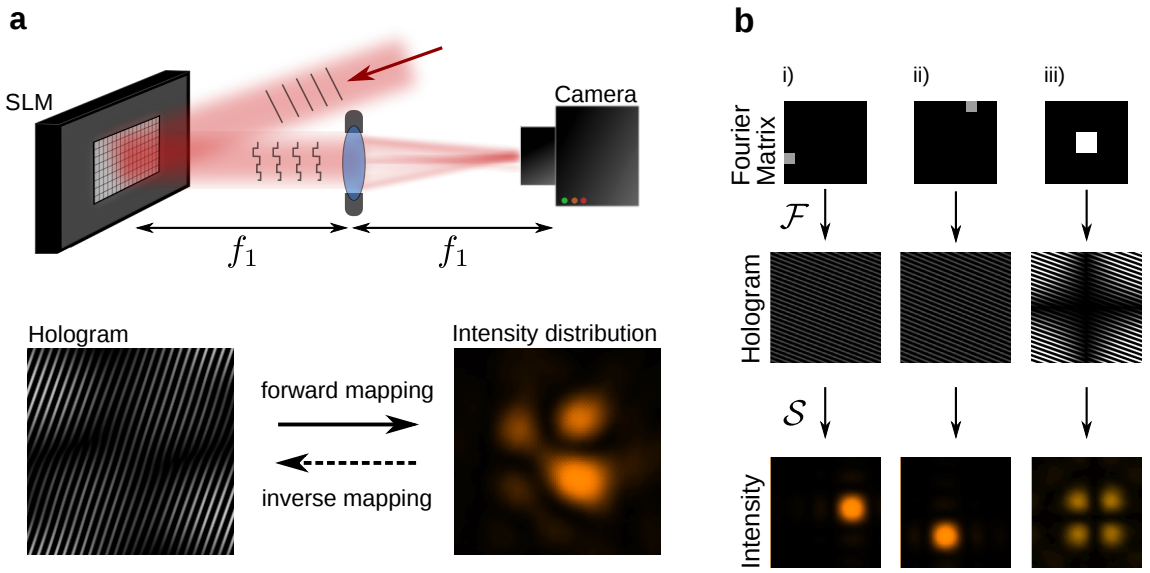

Fig. 2.2 A Fourier-optics description of digital holography. **a** Simplified sketch of waveform-shaping through phase-control. A lens with a focal length f_1 is placed in a 2f-configuration with respect to the SLM and the camera. **b** A few examples of the relation between (real-valued) Fourier-matrices, holograms and the resultant intensity distribution. The third example, iii), in comparison with i) and ii) shows the non-linear relation between Fourier-components and the intensity: Despite higher and more non-zero input components, the output intensity of each spot is reduced. The intensity is not necessarily conserved compared to single-spot examples, since additional (ghost) spots may form outside the field-of-view upon addition of frequencies. The mappings F and S are introduced later in this thesis. Pictures of intensity distributions are contrast enhanced for better visibility.

advancement, which can be mapped to actual phase-changes $\mathrm{SLM}_{ij} \rightarrow \Delta\varphi_{ij} \in [-\pi, \pi]$ of the impinging wavefront.

The action of the SLM on an incoming, planar laser beam is best understood within the framework of Fourier optics.

This framework applies when the relationship between two conjugate wavefronts is considered. Here, I consider the electrical-field value E_{SLM} of a wavefront just after the interaction with the SLM and its conjugate wavefront impinging on the camera E_{cam}, as sketched in Fig. 2.2**a**. Fourier optics relies on Fraunhofer's approximation of a monochromatic wave that is propagated through a lens in a 2f-configuration [121]

$$E_{\mathrm{cam}}(r_x, r_y) \approx -i\frac{\kappa}{f_1} e^{i4\pi\kappa f_1} \int \mathrm{d}\tilde{r}_x \int \mathrm{d}\tilde{r}_y \, E_{\mathrm{SLM}}(\tilde{r}_x, \tilde{r}_y) e^{-i2\pi\frac{\kappa}{f_1}\left(r_x\tilde{r}_x + r_y\tilde{r}_y\right)}. \tag{2.3}$$

where κ denotes the wave-number of the beam. The wave-number can be put into relation to the wave-length λ, the speed of light c, the refractive index of the medium n

and the temporal frequency of the wave ν, $\kappa = n/\lambda = n/c\nu$. The other new variable in Eq. (2.3), f_1 represents the focal length of the lens in Fig. 2.2. Importantly, the Fourier transformation runs over the lateral degrees-of-freedom r_x, r_y, that are perpendicular to the beam direction.

Since the SLM in my setup can only affect the phase, its action can be condensed into a unitary operator $p(r_x, r_y) = e^{i\Delta\varphi(r_x,r_y)}$, which transforms the incoming wave E_0 into

$$E_{\mathrm{SLM}}(r_x, r_y) = e^{i\Delta\varphi(r_x,r_y)} E_0(r_x, r_y). \tag{2.4}$$

The field is advanced to the camera plane using Eq. (2.3) and the intensity that the camera then perceives is given by

$$I_{\mathrm{cam}}(r_x, r_y) = ||E_{\mathrm{cam}}(r_x, r_y)||^2. \tag{2.5}$$

The above Eq. (2.5) shows the fact that any global phase-offset $e^{i\Delta_{\mathrm{global}}\varphi}$ in the SLM plane will not affect the intensity in the camera plane. For this reason one can neglect any laterally-constant phase, such as $i4\pi\kappa f_1$ in Eq. (2.3). The phase invariance plays an important role in holography and I come back to its implications in Sec. 2.2.5.

Of course, the phase-change $\Delta\varphi(r_x, r_y)$ is piece-wise constant since the SLM is pixelised. However, effects of (1) pixelation, (2) digitisation of phase-changes $n \in [0,255] \rightarrow \Delta\varphi_n \approx n2\pi/255 - \pi$, (3) the pixel area and pitch, and (4) effects related to non-linearities in the grey-value-to-phase-shifting relation of the liquid crystals in the SLM are manifold and mostly beyond the scope of this thesis.

A good summary of the various aberrations haunting holographic optical tweezers can be found in Alexander Jesacher's PhD thesis [122]. In my setup, instead of a single, one needs to apply three consecutive Fourier-transformations of the electrical field starting at the SLM: Fourier space (SLM) $\xrightarrow{1}$ Real space (mid plane) $\xrightarrow{2}$ Fourier space (obj. back aperture) $\xrightarrow{3}$ Real space (microscope plane), as shown later in Fig. 2.6. However, apart from a spatial rescaling due to magnifications, the logic described above still applies.

So how do holographers create the intensity distribution on the camera in Fig. 2.2, that they want?

This question is more involved as the relevant Eqs. (2.3)-(2.5) at first suggest. The problem arises from the fact that the variable of control is the field of phase-shifts $\Delta\varphi(r_x, r_y)$, which must be purely real. The situation would be much simpler if $p(r_x, r_y)$ was non-unitary and could affect amplitudes and phases alike, since finding a suitable hologram would then reduce to applying an inverse Fourier transformation. For instance, a diffraction-limited spot in the image plane $E_{\mathrm{cam}}(r_x, r_y) \propto \delta(r_x - k_x)\delta(r_y - k_y)$ requires

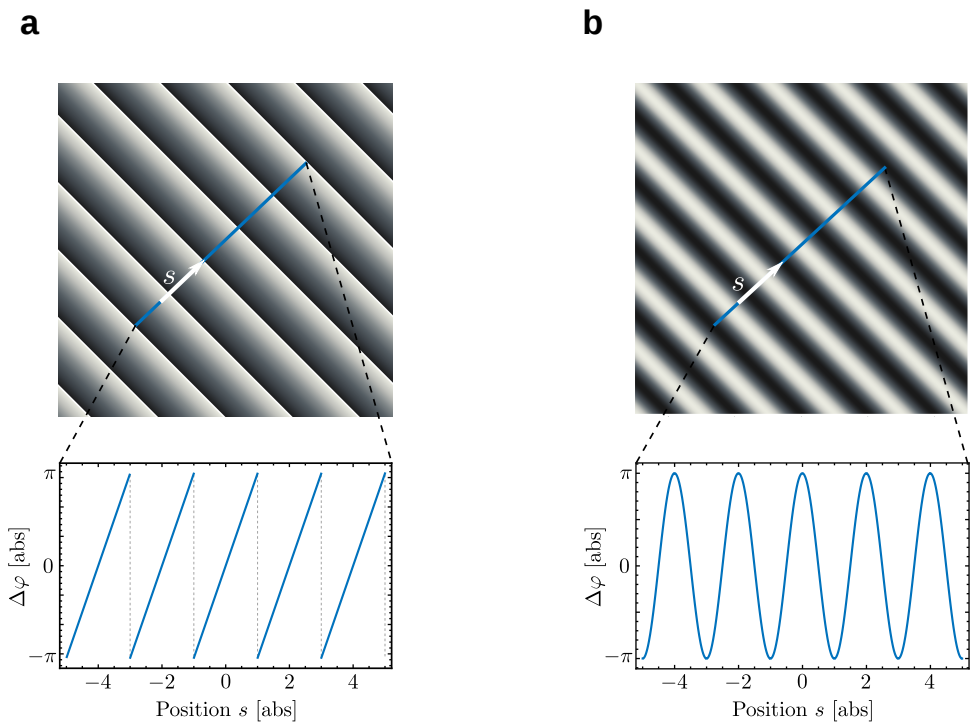

Fig. 2.3 Two types of holographic gratings used in this thesis. **a** Blazed grating for a single point trap. **b** Fourier grating with similar spatial frequency.

$E_{\mathrm{SLM}}(r_x, r_y) \propto e^{i\vec{k}\cdot\vec{r}}$, if I ignore the scaling factor f_1, the aperture of the SLM, and the incident wave structure $E_0(r_x, r_y)$ for a moment. This in turn would require $p(x,y) \propto e^{i\vec{k}\cdot\vec{r}}$ and hence $\Delta\varphi(\vec{r}) = \vec{k}\cdot\vec{r}$, a linearly increasing phase shift. Since the SLM values are limited to a range of (idealised) $[-\pi, \pi)$, one instead takes the complex argument, i.e. $\Delta\varphi(\vec{r}) = \arg\left[e^{i\vec{k}\cdot\vec{r}}\right] = \left(\vec{k}\cdot\vec{r} \bmod 2\pi\right) - \pi$. The resultant pattern is known as blazed grating and is shown in Fig. 2.3**a**. In an experiment, a blazed grating will of course result in a laser dot of finite size since the incoming wave E_0 will have a Gaussian-beam lateral structure and because of point-spreading properties of the optical system. More complicated patterns, however, require complex algorithmic approaches (see Sec. 2.2.5 and chapter 4).

In practice, it may be simpler to display a Fourier frequency $\Delta\varphi(\vec{r}) = \pi \mathrm{Re}\left[e^{i\vec{k}\cdot\vec{r}}\right]$, instead of a blazed grating as shown in Fig. 2.3**b**. Even though this is a crude approximation and Fourier gratings lead to undesired off-target traps, this may not be a limiting factor if the angle to the zeroth order is high enough, such that these beams do not enter the field of view. This point is discussed in further detail in Sec. 2.2.3. The examples in Fig. 2.2**b** show that the spatial frequency of a Fourier grating still determines the position of the trap.

Shape-phase holography

Applications such as optical tweezing often require more complex trap shapes than single points, such as line or circular traps. The method of shape-phase holography [118] is tailored to producing such simple patterns. While the spatial frequency displayed on the SLM (see Fig. 2.3) determines the centre position of the trap, the outline of the area on the SLM displaying this spatial frequency can be effectively used as an aperture. The aperture is not real in the sense that no rays are blocked by it. Rays that impinge on the SLM outside of the aperture region, simply get diffracted into a different direction.

In order to simulate the effect of a Gaussian-shaped aperture overlaying the spatial frequency as shown in Figs. 2.4**a**-**b** and **d**-**f**, I multiply the unitary phase-factor $e^{i\Delta\varphi(r_x,r_y)}$ with a real-valued Gaussian despite the SLM being a phase-only device. The result reads,

$$p_{\mathrm{SP}}(r_x,r_y) \propto e^{-\frac{(rx-\mu_x^2)}{2\sigma_x^2}-\frac{(r_y-\mu_y^2)}{2\sigma_y^2}} e^{i\arg\left[e^{i(r_x\frac{2\pi}{l_x}+r_y\frac{2\pi}{l_y})}\right]} \tag{2.6}$$

where SP stands for "shape-phase". The convolution is turned into a multiplication by the Fourier transformation so that I obtain

$$E_{\mathrm{cam}} \propto e^{ir_x\mu_x+ir_y\mu_y} e^{-r_x^2\frac{\sigma_x^2}{2}} e^{-r_y^2\frac{\sigma_y^2}{2}} \delta\left(\frac{2\pi}{l_x}\kappa+r_x\right)\delta\left(\frac{2\pi}{l_y}\kappa+r_y\right), \tag{2.7}$$

where unitary factors $e^{i\cdots}$ can be disregarded as far as the intensity is concerned. The entire result is pathological because I assumed a flat lateral structure of the incoming beam, but the logic is perhaps clear: Shape-phase methods allow holographers to shape the light intensity located around the position that corresponds to the spatial frequency $(1/l_x, 1/l_y)$.

I give six examples of shape-phase holography in Fig. 2.4. The holograms needed for most of the work in this thesis make heavy use of shape-phase methods, for instance to create line-shaped traps (see Fig. 2.4**d** - **f**).

The *Red Tweezers* program is the holography engine of choice to create phase-shaped light patterns [123]. The program leverages the computational power of a Graphical Processing Unit (GPU) to achieve real-time recalculation of the hologram required to steer independently addressable point- and line traps to desired places. Importantly, the program also offers a straightforward control of the location of the envelope on the SLM, i.e. parameters μ_x, μ_y in Eq. (2.7), which, even though they do not influence the intensity in the focal plane (see Fig. 2.8), have a measurable effect on tweezing forces.

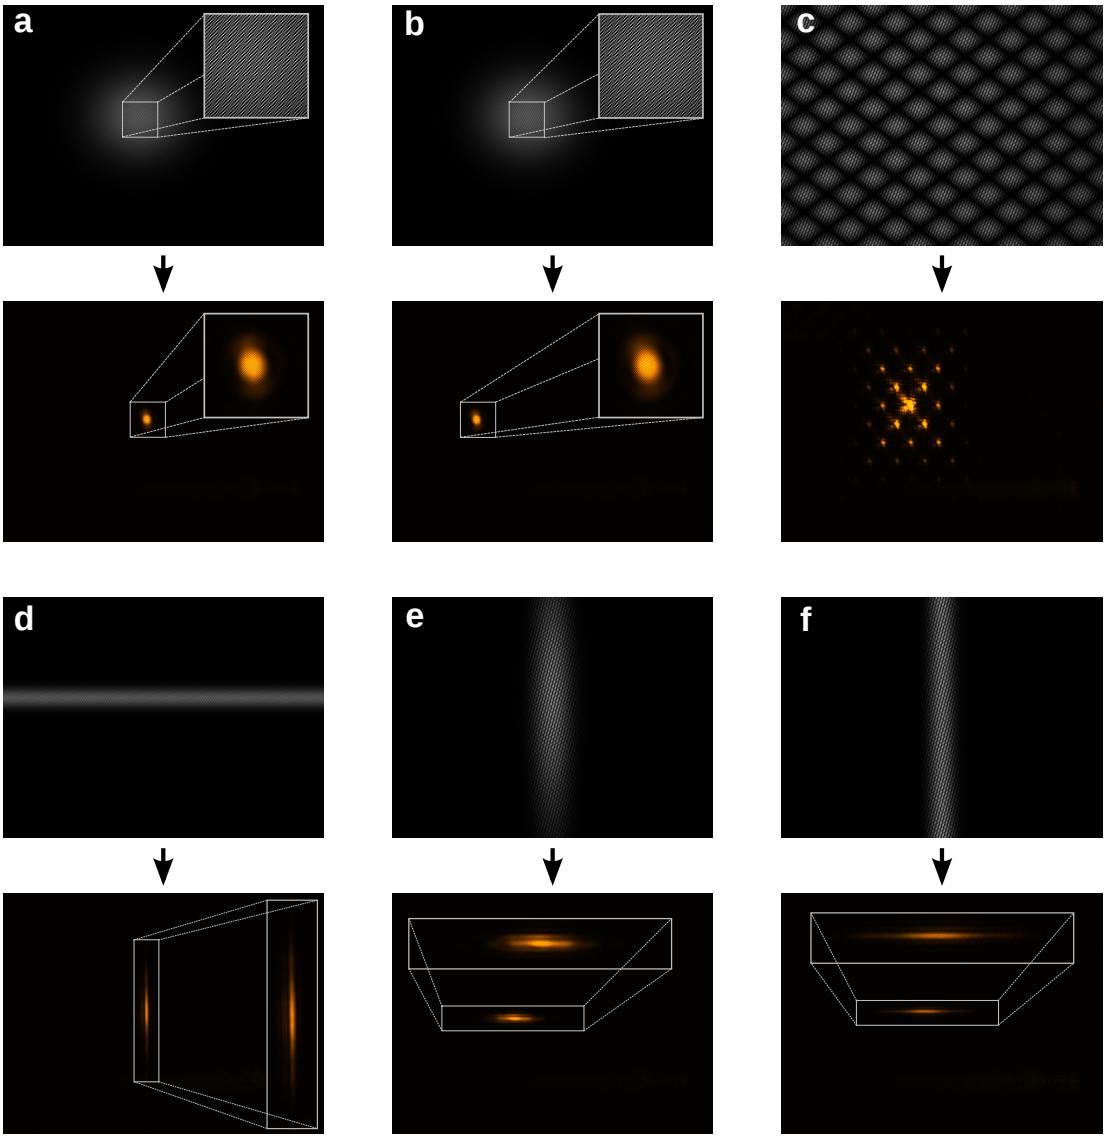

Fig. 2.4 Experimental examples of shape-phase holography. Each panel shows the 800 × 600-sized hologram in grey (top) and the resultant intensity distribution in red (bottom). The intensity distribution was recorded on CMOS 3 on 1024 × 768 pixels (see Fig. 2.6). The white-framed box in the top-right corner in each picture shows a magnification of the hologram and laser pattern. With the exception of panel **c**, all holograms were calculated using Eq. (2.6): A sinusoidal pattern with a selected spatial frequency is multiplied with a slowly-evolving Gaussian. The fine ripples constituting the spatial frequency in the holograms may not be visible in print. **a** Isotropic Gaussian as envelope aperture. **b** Isotropic Gaussian as envelope aperture at same position, but shaping a different spatial frequency. **c** Holograms with repetitive symmetry result in ghost traps. **d** Elliptic Gaussian with one standard deviation sent (practically) to infinity, results in a line-shaped trap. **e** Less extreme elliptic Gaussian, resulting in an elliptic intensity distribution elongated along the thin direction of the Gaussian. **f** Line trap, but orthogonally to panel **d**.

Phase-gradients in elongated optical traps

In this section, I introduce a crucial effect which is of particular importance in elongated line traps produced by holographic optical tweezers: phase-gradients [124].

In the small-angle approximation of ray optics, the effect of a lens can be described by $\mathbf{M}\vec{e}$ where $\mathbf{M}$ is a matrix and $\vec{e} = (x,\theta)^T$ denotes a vector composed of the lateral offset x and the angle θ of the ray with respect to the optical centreline. In a 2f-configuration, the action of the lens reads [121]

$$\begin{pmatrix} 0 & f \\ -1/f & 0 \end{pmatrix} \begin{pmatrix} x \\ \theta \end{pmatrix} = \begin{pmatrix} f\theta \\ -x/f \end{pmatrix}. \tag{2.8}$$

In other words, a lateral offset to the optical centreline of an incoming ray translates into an outgoing angle and vice-versa. The offset position of the aperture envelope in the previous section (μ_x, μ_y) therefore controls the angle of the laser beam when it leaves the objective (within the realm of applicability of this approximation).

Since the scattering force described in Eq. (2.2) is poynting in the direction of beam propagation, it can be used to exert forces with vector components that lie in the image plane if the angle of incidence is sufficiently high.

Roichman *et al* [124] provide a quantitative expression for the in-plane component of the force

$$\vec{F}_{\text{phase-grad}}(\vec{r}) = \frac{\kappa}{n_m \mu_r \mu_0 c} I(\vec{r}) \vec{\nabla}\varphi(\vec{r}), \tag{2.9}$$

where φ denotes the in-plane phase profile (not to be confused with the SLM phase shift). Since phase-gradient forces ultimately result from scattering, they can be shown to be non-conservative. Indeed, as Roichman *et al* demonstrate $\vec{\nabla} \times \vec{F}_{\text{phase-grad}} \propto \vec{\nabla} I \times \vec{\nabla}\varphi \neq \vec{0}$.

While scattering forces in optical traps can be significant, if not dominant, line-trap phase-gradients typically only consists of a small in-plane fraction which can be readily controlled by adjusting μ_x, μ_y. In this thesis, I use phase-gradient forces to nudge Brownian particles along line-traps where both the intensity I and the phase-gradient $\vec{\nabla}\varphi$ stay constant along the trap. As I show in chapter 3, the resultant forces are on the order of a few femtoNewton and thus in the right regime for my experiments.

2.2.2 Principles of phase-only spatial light modulators

The functionality of spatial light modulators derives from birefringent liquid crystals that are reactive towards externally applied electrical fields and sit in a sandwich-structure

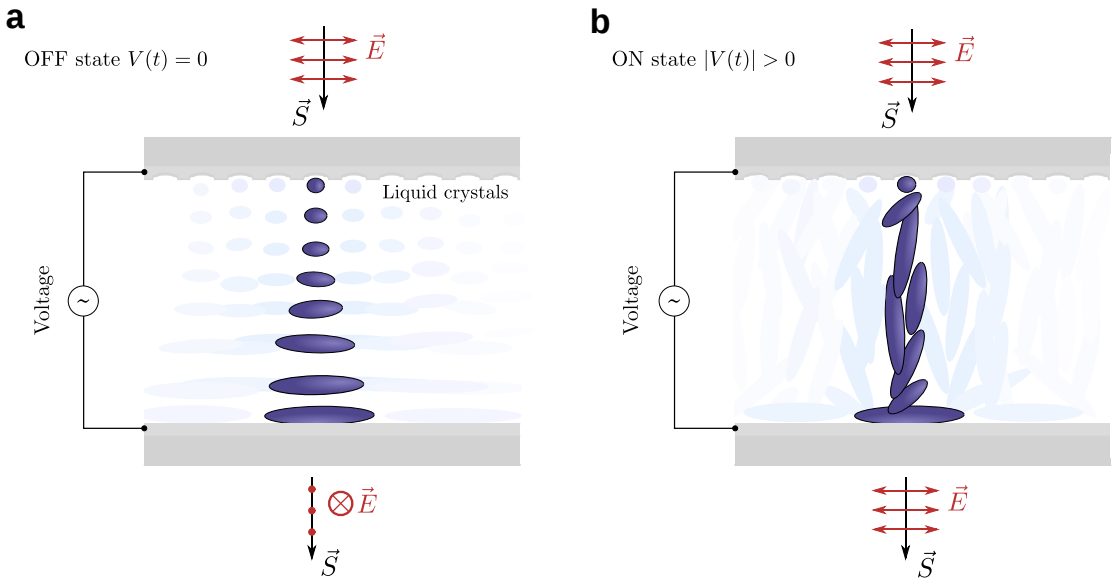

Fig. 2.5 Sketch of the functionality of SLMs, adapted from A. Jesacher's PhD thesis [122]. Liquid crystals are sandwiched between two electrodes with orientation grooves forcing the crystals near the electrodes into a particular orientation. **a** OFF-state in which liquid crystals reorient the $\vec{E}$-field vector of incoming light rays. **b** ON-state in which liquid crystals orient along the externally field and consequently do not rotate the field polarisation.

between two orthogonal polarizers as shown in Fig. 2.5. In the absence of electrial fields (OFF), the crystals are orientated in a helical structure (twisted nematic) such that they turn the field vector of incoming light 90°. By contrast, under the influence of a certain voltage, the crystals align along the field direction, loosing their polarization-changing properties (ON), which leads to incoming light getting attenuated at the second polarizer.

Crucially, the liquid crystals directly act on polarization, not the phase. The reason why SLMs can modulate the phase is the *Pancharatnam-Berry-phase* or *geometric phase* effect [122]. The SLM changes and then restores the polarization of the incoming field, so that the field vector undergoes a certain amount of rotation in between. The amount of rotation can be put in direct correspondence with the resultant phase change $\Delta\varphi$ [122].

Phase-only SLMs have the advantage that they do not need to absorb any of the incoming radiation. This reduces radiation-induced heating and thus higher laser power can be used.

A drawback of this technical approach is the inherent rate-limitation of the SLM due to relaxation and alignment times of liquid crystals. Moreover, SLM-generated light fields may suffer from flickering since the voltage applied is alternating to prevent electrolysis of

the LC-solution. The frame rate of the SLM used in this thesis (LCOS X10468, Hamamatsu) is fixed at 60 fps.

2.2.3 Holographic optical tweezers setup

Fourier optics require lenses placed in 2f-configurations before and after the SLM such that the optical transformation described in the previous chapter applies. The four lenses L1-L4 in Fig. 2.6 and Tbl. 2.1 are placed in a 4f-configuration, sandwiching the SLM. The focal lengths $f_{L1} - f_{L4}$ were chosen to (1) overfill the active area of the SLM such that the incoming intensity profile is as laterally flat as possible, (2) overfill the back aperture of the objective in order for gradient tweezing forces to be able to overcome radiation pressure [125], and (3) ensure the design can be physically realized. Short focal lengths, for instance, make it difficult to accomodate all required optical elements, such as dichroic mirrors, in the beam path.

Together with my colleague, Y. Tan, I redesigned the layout of the HOT setup during the course of my PhD. The setup had to be realigned from scratch, the 4f-lens system had to be changed, and several components, including the mount of the microscope stage had to be manufactured to custom dimensions. Crucial differences to the previous version of the HOT include (1) a new CMOS Camera (Microtron 1362) supporting vastly higher framerates (up to $\sim 180k$ fps), (2) a larger fill factor of the SLM by the incoming laser beam, and (3) an angular offset of the zeroth order diffraction of the SLM. The latter point is relevant in so far as the previous version of the setup suffered from the problem that the zeroth diffraction order, even though it was partially occluded, entered the field of view in the microscope in a central location, heating the sample, distorting other optical traps, and resulting in undesired tweezing.

The drawback of the new setup design is that the spatial frequencies of all holograms need to be offset by a fixed amount in order for the diffracted light to enter the microscope. This disadvantage, however, is also an advantage: the angular distance between adjacent higher diffraction orders in the new design is just large enough such that they are usually not diffracted into the angular range of the pinhole $P3$ in Fig. 2.6 and hence do not disturb the experiment.

The Ytterbium fiber laser in Fig. 2.1 has many desirable properties from a tweezing perspective, since it does not suffer from beam pointing instabilities and only minor power fluctuations ($\pm$2% [126]). Its output mode is a TEM_{00} and it has hence a Gaussian intensity profile with an output $1/e^2 - \varnothing$ of 5 mm ($1/e - \varnothing = 2.5$mm) downstream from the fiber collimation optics (not shown; part of the Laser-symbol in Fig. 2.6).

Abbreviation	Optical Element	Property
M1- M4, M7, M10	AR-coated dielectric mirror EO3	$\varnothing = 1$”
M5, M6	AR-coated dielectric mirror E03	$\varnothing = 2$”
M8	Silver elliptical mirror	$\varnothing = 1$”
P1	Zero-order $\lambda/2$ wave plate	$\lambda = 1064$ nm
P2	Pol. beamsplitter cube	$\lambda = 900 - 1300$ nm
P3	Pinhole	
L1	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 30$mm
L2	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 150$mm
L3	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 500$mm
L4	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 200$mm
L5	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 150$mm
L6	Plano-conv. lens	$\varnothing = 1$”, $f = 500$mm
L7	Plano-conv. lens	$\varnothing = 1$”, $f = 50$mm
L8	Plano-conv. lens	$\varnothing = 1$”, $f = 30$mm
L9	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 200$mm
L10	Plano-conv. lens (IR AR-coated)	$\varnothing = 1$”, $f = 100$mm
L12	Air objective	M= 20X
P4, P5	Pellicle beamsplitter	$R/T = 8/92$ $\lambda = 0.4 - 2.4\,\mu$m
D1	Lowpass dichroic	$\varnothing = 2$”, $\lambda_{\text{cutoff}} = 1000$ nm
D2	Highpass dichroic	$\varnothing = 1$”, $\lambda_{\text{cutoff}} = 505$ nm
F1	UVFS hot mirror	$\varnothing = 1$”
F2	IR bandpass	$\varnothing = 1$”
F3	IR bandpass	$\varnothing = 1$”
CMOS 1	Mikrotron EoSens Cl (MC1362)	1280×1024 pix
CMOS 2	Imagingsource, USB 2.0 monoch.	1280×1024 pix
CMOS 3	Imagingsource, USB 2.0 colour	1280×1024 pix
Objective	Olympus UPLSAPO 100XO	M= 100X, $f = 1.8$mm back foc. pl. $\varnothing = 3.78$mm
Piezo stage	PI P-561 (stage), E-754 (control)	Range (x, y, z): 100μm
SLM	LCOS X10468, Hamamatsu	800×600 pixel, 60 fps
Laser	YLM-5- 1064-LP, IPG Photonics	$P = 0 - 5$ W, $\lambda =$1064 nm

Table 2.1 Abbreviations and properties of optical elements in Fig. 2.6.

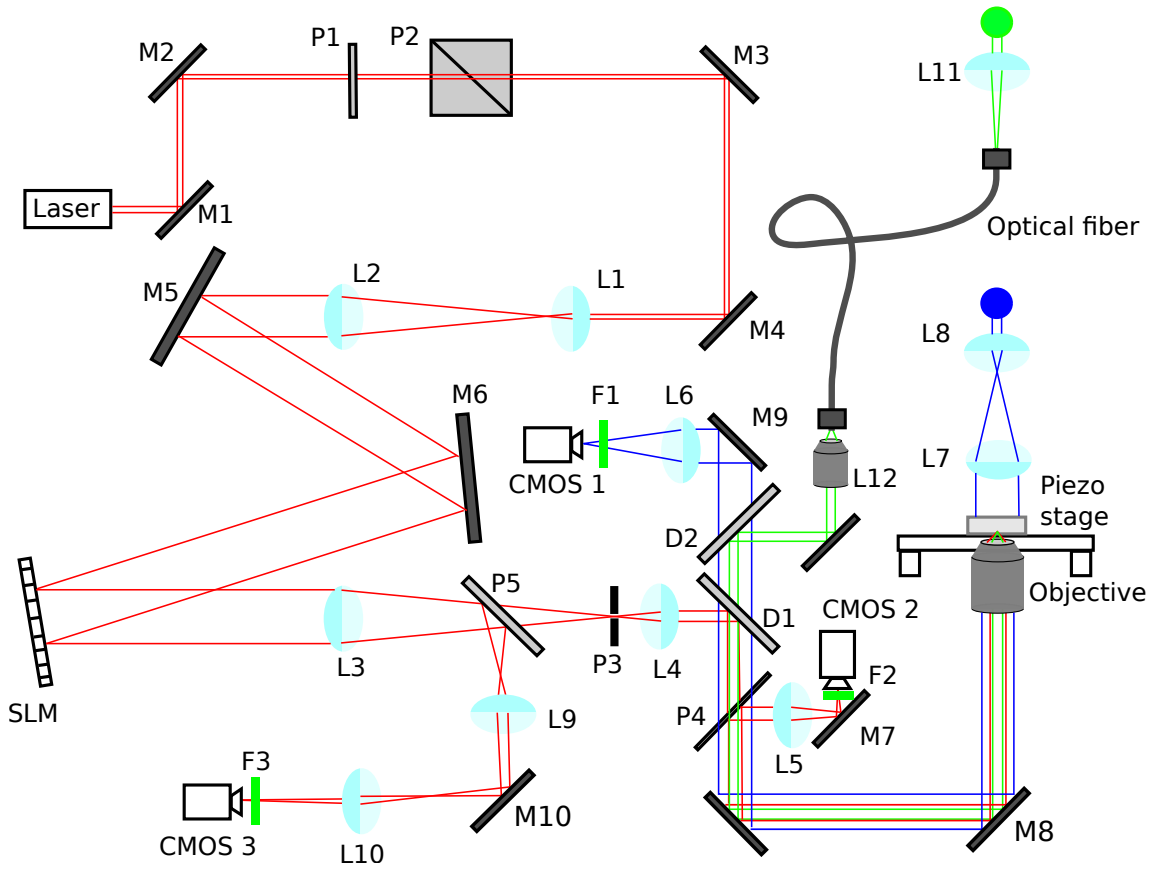

Fig. 2.6 Detailed sketch of all optical components in the HOT setup. Updated sketch of the setup from [127].

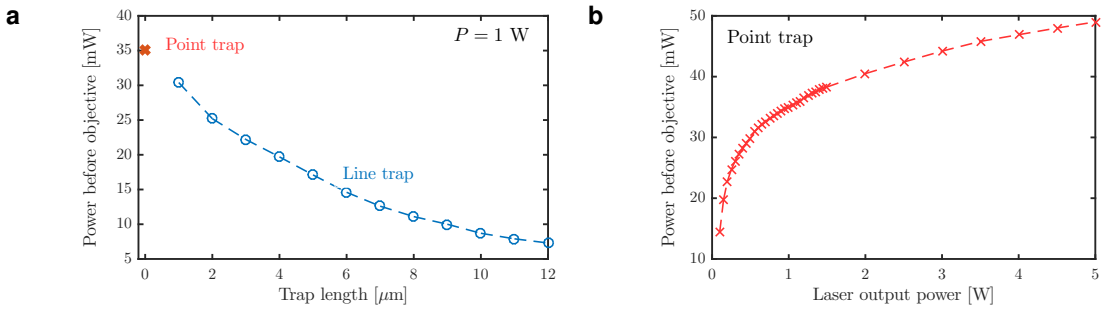

Fig. 2.7 Laser intensity before entering the objective measured using a power meter. **a** Total power allocated to point and line traps of various lengths. **b** Non-linear dependence of laser intensity entering the objective aperture and total laser power.

The relay optics $L1-L2$ then increases the $1/e^2$-diameter by a factor of $f_{L2}/f_{L1} = 5$ to 25 mm. The active area of the SLM has dimensions 16×12 mm$\times$ mm. This means that we are overfilling the active area roughly twice in the vertical direction with the $1/e^2$-diameter, providing a slowly varying intensity across all active pixels. In order to achieve stable tweezing, the gradient forces Eq. (2.1) have to overcome the scattering forces in all directions, especially in the direction of beam propagation. This is typically achieved by overfilling the back aperture of the objective [125] with the incoming laser. The rays that provide the restoring force in the axial direction originate from the rim of the back aperture, as indicated in Fig. 2.1**b**. In the case of holographic tweezers, 'overfilling the back aperture' requires the image of the active area of the SLM to be greater by some margin than the back aperture diameter of the objective (listed in Tab. 2.1). If, on the other hand, the active area of the SLM is magnified too much, the loss of power and imaged pixel area may become prohibitive.

In Fig. 2.7, I plot the laser power that enters the back aperture of the objective as a function of the hologram (panel **a**) and the total input power (panel **b**). The non-linear dependence of laser intensity before the objective is (partially) caused by overfilling the SLM, since with increasing laser power a growing fraction of the intensity hits the SLM outside of the active area.

The telescope $L3-L4$ in my setup was chosen to yield a magnification of $f_{L4}/f_{L3} = 2/5$ such that the image of the central square of the active area would have dimensions 4.8×4.8 mm $\times$ mm when it enters the back focal plane of the objective which has a diameter of $\varnothing = 3.78$ mm, overfilling by approximately 1 mm radially.

On the imaging side, I use a tube lens (L6) with a focal length of 500 mm, which results in a total magnification of $f_{L6}/f_{\text{Objective}} \approx 277$ on the CMOS 1 camera. The pixel size of the

Microtron is 14^2 μm $\times$ μm so that a pixel-length corresponds to 50.5 nm in the imaged plane. For comparison, the Abbe criterion for, say, $\lambda = 560$ nm gives a minimum resolvable distance of 200 nm (NA= 1.4). The corresponding spatial Nyquist frequency thus amounts to 100 nm. This means that the microscope/camera-system oversamples by a factor of approximately 2.

The total sensor area of the LUPA1300-2 as provided by Mikrotron [128], 17.92 $\times$ 14.34 mm$\times$ mm divided by the pixel number 1280 $\times$ 1024 amounts to precisely 14^2 μm $\times$ μm. The conversion factor 1 pix = 50.5 nm calculated above is important and is used throughout this thesis. It is also a mandatory input parameter of the internal holography engine of the RedTweezers programe. The excellent agreement of the RedTweezers-physical model used to create holograms for given trap locations with the observed trap-location corroborates this number. Even just slightly different input parameters lead to a notable divergence of the observed and predicted trap locations. The same holds for the image size of the active area of the SLM 4.8^2 mm $\times$ mm, which, too, is a required input parameter of the holography engine.

Sample illumination is provided by two LEDs, one in trans (depicted as blue circle in Fig. 2.6) and one in epi-configuration (green circle). The role of the latter is to provide excitation for potential future fluorescence-based experiments. The LED in trans configuration is easily exchanged; I used a white light LED in all experiments described here in order to pass the dichroic mirros (D1, D2).

As shown in Fig. 2.6, the setup is endowed with two further CMOS cameras imaging planes that are respectively conjugate to the microscope plane. While CMOS 2 is used to image the back reflection coming from the sample, e.g. to study phase-gradients over different axial positions, CMOS 3 is used to map the created light landscape in a quantitative fashion. Because of this, CMOS 3 is of primary importance in my inverse holography project, the underlying physics of which I describe in Sec. 2.2.5.

2.2.4 Particle localisation routine

The core innovation of this thesis lies in the software development on the HOT setup. The control software my colleagues and I use to carry out Brownian dynamics experiments is an extension of the RedTweezers holographic control program mentioned in the previous section. I extended or refined the software with (1) a robust autofocus routine, (2) a realtime colloid localisation, (3) a dynamic voltage control with localisation-feedback, (4) a laser-control, and (5) a robust automation routine with localisation-feedback. The automation routine is capable of performing experiments in parallel and unsupervised, automatically cycling through a number of parameters and detecting as well as reacting

to error-states. In the appendix in Sec. A.1, I give a detailed account of the core algorithm of the automation routine.

In addition, I optimised the performance of the software such that it can analyse each individual frame at frame rates of up to around 2000 fps without missing any frames coming from the Mikrotron (CMOS 1). Upon arrival, images are stored into an image queue in a dedicated loop, which has no other tasks. Another loop, which is carried out concomitantly then extracts and processes each frame from the queue and stores the data into a binary format (tdms). Thus, frames are not stored on the harddrive; only the extracted information is. However, the user still has the option to record videos at all times. At the end of every recording, all experimental parameters such as trap positions or laser power are also stored in the file.

In the following I discuss a few core algorithmic components of the localisation routine. The list below contains a step-by-step explanation of localisation:

- Before experiments: Gather $\sim$ 500 images and average.
- Drag colloid into ROI and begin experiment.
- For each frame, subtract average frame and threshold the result to obtain a binary image.
- Extract all regions where pixels are one (above threshold).
- Extract the original pixel-values in these regions from the frame.
- Optional: Perform a Gaussian blurring (kernel size is user controlled).
- Perform either (1) centroid or (2) self-convolution algorithm [129] to find centre of extracted region (colloid).

2.2.5 Digital holography as an inverse problem

As I indicated throughout this chapter, a given intensity distribution in the focal plane will have several corresponding holograms that would create it. The associated algorithmic problem of finding any or even all phase-patterns that give rise to a given intensity distribution is classical and has inspired numerous publications in various contexts, notably diffraction-based techniques such as electron microscopy [130, 131, 122].

The phase-problem encountered in shape-phase holography is illustrated in Fig. 2.8, where two holograms with different envelope locations but similar spatial frequencies are shown that give rise to approximately identical intensity distributions. Differences

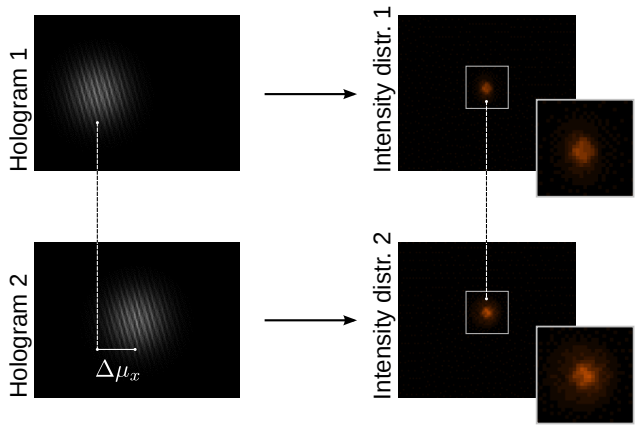

Fig. 2.8 The phase problem in shape-phase holography. Two holograms displaying the same spatial frequency, that are overlaid with different Gaussian apertures, separated by a horizontal distance $\Delta\mu_x$, result in very similar intensity distributions.

between the two intensity distributions mostly arise from aberrations and spatial variation in illumination intensity of the SLM (see Sec. 2.6).

However, the phase-problem arises in many more ways in holography.

As I want to briefly discuss here, phase-invariances also appear in restricted holograms containing only a few or even just a single spatial frequency. In Fig. 2.9**a**, I demonstrate that the intensity of a point-trap resulting from a hologram that contains only a single spatial-frequency is a constant over the complex phase φ of that frequency. The holograms shown in this case are the result of taking the real part of a Fourier transformation of complex-valued 8×8-matrices $\mathbf{f}$ (the same matrices are shown in Fig. 2.2**b**). Moreover, as shown in Fig. 2.9**b**, if I vary the amplitude r of the same Fourier component instead of its phase, the intensity of the spot describes a non-invertible bell-shaped curve. The bell shape is caused by the circular nature of phase-shifts: From a certain shift value (Fourier amplitude) onwards, the intensity will adopt a previously held value.

The red curve in panel **b** is the mean over a collation of 20 measurements at the same Fourier amplitudes, the grey envelope shows the standard deviation. The data was recorded at a laser output of $P = 0.3$ W on CMOS3 (see Fig. 2.6) with an exposure time of 10 ms and a gain of $8x$.

The phase invariances in both panels **a**, **b** are shown for a single spatial frequency and will likely combinatorically multiply if holograms consisting of several Fourier frequencies are considered.

The plot also introduces an important concept: Instead of thinking about the relationship S between a hologram $\mathbf{h}$ and the intensity distribution $\mathbf{I}$, one could consider the

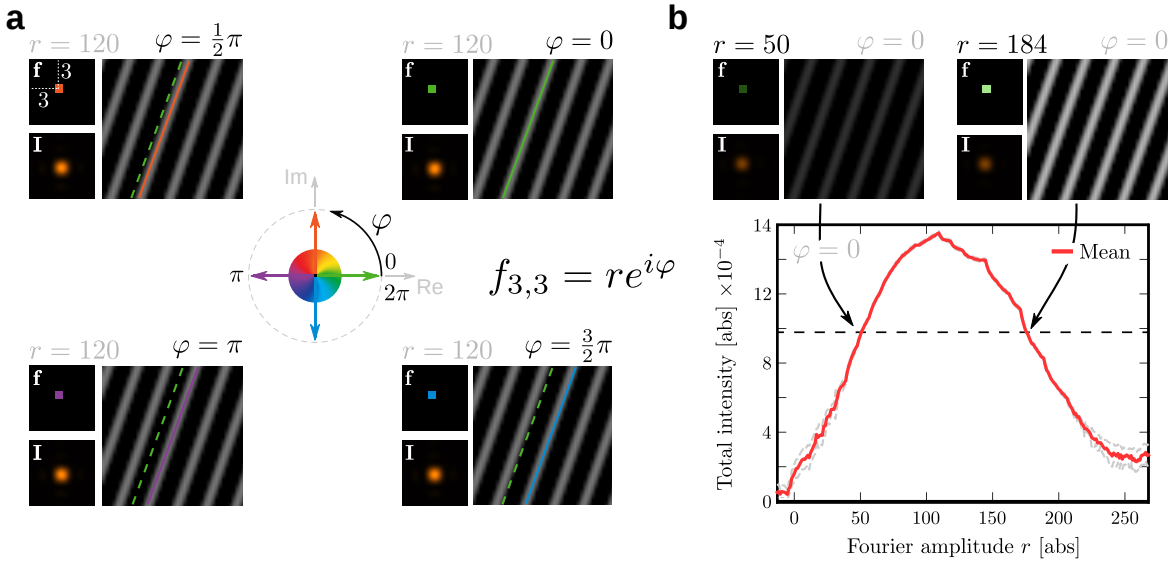

Fig. 2.9 Phase-intensity invariances for a single Fourier component, here exemplified using $f_{3,3}$ from a preselected 8×8-Fourier matrix **f**. **a** Variation of phase: The measured intensity is invariant over the phase $\varphi \in [0, 2\pi)$ of the complex-valued Fourier component 3,3 (colour coded). The invariant laser intensities **I** are shown below the Fourier matrices. **b** Variation of amplitude: The intensity is a non-invertible function of the amplitude $r \in \{0, 1, \ldots, 255\}$ of the Fourier component.

mapping between the Fourier frequencies **f** that underlie the hologram **h** and the ensuing intensity distribution **I**, i.e. $S \circ F : \mathbf{f} \to \mathbf{I}$. I go into greater detail on this in chapter 4.

As mentioned before, examples of the two mappings S, F are shown in Fig. 2.2**b**. Crucially, the dimensionality of the space of relevant Fourier frequencies is much smaller than holograms such that this space is more amenable to numerical treatment.

Gerchberg-Saxton algorithm

Arguably, the best-known phase-retrieval algorithm is the Gerchberg-Saxton [130], summarised in Fig. 2.10: Initially, the phase is set to random. It is paired with the target intensity distribution in a field of complex numbers, where intensities are recorded as amplitudes. A sequence of back-and-forth Fourier transformations, during which the amplitudes are replaced by the target intensity and constant matrices then converge the phase-field to the final pattern. I give two examples in Fig. 2.10**b** and **c**. Algorithmic phase-retrieval is subject of ongoing research and several variants of the original GS algorithm, geared towards their respective application, have been proposed in the literature [132–134].

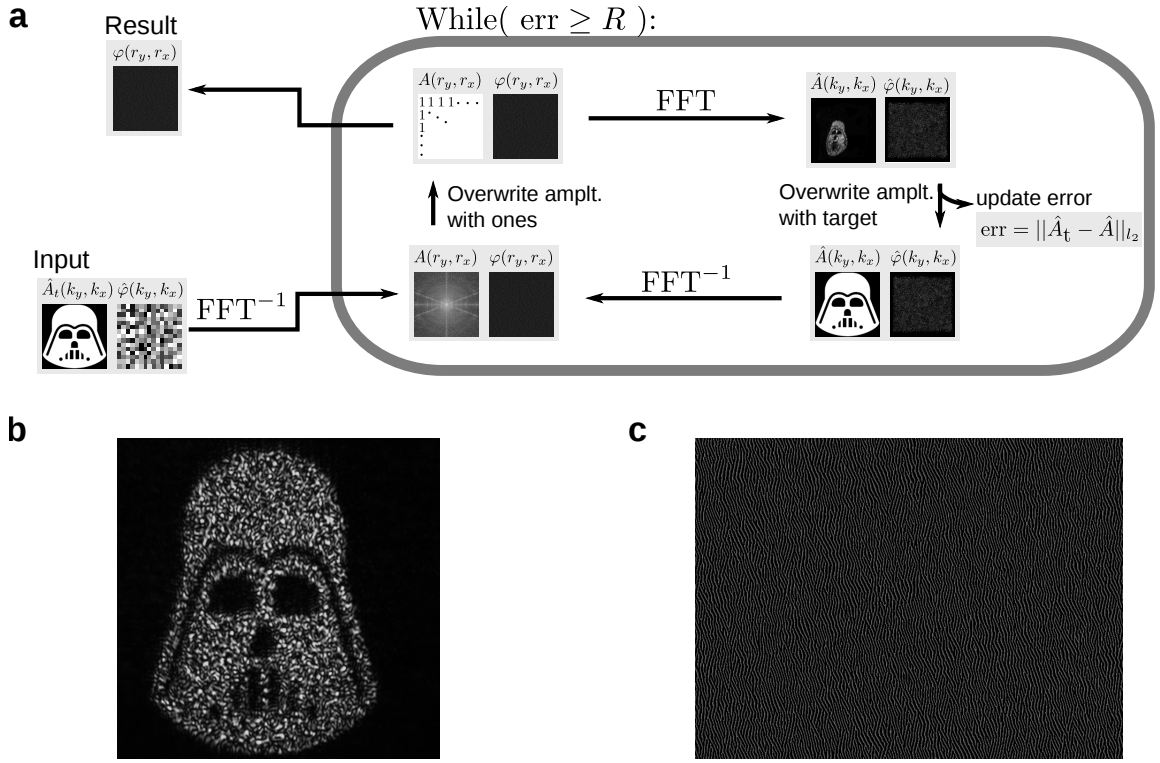

Fig. 2.10 Gerchberg-Saxton Algorithm. **a** Sketch of the algorithm. Grey boxes represent matrices of complex numbers $z = Ae^{i\varphi}$. **b** Example of an intensity distribution created with the GS algorithm on the HOT setup (measured on CMOS 3). **c** The corresponding SLM pattern.

Chapter 3

Results of path-time experiments & simulations

This chapter presents my main experiments on various flavours of first-passage times carried out on the HOT setup described in Chapter 2. In addition, I summarise results from Brownian dynamics simulations used to study selectivity in channel transport. The results presented here have been published at the time of writing in at least one paper [94] about transition-path times in non-equilibrium systems (Sec. 3.1). A second publication about optimised Brownian escape times (Sec. 3.2) has been submitted. Efforts to measure and publish first-passage times in corrugated channels, which offer a perspective on entropic forces, have seized following a very similar publication by a competing group [135].

3.1 Experimental evidence of path-time symmetry and its breakdown

The following section contains the results of the aforementioned publication which I authored [94]. Whenever data from my coauthors is considered, I indicate this in the following text.

Classical thermally activated reactions are ubiquitous in nature and technology with a wide range of examples including folding transitions of proteins [136, 62, 137, 65] and DNA [138, 72], transitions of colloidal particles between optical traps [139], and the dynamics of molecules in membrane channel proteins [140–142], artificial nanopores [143, 144], and channels [135, 145]. Depending on the system, these transitions may proceed along a single, multiple, or even a continuum of pathways in phase-space that connect the initial and final states. I mentioned the debate around the question of pathway multiplicity in

Sec. 1.1.2. In Sec. 1.2.2, I defined the transition-path time τ as the time it takes to travel from one thermodynamic state to another.

From the definition, it follows that in a two-state system, the transition-path time $\tau_{\mathrm{I}\to\mathrm{II}}$ can be measured as follows: whenever the system leaves the area of state I, a stopwatch is triggered. It is stopped when the system either returns or transitions to state II. In the latter case, this constitutes a single realisation of the transition path (see Fig. 1.2).

Importantly, as I explain in Sec. 1.2.3, transition-path times τ do not directly determine the rates $k_{\mathrm{I}\to\mathrm{II}}$ and $k_{\mathrm{I}\leftarrow\mathrm{II}}$ of the reaction [99, 146]. This is because, rate coefficients account for all prior unsuccessful attempts at leaving the state in addition to the actual time of travel of the successful attempt. Transition rates can therefore be strongly asymmetric if one state is thermodynamically favourable over the other. By contrast, transition-path times are expected to be statistically symmetric in equilibrium, in accord with the principle of microscopic reversibility [147, 91] (see Sec. 1.1.3).

In many systems, it has been challenging to resolve individual transitions. However, technological progress has now advanced to a point where information about folding events of polymers can be gathered using optical techniques, such as FRET, which has sparked considerable interest in transition-path times [62, 137, 65], as mentioned in Sec. 1.1.2. In electrophysiology and pulse-sensing experiments, molecules are interrogated by voltage-driven transport often proceeding along a single pathway through membrane channels [148]. This technique detects changes in ion flow due to blockage by solutes of interest. However, for small, uncharged molecules (such as some antibiotics), these measurements are often not sensitive to the direction of travel or the orientation of the channel [149]. Direction, however, matters in biological membrane channels, which usually have an asymmetric structure. Channel asymmetry can, for instance, give rise to ratcheting effects that rectify diffusive currents under the influence of fluctuating forces [150]. Despite their importance, most of the thermodynamic principles of transition-path times have not been studied systematically, especially in experiments outside of thermodynamic equilibrium.

3.1.1 Uphill and downhill exit-path-time symmetry

Here, I use the HOT setup in conjunction with confining microchannels to physically simulate the escape of a Brownian particle from a cavity, reminiscent of the escape of solutes from membrane channels [151]. The movements of solutes such as ions in membrane channels or nanopores often follow thermodynamic gradients. Such a gradient is modelled here with a phase-gradient force f in a laser line trap as sketched in Fig. 3.1**a**. I plot a selection of trajectories of Brownian particles acquired from automated drag-and-drop

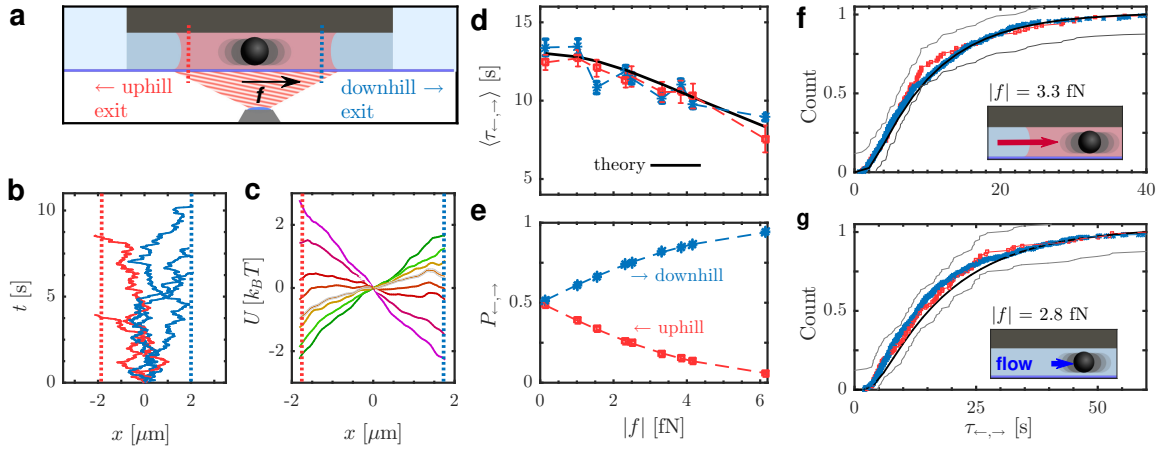

Fig. 3.1 Exit-path times in the uphill and downhill direction are identically distributed. **a** Exit-path times of a colloidal particle from a predefined interval are measured for various phase-gradient forces $f \approx 0-6$ fN. The sketch shows a particle in a microfluidic channel, confined from above by PDMS and from below by the cover slip. The laser enters the channel through the cover slip and forms a potential landscape for the colloid inside the channel. Both exists, down- and uphill, from the interval considered in the other panels are marked by red- and blue-dotted lines respectively. **b** Examples of five trajectories and exit-path times. **c** Inferred potentials (see Sec. 3.1.6). The standard error of the mean is given by the line width. **d** Mean exit-path times in the uphill (red) and downhill (blue) directions agree, but decrease with increasing force. Error bars indicate the standard error of the mean. **e** Probability to exit the interval in the uphill (red) and downhill (blue) direction. Error bars (on the order marker size) show the standard error of the mean. **f** Cumulative distribution of first exit-path times in the uphill (red) and downhill (blue) direction for the 3.3 fN-data point and boundaries of the Kolmogorov-Smirnov test (light grey) (see Sec. 3.1.7). The exit times were measured from a predefined interval similar to panel **a** (dotted lines). While the lines represent the full data set, the markers only show a selection of the data to avoid crowding of the figure. The black line in **d** and **f** show Fokker-Planck-based predictions as discussed in the text. **g** As previous panel, but with pressure-driven flow as force.

experiments in panel 3.1**b**: At $t = 0$, the particle is positioned at the centre $x_0 = 0$ of a predefined spatial interval within a microchannel. The particle is released, a stopwatch is triggered, and a laser line trap with a prescribed phase-gradient is turned on (see Fig. 1.2**a** and Sec. 3.1.6). Once the particle leaves the interval, the stopwatch and the measurement are stopped. I mention here again, that a more detailed description of the automation routine can be found in the appendix in Sec. A.1.

I find that the probability density of positions $\rho(x)$ recorded in an ensemble of repeats goes to zero at the interval boundaries $x_{\leftarrow}$, $x_{\rightarrow}$. These points can therefore be considered as absorbing in the Fokker-Planck picture. For each value of force f, I gathered around 1000 trajectories of which $950-980$ were free of incidents such as other particles entering the channel. In order to check for static bias, I applied both positive and negative forces (see Fig. 3.1**c**). Such a bias is often caused for example by weak latent flows of water. Within experimental resolution, I find that the diffusion profile $D(x)$ along the channel is not affected by the applied phase-gradients (see Sec. 2.2.1 and Sec. 3.1.6), such that any difference in dynamics must be attributed to the difference in force. The inference method used to estimate potential $U(x)$ and diffusion profile $D(x)$ is discussed in Sec. 3.1.6.

The central result of this experiment is the equivalence of the mean left and right exit-path times $\langle\tau_{\leftarrow}\rangle$, $\langle\tau_{\rightarrow}\rangle$, shown in Fig. 3.1**d**. Moreover, shorter exit-path times for higher absolute forces $|f|$ indicate a speed-up of both uphill and downhill trajectories. By contrast, exit probabilities behave intuitively: exits against the force (uphill direction) become increasingly unlikely with an increasing force magnitude $|f|$ as shown in Fig. 3.1**e**.

The theoretical mean exit-path time (black line in Fig. 3.1**d**) was obtained from a solution of Eq. (1.36) with boundaries at $x_{\leftarrow} = -L/2$ and $x_{\rightarrow} = L/2$ with $L = 3.7$ µm. I want to note here that L stands for the length of the interval, not the channel. Due to the observed exit time symmetry, the equation for the mean exit time can be solved for any exit $x_{\leftarrow}$, $x_{\rightarrow}$. Interestingly, this symmetry extends beyond a simple equivalence of the mean. In fact, the distribution of exit-path times agree in the uphill and downhill direction as shown in Fig. 3.1**f**. This holds even for forces of different nature such as hydrodynamic drag as shown in panel **g**. Trajectories that manage to exit against the flow do so at precisely the same drift speed $\langle\dot{x}\rangle$ as the ones that follow the flow. The statistical significance of similarity between two given cumulative distributions is asserted by the Kolmogorov-Smirnov test (see Sec. 3.1.7). Throughout this study I require a significance of 0.5. Theoretical distributions of exit-path times shown in Fig. 3.1**f**,**g** to, for example, the left side $\rho_{\tau_{\leftarrow}}(t)$ were obtained from a numerical solution of the Fokker-Planck equation

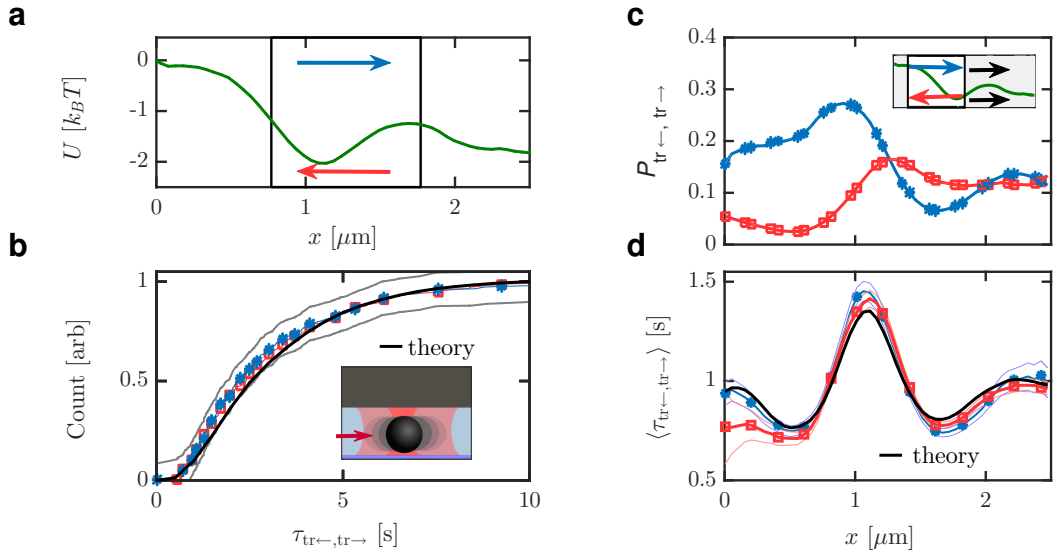

Fig. 3.2 Transition-path-time symmetry of particles in force landscapes. **a** Free energy $U(x)$ calculated from inferred forces (see Sec. 3.1.6). The standard error of the mean (not shown) is smaller than the line thickness. **b** Cumulative probability of transition-path times across the region defined by the black box in **a** (compare with Fig. 1.2**b**). Kolmogorov-Smirnov test boundaries are plotted in light grey. **c** Transition probabilities across a 1 µm region, which is continuously moved along x. The standard error of the mean (not plotted) is on the order of the marker size. **d** Transition-path times over the same spatial interval. Error envelopes show the standard error of the mean. While the lines in panels **b** - **d** show the full data set, the markers only highlight selected data points.

(see text leading up to Eq. (1.24) for derivation)

$$\partial_t \rho(x,t|x_0) = -\partial_x j(x,t|x_0) \tag{3.1}$$

with $j(x,t|x_0) = f\rho(x,t|x_0) - D(x)\partial_x\rho(x,t|x_0)$ denoting the current of probability. The initial density of colloid positions $\rho(x,t_0=0) = \rho_0(x)$ was modelled as a sharp peak at the channel centre, $x_0 = 0$. Once $j(x,t|x_0)$ is obtained, the exit-path time distribution is given by the one-dimensional version of Eq. (1.41), i.e. $\rho_{\tau_\leftarrow}(t) = j(x_\leftarrow,t|x_0)/P_\leftarrow(x_0)$ with $x_\leftarrow$ denoting the x-position at the left boundary. Both boundaries were treated as absorbing, i.e. $\rho(x_\leftarrow,t) = \rho(x_\rightarrow,t) = 0$. The two exit probabilities $P_\leftarrow(f,x_0=0)$ and $P_\rightarrow(f,x_0=0)$ are given by Eqs. (1.44) and (1.45). $\rho_{\tau_\rightarrow}(t)$ can be obtained by exchanging $x_\leftarrow$ for $x_\rightarrow$ and $P_\leftarrow$ for $P_\rightarrow$.

3.1.2 Uphill and downhill transition-path-time symmetry

The robust symmetry observed experimentally in exit-path times is also found in direct transitions between any two points x_{L}, x_{R} in a quasi-one-dimensional microchannel that is filled with an optical landscape. I deliberately choose left (L) and right (R) subscripts here to contrast transition-path times from exit-path times (see also Fig. 1.2**a**,**b**). The landscape considered here consists of a mixture of a point trap and a line trap with a positive phase-gradient force created by the HOT (see inset in panel 3.2**b**). I employ the same HOT automation routine as before to observe around 500 uninterrupted colloid trajectories. The energy potential inferred from this ensemble of particle trajectories is plotted in panel 3.2**a**. The transition-path times $\tau_{\mathrm{tr}_{\leftarrow}}$, $\tau_{\mathrm{tr}_{\rightarrow}}$ across the interval shown as a black box in 3.2**a**, are identically distributed as can be seen in 3.2**b**. Based on a spline interpolation of the inferred potential $U(x)$ and a spatially dependent diffusion coefficient $D(x)$, I calculated the theoretical distribution of transition-path times $\rho_{\tau_{\mathrm{tr}}}(t)$. Again, I treat both boundaries x_{L} and x_{R} as absorbing. I compute $\rho_{\tau_{\mathrm{tr}}}(t)$ for an initial density $\rho_0(x)$ which is sharply peaked close to the initial exit, following Sec. 1.2.2. For this purpose, I use a one-dimensional Gaussian $\mathcal{N}(x_0,\sigma^2)$ with $x_0 = 4\sigma$ measured from the initial exit with a $\sigma \sim 10^{-7}\mu$m only just large enough to avoid underflow in Mathematica.

For the sake of this example, I choose the direction left to right and thus set $x_0(\epsilon) = x_{\mathrm{L}} + \epsilon$. The current density reads $j(x_{\mathrm{R}},t|x_0) = \partial_x(U(x)\rho(x,t|x_0))|_{x=x_{\mathrm{R}}} - D(x_{\mathrm{R}})\partial_x\rho(x,t|x_0)|_{x=x_{\mathrm{R}}}$ at the right boundary x_{R}. I normalise the distribution by the overall probability to exit through x_{R}, $P_{\rightarrow}(x_0)$ (see Eq. (1.45)), assuming x_0 as the initial position. Finally, I obtain the distribution of transition times $\rho_{\tau_{\mathrm{tr}}\rightarrow}(t)$ from Eq. (1.53). A plot of $\rho_{\tau_{\mathrm{tr}}\rightarrow}(t)$ is shown in Fig. 3.2**b** (black).

In panel 3.2**c** I plot the probability of direct transition across the same interval length, when this interval is continuously moved along the channel. For each position of this interval, the transition probabilities and times are recorded. As can be seen in panel 3.2**d**, the mean transition-path times calculated in this way in both directions are sensitive to the local force, especially when the transition interval touches the optical point trap. Despite this sensitivity, transition times in both directions are in excellent agreement. The theoretical prediction for the mean transition-path time $\langle \tau_{\mathrm{tr}_{\rightarrow}} \rangle$ plotted in black in Fig. 3.2**d** was calculated using Eq. (1.52).

3.1.3 Breakdown of path-time symmetry under external coloured noise

The question that arises is whether and how this symmetry can be broken. In Sec. 1.1.1, I hint at "thermodynamical consequences" of enlarging phase-space from one to two di-

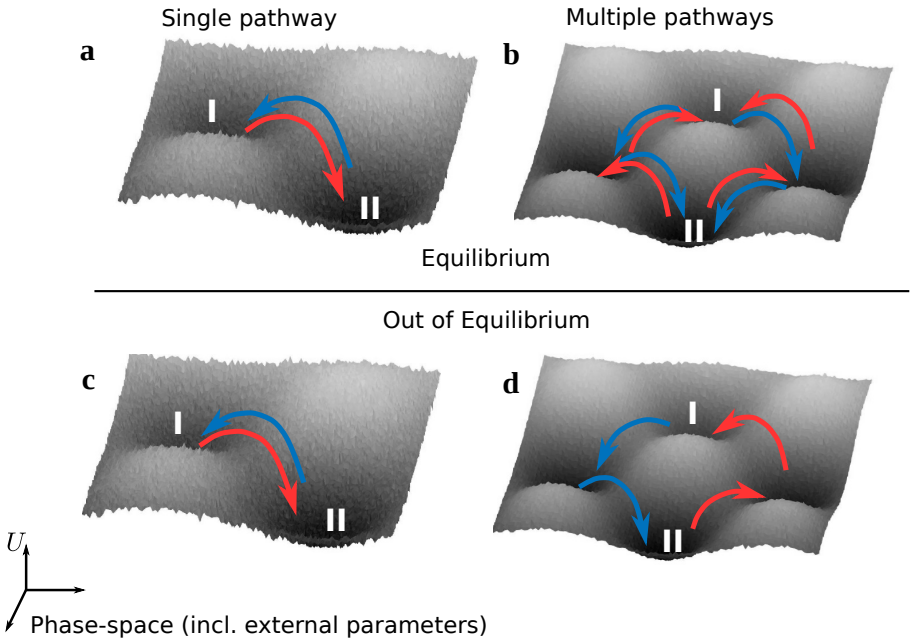

Fig. 3.3 Sketch of the topological perspective on equilibrium and out-of-equilibrium state transitions. **a** Single pathway in equilibrium. **b** Multiple pathways in equilibrium. Detailed balance fulfilled everywhere. Transition path-time statistics must necessarily be symmetric, but potentially multimodal transition-path time distribution. **c** Non-equilibrium transitions confined to a single-pathway. This scenario arises, for instance, when an external white-noise is applied to a confined system. **d** Breakdown of detailed balance in multi-pathway systems. In steady-state state this system will sustain circulatory and self-avoiding probability fluxes. Transition-path times can be asymmetric.

mensions. Indeed, detailed balance cannot be broken in steady-state in one-dimensional systems with continuous variables (or discrete systems with only two states). It can break down, however, as soon as circular, self-avoiding fluxes become possible from a topological point-of-view [61, 56]. I sketch and summarise the thinking behind the topological perspective on state transitions in Fig. 3.3.

In the following section, I describe the effect of external forces $f_{\text{ext}}(t)$, that stochastically switch between two levels $+f_0$ and $-f_0$ with exponentially distributed switching times. Such two-state switching processes are generally referred to as telegraph noise. The time between two switches is, as mentioned, exponentially distributed with a decorrelation rate α, such that $\langle f_{\text{ext}}(t+\Delta t) f_{\text{ext}}(t)\rangle \propto e^{-\alpha \Delta t}$ for $\Delta t > 0$. As can be seen in this example, non-Markovian variables have internal states, which are hidden from the system. In this case, the topological effect is the temporal persistence of the applied force value, which enables state transitions across, effectively two different potential barriers with force-fields $f_1(x) = -\partial_x U(x) + f_0$ and $f_2(x) = -\partial_x U(x) - f_0$.

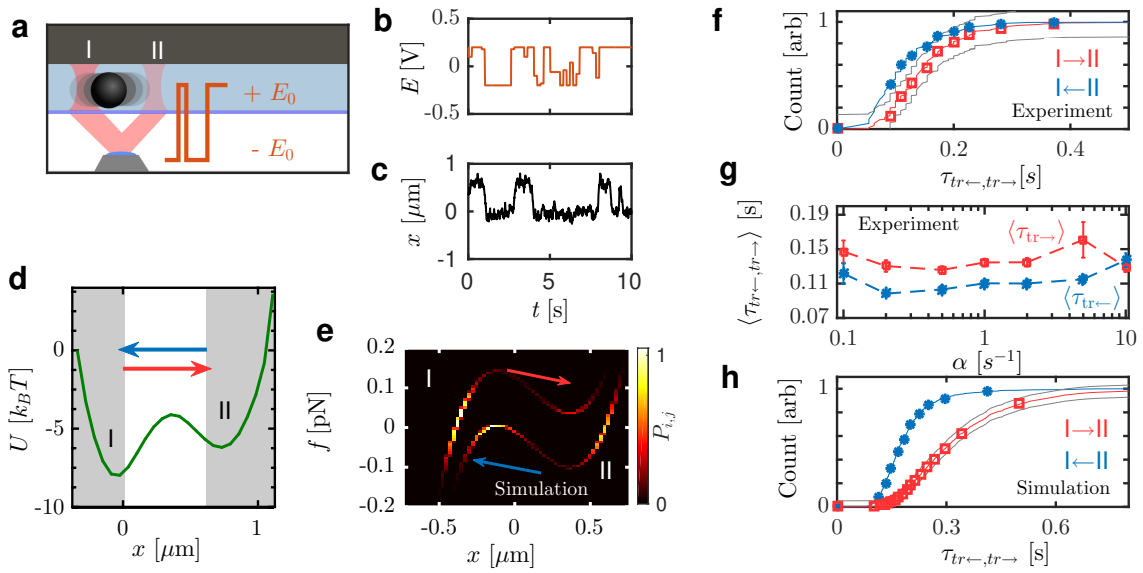

Fig. 3.4 Transition-path times of particles between two potential minima in microfluidic experiments and simulations. **a** Sketch of the experimental situation: a bistable potential is created by a strong and weak optical trap under stochastically switching voltage $\pm E_0$. **b** Applied voltage $E(t)$ over the first ten seconds. Brief jumps in the telegraph process cannot be resolved and appear as averages between the two voltage levels. **c** Trace of the first ten seconds of the position variable $x(t)$. **d** Mean potential U inferred from the trajectories. Error is on the order of the marker size (not shown). **f** Symmetry breaking in the cumulative distribution of transition-path times I $\rightarrow$ II (red) and I $\leftarrow$ II (blue). **e** Colour map representing the relative occupation of states in the force$\times$position plane under the influence of an external, coloured force. f here denotes the total force $f = f_{\text{ext}} - \mathrm{d}U/\mathrm{d}x$. **g** Mean transition-path times $\langle\tau_{\leftarrow,\rightarrow}\rangle$ over the decorrelation rate α of the electrical field process. Error bars show the standard error of the mean. **h** Cumulative distribution of transition-path times obtained in Brownian dynamics simulations.

The application of external telegraph forces drives the system into a non-equilibrium steady state (NESS). I create a NESS on the mesoscale by combining a bistable optical potential consisting of two point traps with different trap strenghts created with our HOT and randomly sign-switching electrical fields (see panel 3.4**a**). A trace of the position variable $x(t)$ and the electrical field $E(t)$ are shown in panels **b** and **c**. I set the traps apart by 0.7 μm and direct 50% more light to the left trap than to the trap on the right, while operating at an overall laser power of $\sim$ 50 mW to avoid heat-induced convection. The minima of the two traps correspond to the two states I and II the transition time is measured over the white area in panel 3.4**d**. Different trap strengths result in a difference in curvature between the two traps; the transition barrier loses its symmetry with respect to the centreline separating the two states.

Indeed, as shown in panel 3.4**f**, I observe a statistically significant difference in the distribution of transition-path times $\tau_{\mathrm{I}\to\mathrm{II}}$ and $\tau_{\mathrm{I}\leftarrow\mathrm{II}}$ for $\alpha = 0.5\,\mathrm{s}^{-1}$. Interestingly, the difference in mean transition times stays fairly constant over a range of decorrelation rates α, as shown in panel 3.4**g**. Towards higher noise decorrelation rates α (note the logarithmic scale), the symmetry is restored. For high frequencies of sign switches, the telegraph force approaches a white noise process and the system should approach the scenario sketched in Fig. 3.3**c**.

To study further the underlying mechanism that led to the breaking of this symmetry, I recreate my experiments in one-dimensional Brownian dynamics simulations (details are discussed in Sec. 3.1.8). The motion of a colloidal particle under the influence of external telegraph forces is well described by the following equation

$$\gamma\dot{x}(t) = f_{\mathrm{ext}}(t) - \frac{\partial U}{\partial x}(x) + \sqrt{2k_BT\gamma}\xi(t), \tag{3.2}$$

where γ denotes the friction coefficient of a sphere and $f_{\mathrm{ext}}(t) = f_0\mathscr{T}(t)$ denotes the force exerted by the electrical field. $\mathscr{T}(t)$ represents a random telegraph process that switches between 1 and -1. $U(x)$ corresponds to the free energy and $\xi(t)$ is a Gaussian white-noise process with zero mean and unit variance $\langle\xi(t)\xi(t')\rangle = \delta(t-t')$. I did not attempt to model every parameter of the experiment quantitatively, but rather test the generality of the observed split of transition pathways. The distribution of the system state in the NESS is shown in panel 3.4**e**. The transition pathways indeed split up. The distribution $P_{i,j}$ of the system state along these pathways in this force$\times$position plane turn out to be visibly different on each leg. Colloids transitioning into one direction will therefore likely experience different force distributions along their pathway than colloids transitioning into the opposite direction. The red and blue arrows in the figure indicate the preferred

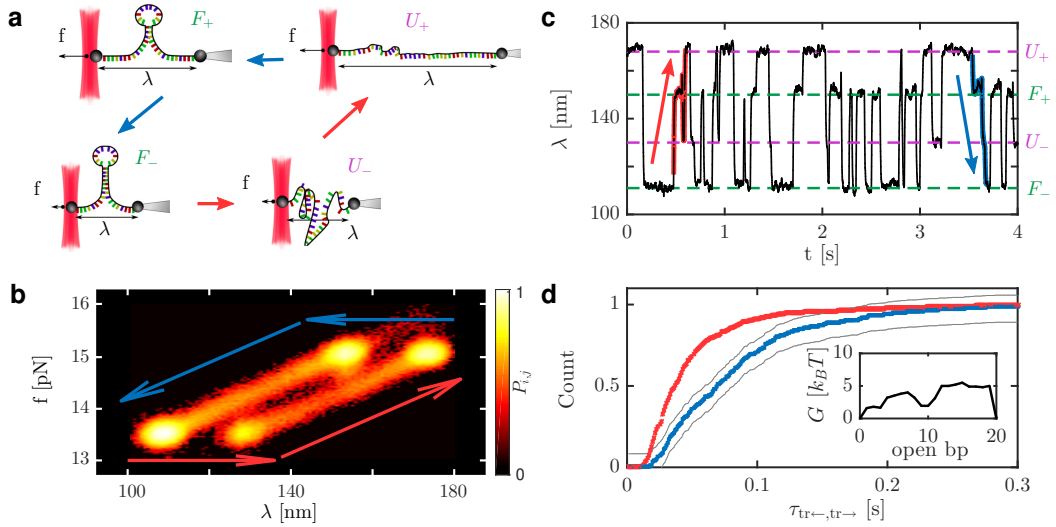

Fig. 3.5 Folding and unfolding transitions of DNA hairpins under fluctuating external forces. **a** Sketch of the experiment: Each end of the DNA-hairpin strand is grafted onto a colloid. While one colloid is firmly attached to a pipette, the other is held in force-measuring optical tweezers. The position of the different states F_+, F_-, U_+, U_- are indicated in the $f \times \lambda$ plane (force×extension plane). **b** Relative occupation of states in the $\lambda \times f$ plane. **c** Excerpt of the dynamics of the hairpin. **d** Cumulative histogram of transition times from $U_+ \rightarrow F_-$ (blue) and vice-versa (red). The inset shows the free energy G of the barrier in equilibrium. The experiments were carried out by our collaborators Dr M Ribezzi-Crivellari and Prof F Ritort. The figure was created jointly by Dr M Ribezzi-Crivellari and me.

sense of transition direction along the two pathways; detailed balance is indeed broken in this two-dimensional space [61] and transition-path times differ as a consequence (see Fig. 3.4**f**).

3.1.4 Folding-time symmetry and its breakdown in DNA-hairpins

The following section draws on data provided by my collaborator Dr Marco Ribezzi-Crivellari, who measured folding dynamics of DNA hairpins in Prof Felix Ritort's lab at the University of Barcelona. The data provides a molecular analogue of my microfluidic experiments.

Having established broken symmetry on the mesoscale I now show the generality of the effect with an experimental realisation on the molecular scale. My collaborators measured folding and unfolding transition times of short 20-bp DNA-hairpins under the influence of telegraph forces using an optical tweezers-based force spectroscopy

setup [152] (see Fig. 3.5**a**). In these experiments, the hairpin is grafted onto a colloid on each end. While one colloid is firmly attached to a pipette, the other colloid is held in force-measuring optical tweezers and subject to a feedback-controlled force. If the force is kept constant in time, the hairpin thermally transitions between two main states folded (F) and unfolded (U), which differ in molecular extension and thus in trap position λ. I here describe experiments performed using a non-equilibrium protocol. The system is subject to a telegraph force and each of these two (U,F) states splits into a doublet: high (F_+, U_+) and low force (F_-, U_-). This is shown in panel 3.5**b** using the density $P_{i,j}$ of states in a coarse-grained space spanned by the force measured by the optical tweezers f and the trap position λ. Importantly, as the arrows indicate, the telegraph force not only leads to a splitting of states, but also causes transition pathways to diversify. The system is more likely to unfold $F \rightarrow U$ during extended periods of high force ($+$), than during periods of low force ($-$). As a consequence, transitions from state F_- to U_+ through U_- (red arrows in Fig. 3.5**b**,**c**) are more likely than transitions from state U_+ to F_- through F_+ (blue arrows in Fig. 3.5**b**,**c**).

A typical trajectory of the system is shown in panel 3.5**c**, highlighting transitions from U_+ to F_- (blue) and vice versa (red). This split in pathways results in the visible difference of cumulative distributions of folding (red) and unfolding (blue) transition times in panel 3.5**d**. However, this difference is not necessarily always present: The difference between back and forth transition-path times can become arbitrarily small under certain conditions that we describe in further detail in our paper [94]. We conclude, that an observed transition-path-time asymmetry in an overdamped system indicates that the system under study breaks detailed balance and is hence out of equilibrium. Importantly, the reverse conclusion does not hold: a system which satisfies transition-path time symmetries is not necessarily in equilibrium.

Overall, the transition-path times of our DNA-hairpin are significantly longer than previously reported values [72], because our system transitions via intermediate states (U_- or F_+). The time spent in corresponding minima affects the overall transition-path time in a path-dependent way and thus amplifies the asymmetry. By contrast, on the mesoscale, transitions are slow enough such that we could resolve the asymmetries shown in Fig. 3.4**f**, which directly originate from asymmetries in the barrier shape.

In conclusion, I find that the overarching topological picture indeed applies to the molecular scale; the dimensionality of the space of folding is effectively increased by at least one due to the external coloured noise. In this increased phase-space, a breakdown of detailed balance results in a diversification of transition pathways, which causes a transition-path-time asymmetry. Importantly, all participating degrees of freedom, in-

cluding internal variables of external forces (see the example in Sec. 1.1.3), have to be considered in the analysis.

3.1.5 Conclusion

In this section, I present experimental evidence of a fundamental transition-path-time symmetry in Brownian transitions and its breakdown on the meso- and molecular scale under the influence of stochastic external forces. In accord with intuition, my colleagues and I find that uphill transitions become less likely, as the potential gradient between the initial and end state becomes steeper. Uphill and downhill transition-path times, however, are identically distributed under steady-state conditions. Conceptually, I show that transition-path times connecting any two points in the space of the system are thermodynamically well-defined quantities. Indeed, I find that in a time-constant force landscape, measured transition-path times agree with theoretical predictions that assume absorbing boundary conditions at both ends of the transition interval. It is important to note that boundaries can be located anywhere in the potential landscape and do not need to coincide with minima of the potential.

In contrast to transitions driven by thermal forces, I find that the transition-path-time symmetry can break down under the influence of coloured noise. The additional timescale of external telegraph forces in the microfluidic systems changes the topology of transition dynamics. I therefore uncover a diversification of transition pathways in the extended phase-space, which includes the external force. Back and forth reactions follow, on average, different paths, breaking detailed balance and the transition-path-time symmetry. Specifically, I show that transition-path times of a colloid in an asymmetric double well potential become measurably asymmetric, when perturbed by randomly switching electrical fields. The asymmetry is sensitive to the frequency of field reversals and disappears for frequencies that are much higher than the barrier crossing time. Similarly, a DNA-hairpin that is driven out of equilibrium by a force that switches randomly between two levels, exhibits asymmetric folding-/unfolding-path times. The observed asymmetry in transition-path times, however, is a result of an implicit projection of the system state onto a one-dimensional reaction coordinate. A breakdown of transition-path-time symmetry does therefore not imply a breakdown of microscopic reversibility. I note that all systems studied here are overdamped and effects related to inertia can be neglected.

My results have direct implications for the study of transitions in membrane channels or nanopores. Translocation times of solutes, like antibiotics, through membrane channels, are of interest in electrophysiological measurements [148, 149]. Due to a lack of any direct optical access in these experiments, the shape of the current signal during

translocation is the only source of information about the channel-solute interaction. The results show furthermore that it should be possible to infer the direction of travel solely from first-passage times. A reversal of the electrical potential in such an experiment should result in a distinct translocation time distribution if the solute-channel interaction landscape is asymmetric. The combination of a (sign-flipped) field and interaction potential can be interpreted as the limit of infinite switching times in Fig. 3.4, which amounts to the simple case of back and forth translocating solutes experiencing different time-constant force landscapes. Furthermore, in studies of molecular motors, transition-path time measurements could enable one to discriminate between power stroke and ratchet mechanisms, beyond thermodynamic considerations [24]. Arguments based on first-passage-time symmetries have already been used to question the thermodynamic consistency of interpretations of Kinesin motility experiments [153, 93].

Moreover, in systems driven by ratcheting [23, 154, 25, 145], unbiased coloured noise rectifies Brownian dynamics around points of asymmetry of the energy landscape. The asymmetry of transition-path times demonstrated in this study could be used as experimental evidence of this effect. Transition-path-time asymmetries could therefore be helpful in identifying and quantifying non-equilibrium dynamics in biological and molecular systems and complement recently discussed techniques such as broken detailed balance in active matter [155] and filament fluctuations [61]. I note that a breakdown of the transition-path-time symmetry can be diagnosed by tracking only one degree-of-freedom, whereas diagnosing a breakdown of detailed balance requires a minimal dimension of two in continuous coordinates [155]. This might be particularly helpful in FRET experiments, where usually only a single degree-of-freedom, the FRET efficiency, is accessible.

The path-time symmetries I explore come in two flavours: an exit-path-time symmetry (see Fig.3.1) and a transition-path-time symmetry (see Fig.3.2). Recent theoretical advances [85, 83] point to a connection between the two flavours, which lies in a symmetry of the first-passage time of the entropy produced during the transition. Since a breakdown of transition-path-time symmetry is a sufficient, but not necessary condition for non-equilibrium dynamics, it is in general not possible to deduce entropy production from observed transition-path times as my coauthors and I discuss in Supplementary Note 4 and Supplementary Figure 4 of our publication [94].

3.1.6 Experimental methods

The colloidal particles used in the experiments described above, consist of polycarbonate with a COOH-functionalized surface, with a diameter of 0.5 μm, and were purchased from

Polysciences Inc. The particles were suspended in 0.5× TRIS-EDTA buffer at pH 8 and additionally 3 mM KCl to screen potential charges on the walls of the channel. Prior to each experiment, I ascertained that no detectable hydrodynamic flow was present by comparing left and right exit probabilities of colloidal particles initialised in the centre of the channel.

The geometry and fabrication of the $5 \times 1 \times 1$ μm^3-microfluidic-channel mask used here is explained in [156]. The frame rate was set to 80 fps in all experiments. Details of the experimental protocols and setup used by my coauthors in this section can be found in our publication [94].

Colloids as hard-sphere Brownian walkers

The Debye length which characterizes the length scale of eletrostatic interactions in aquaeous solutions is here on the order of $\lambda_D \approx 4$ nm given the KCl and TRIS-EDTA concentrations described above (see also Chapter 5). The channel geometry, however, is on the order of micrometers and electrostatic effects should therefore be negligible for my purposes. From an electrostatics perspective, isolated colloids here should behave as hard spheres which do not feel wall repulsion unless they are in immediate contact with the channel walls.

In the presence of time-varying electrical fields, electrodynamic effects have to be considered. In my experiments discussed in Fig. 3.4, I use switching electrical fields to drive colloidal particles trapped in optical potentials out of equilibrium. In such circumstances, the electrophoretic mobility μ_{el} of the colloids determines the strength of the electrophoretic force the fields exert on them. In bulk, COOH-labelled colloids with a nominal diameter of 380 nm dispersed in a solution with a salt concentration of $c_{\text{KCl}} = 10$ mM have been measured to have an electrophoretic mobility of $\mu_{\text{el}} \approx 10^{-12}\,\text{m}^2\text{V}^{-1}\text{s}^{-1}$ at pH 7 [157]. This value is not expected to be significantly different at pH 8, as measurements indicate [157]. However, as I explain in more detail in Chapter 5, electrophoresis in channels typically occurs in conjunction with counteracting electroosmotic flows.

Detailed measurements of colloidal dynamics conducted in and nearby channels with similar depths and length scales show that the interplay of these two effects can cause complex colloidal dynamics at channel entrances [157]. It is for this reason that I chose to conduct my experiments only within the interior of channels so that entrance-geometry-related effects, such as those described in the paper, do not apply. For the purposes of this study it is furthermore irrelevant if the colloids are driven by electrophoretically- or electroosmotically-induced forces (or a mixture thereof). It is sufficient that the net-exerted force swiftly switches direction with the sign of the applied voltage.

Inference of Brownian motion parameters

A statistical tool which is of particular importance to my thesis are force and diffusion-coefficient estimators. Ideally, these estimators do not require additional measurements, but can simply be applied to the trajectories measured during the experiment. There are several published methods that can be used to infer forces from trajectories, including the use of splitting probabilities [158], and Bayesian methods [159]. The method, I adopted in all of my experiments is simpler: In order to infer steady-state forces and diffusion coefficients in a one-dimensional interval $[a,b]$, the interval is first subdivided into N bins with length $\delta = (b-a)/N$. Then, the displacement statistics $\{\Delta x_t\}$ are calculated, i.e. $\Delta x_t = x_{t+1} - x_t$. The displacements are assigned to a bin k where the step originates from, that is, $k = \mathrm{floor}((x_t - a)/\delta)$ for $k \in \{0,\ldots,N-1\}$.

Next, a Gaussian is fitted to the displacement statistics in each bin, $\rho_k(\Delta x) \sim \mathcal{N}\left(\mu_k, \sigma_k^2\right)$. The diffusion coefficient is then obtained from the variance

$$\hat{D}_k = \sigma_k^2/(2\Delta t) \tag{3.3}$$

where Δt is the time step between two consecutive camera frames. For the force, the friction coefficient γ is required. I here make the assumption that the Einstein-Smoluchowski relation holds and write $\hat{\gamma}_k = k_B T/\hat{D}_k$. The force then follows using

$$\hat{f}_k = \hat{\gamma}_k \mu_k/\Delta t. \tag{3.4}$$

Confidence bounds for $\hat{f}$ and $\hat{D}$ can be obtained by propagating the bounds of the Gaussian fit through Eqs. (3.3) and (3.4). Potentials $U(x)$ are obtained from $\hat{U}(x) = -\int_a^x \mathrm{d}x' f(x') \approx -\sum_{l=0}^{k}(\hat{f}_l + \hat{f}_{l+1})/2$.

Phase-gradient calibration

I give a short description of phase-gradients in Sec. 2.2.1. In my explanation, I mention the relationship between tweezing-beam-angle and position of the hologram-envelope (μ_x, μ_y). This relationship can be observed experimentally as shown in Fig. 3.6. Instead of going through a single 2f-system, the beam in my setup crosses three lenses (the last one being the objective).

In Fig. 3.6 I show a scan of the back reflection for different axial positions of the cover slip using the Piezo-stage. The angle of the beam can be measured from these pictures by observing the lateral movement of the centre of intensity, $\vec{c}(z) = (x(z), y(z))$. The angle then follows from $\alpha(p) = \langle \arctan(\Delta z/\Delta x) \rangle$ where Δz is the axial-step of the piezo, Δx is

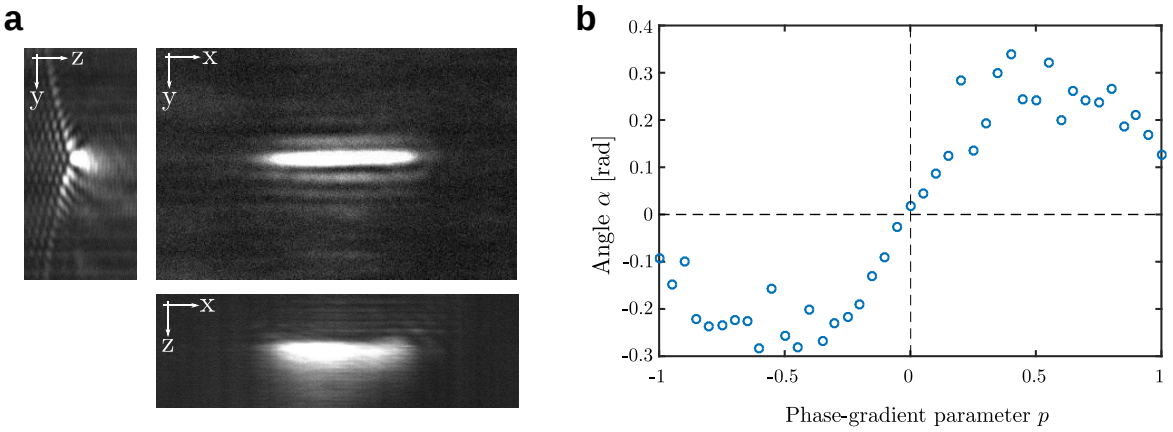

Fig. 3.6 Visualisation of phase-gradients. **a** Back-scattered image of a line-trap taken at different axial positions. $X-Y-Z$-cuts of the three-dimensional intensity profile along the respective central plane. **b** By measuring the movement of the centre-of-mass over z, one can infer the angle of the beam.

the inferred lateral movement of the centre of intensity, and the average runs over the axial image stack. Whenever I need to observe the back reflection from the microscope cover slip, I use CMOS 2 in Fig. 2.6. The shape of the created intensity pattern is otherwise better observed on CMOS 3 (see Fig. 3.7**d**).

Calibration means inferring the relationship $f(p)$ of the forces that are created in a line-trap with phase-gradient parameter p (see Fig. 3.7**a**, **b**). This can be done with the experimental data itself; there is no need for additional experiments. By solving Eqs. (1.44) and (1.45) for f, I can compute the force f from the the exit probabilities P_R, P_L. Since the measured exit-times in Fig. 1.4 are in excellent agreement with theory based on these force-estimates, the theoretical framework is shown to work in a self-consistent way.

3.1.7 The Kolmogorov-Smirnov test

A visual inspection of similarity between two cumulative distributions can be deceiving if no quantitative scale of the expected deviation is established. For instance, what level of similarity is sufficient in the case of Fig. 3.1, where two cumulative samples of exit-path times are presented, in order to conclude that the samples have indeed been drawn from the same underlying distribution? This question can be answered using methods from hypothesis testing, specifically the two-sample Kolmogorov-Smirnov test (KS test) [160]. In the following, I give a brief practical outline of how the KS test is applied to two empirical distributions of path times. The test relies on a substantial amount of theory, which is explained in the literature [160]. The null-hypothesis that is tested here at a chosen level

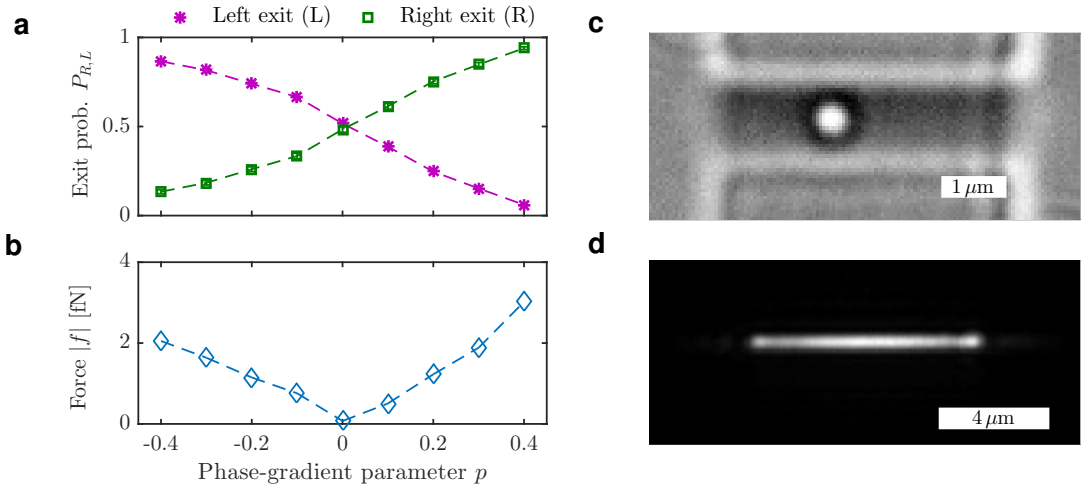

Fig. 3.7 Calibration procedure for line traps. **a** Measured conditional exit probabilities P_L, P_R. **b** Self-consistent force estimate inferred from panel **a** using Eqs. (1.45) and (1.44). **c** Microscope picture of the channel and particle used for the experiment (CMOS 1). **d** Mid-plane picture of the line trap (CMOS 3).

of significance $\alpha = 0.5$ is that both empirical distributions are in fact samples from the same underlying distribution. I chose the value of α as the highest level of significance rounded to one digit in which the KS test would still hold for all presented experiments.

The two-sample KS test is based on the maximum distance

$$D_{n,m} = \max_\tau |P_{1,n}(\tau) - P_{2,m}(\tau)| \tag{3.5}$$

between two cumulative empirical distribution functions $P_{1,n}, P_{2,m}$ with sample sizes n and m respectively. I implemented Eq. (3.5) by linearly interpolating between the discrete elements of $P_{1,n}$ and $P_{2,m}$. I note that in my case, the two distributions can have substantially different sample sizes n, m due to the difference in uphill and downhill transition probabilities.

Once the distance $D_{n,m}$ is computed, the next step in the KS test is to compare it to a table of critical values [160] in order to determine if the null-hypothesis is rejected at significance α. In my case, for a significance level of $\alpha = 0.5$, the distance must fulfil

$$D_{n,m} < 1.22\sqrt{\frac{n+m}{nm}}. \tag{3.6}$$

The grey lines in Figs. 3.1**f**, **g**, 3.2**b**, 3.4**f**, **h**, and 3.5**d** represent the maximum deviation of either the red from the blue curve or vice versa that would still satisfy the KS test at the

chosen level of significance. Importantly, the result of the test does not depend on which deviation (blue-to-red or vice versa) is chosen.

3.1.8 Brownian dynamics simulations

The Brownian dynamics simulation mentioned in Sec. 3.1.3 is designed to qualitatively model the doublewell experiment. I use a bistable, asymmetric potential of the form $U(x) = a/4x^4 + b/2x^2 + cx$, where c controls the asymmetry around $x = 0$. The coefficients are set to $a = 64\Delta U_1/L^4$, $b = -aL^2/4$, and $c = 2\Delta U_2/L$, with $\Delta U_1 = 5\,k_B T$, $\Delta U_2 = 2\,k_B T$ and $L = 1$ µm. I set the diffusion constant to $D = 0.15\,\mu\text{m}^2/\text{s}$, which is close to the value a 500 nm colloid has in a microchannel. The friction coefficient is obtained, again, using the Einstein-Stokes relation $\gamma = k_B T/D$. In simulations corresponding to the results shown in Fig. 3.4**e**, the decorrelation time of the telegraph force is set to $2s$, while the magnitude of the force change is set to $f_0 = \pm 82$ fN.

The physical parameters such as D, L and the potential parameters a, b, c are chosen such that the simulations model an overdamped colloid of size 500 nm on roughly on realistic length, energy, and time scales. The most relevant theoretical simplification here consists in the assumption that no other forces besides the gradient of the potential $U(x)$ and the telegraph force exist. Furthermore, I assume a linear dependence of the force exerted by the fluctuating voltage on its magnitude. The temporal process used to model for the force here is in fact the process used to generate voltage traces, not the force traces, in the experiment. This means, that I assume that any change in voltage immediately translates into a proportional change in force.

3.2 Optimising Kramers rates by potential shaping

In the previous section, I explore in detail exit- and transition-time-path symmetries in and out-of-thermal equilibrium. By contrast, in this section, I consider Kramers escape times in steady-state potentials and show that the escape from a potential minimum is considerably faster for certain barrier shapes, even when compared to unhindered diffusion. At the time of the final submission of this thesis, the results presented here have just been published [161].

As I discuss in Sec. 1.2.3, H.A. Kramers derived a comprehensive theory for the Arrhenius-like scaling of thermally activated transition rates, $k \propto e^{-\Delta U/(k_B T)}$, and introduced a framework to describe such transitions in an energy landscape $U(x)$. Perhaps surprisingly, influences of barrier shapes on transition rates and conditions of optimality thereof have hitherto not received due attention.

My coauthor, Dr. Marie Chupeau in Prof. Emmanuel Trizac's group in Paris, used a variational approach to optimise static barrier profiles and calculate the corresponding mean time of escape. Interestingly, she finds that the maximum achievable relative escape rate is infinite. The barrier optimisation therefore has to be regularised, for instance, by placing an upper bound on the barrier height or curvature. Furthermore, my colleagues show that the rate-boost applies over a range of friction values, extending from the overdamped into the inertial regime (see Sec. 1.1.2). Importantly, the need for regularisation disappears in the underdamped regime.

Guided by this theoretical prediction, I demonstrate in this section experimentally that higher, optimised barriers can paradoxically result in increased escape rates, in contrast to intuition based on Kramers law in Eq. (1.59).

Specifically, I experimentally demonstrate a doubling of escape rates compared to unhindered Brownian motion, which proves that the required barrier profiles can indeed be realised. My results indicate that certain fine-tuned free-energy landscapes of higher amplitude increase reaction rates. In the context of protein folding, a carefully rate-optimised free-energy landscape may thus well exhibit a larger number of intermediate states in spite of additional necessary escapes [162]. In the following, I describe the necessary elements of the theoretical work and my approach to experimental realisation.

As I explain in Sec. 1.2.3, the reciprocal of the escape rate corresponds to the time of first-passage to leave the initial state in Kramers problem. A lower bound for the achievable first-passage time, e.g. of the reaction coordinate of a folding molecule, therefore corresponds to a speed-limit of the ensemble reaction rate [163].

To my knowledge, the idea that the introduction of barriers might speed-up escapes rather than slow them down goes back to Palyulin and Metzler [164, 165].

In our publication, my coauthors and I investigate this idea systematically from a theoretical and experimental perspective. As Palyulin and Metzler show, the speed-up effect already appears for relatively simple barrier shapes [164], such as the triangular-shape shown in Fig. 3.8**a**. Without a reference time, the term "speed-up" is, of course, meaningless. The optimised escape time must be compared to a different escape time. In our case, this reference is given by the free diffusion time $\tau_0(x) = x^2/(2D)$ (set $U(x) = 0$ in Eq. (1.57)). Crucially, the barriers we consider end on the same energy where they begin, such that the process of crossing the barrier does not alter the overall energy balance.

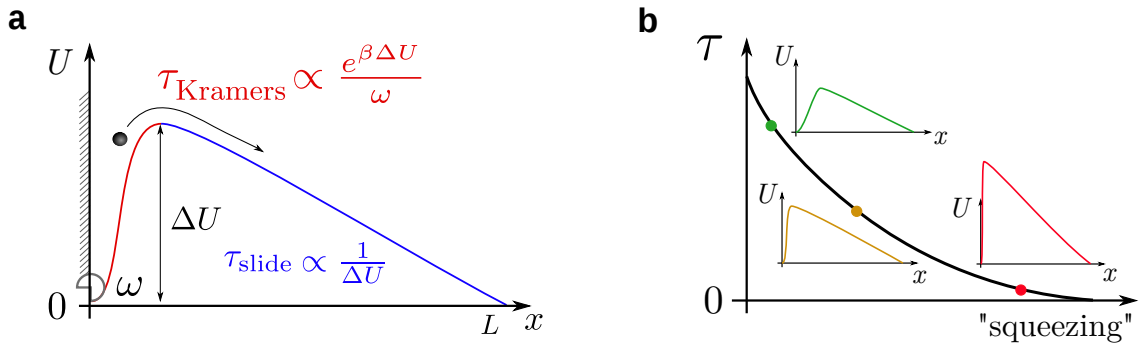

Fig. 3.8 Escape rate-boost effect in a simple potential. **a** Triangular potential defined by the barrier height ΔU and curvature ω of the intial well. β here denotes the inverse temperature, i.e. $\beta = 1/(k_B T)$. Provided ΔU is large enough, the exit-path time to a distance $x = L$ is approximately given by $\tau \simeq \tau_{\text{Kramers}} + \tau_{\text{slide}}$, indicated in red and blue respectively. **b** Squeezing of the initial well (increasing barrier height and initial curvature) results in an arbitrarily short exit-path time. The optimisation problem is thus ill-posed and requires regularisation, at least in the overdamped regime of Brownian motion. The figure was jointly created by Dr. M Chupeau and me.

3.2.1 Variational calculus applied to Kramers problem

The idea of our publication is to take Palyulin and Metzler's analysis one step further and systematically consider the escape time $\tau[U(x)]$ in Eq. (1.57) as a functional of the potential shape $U(x)$.

But before I discuss this variational approach, I want to give some intuition of the speed-up effect.

At $t = 0$, the particle is initialised in the potential well at $x = 0$ (red) in Fig. 3.8. Similar to my discussion in Sec. 1.2.3, the particle is bounded by a reflecting barrier to the left at $x = 0$ and an absorbing boundary to the right at $x = L$. The mean exit-path time from the narrow well is then given by Kramers result $\tau_{\text{Kramers}} \propto \omega^{-1} e^{\frac{\Delta U}{k_B T}}$. Once the particle leaves

the well, it drifts towards the absorbing boundary at $x = L$ within an average slide time $\tau_{\text{slide}} \propto 1/\Delta U$ as follows from gradient descent.

In the high-barrier-limit $\Delta U \gg k_B T$, the overall mean exit-path time τ is approximately given by the sum of τ_{Kramers} and τ_{slide} [164]. This time can therefore be shortened by a simultaneous increase in curvature of the initial well and height of the barrier, see Fig. 3.8**b**. Crucially, there is no lower bound (other than 0) for the exit-path time: further "*squeezing*" will further decrease τ. As a consequence, a sufficiently high and steep barrier yields a mean exit-path time, which is shorter than the corresponding free diffusion time.

In order to simplify the notation, my coauthors rescaled all relevant quantities, i.e. the mean exit-path time, the potential, and the position: $\widetilde{\tau} = D\tau/L^2$, $\widetilde{U} = U/(k_B T)$, $\widetilde{x} = x/L$. The Smoluchowski-Einstein relation $D = k_B T/\gamma$ here relates the temperature T driving the Brownian process to the friction coefficient γ. Tildes are used to denote dimensionless variables, but I drop them hereafter. Furthermore, I also drop the bra-kets around τ, i.e. $\langle\tau\rangle \equiv \tau$ to further simplify the notation.

The mean exit-path time (Eq. 1.56) then reads

$$\tau\left[U(x)\right] = \int_0^1 \mathrm{d}x\, e^{-U(x)} \int_x^1 \mathrm{d}y\, e^{U(y)}. \tag{3.7}$$

Using this functional, I can now ask which potential profile $U(x)$ that satisfies $U(0) = U(x_L) = 0$ leads to a minimal mean escape time τ.

I here follow a slightly different derivation than my coauthors and discretise Eq. (3.7), that is, I set $U(x) \to U(x_i) \equiv U_i$ with $i \in \{0, \dots N\}$.

Under an equidistant discretisation, Eq. (3.7) reads

$$\tau\left(\vec{U}\right) = \frac{1}{N^2} \sum_{i=0}^{N} e^{-U_i} \sum_{j>i}^{N} e^{U_j}. \tag{3.8}$$

Instead of setting the variational derivative $\delta\tau[U(x)]/\delta U(x)$ to zero, I consider the partial derivative $\partial\tau/\partial U_k$,

$$\frac{\partial\tau}{\partial U_k} = \frac{1}{N^2} \frac{\partial}{\partial U_k} \left(e^{-U_0}\left(e^{U_1} + e^{U_2} + \dots e^{U_N}\right) + e^{-U_1}\left(e^{U_2} + e^{U_3} + \dots e^{U_N}\right) + \dots + e^{-U_N}\right).$$

Setting $\partial\tau/\partial U_k = 0$ and de-discretising then leads to Eq. (3.9) below

$$\begin{aligned} e^{-U_k}\frac{1}{N}\sum_{j>k}^{N} e^{U_j} &= e^{U_k}\frac{1}{N}\sum_{j<k}^{N} e^{-U_j} \\ \overset{N\to\infty}{\longrightarrow}\; e^{-U(x)}\int_x^1 \mathrm{d}x'\, e^{U(x')} &= e^{U(x)}\int_0^x \mathrm{d}x'\, e^{U(x')}. \end{aligned} \tag{3.9}$$

The optimal potential has to fulfil Eq. (3.9). My collaborators used this equation to derive the following, simpler equation by differentiating twice and some algebra

$$\frac{\mathrm{d}}{\mathrm{d}x}\left(\frac{1}{U'(x)}\right) = 0. \tag{3.10}$$

The message of the above equation is that any optimal potential will be linear. The crucial question here is how Eq. (3.10) interacts with the regularisation and boundary conditions, $U(0) = U(L) = 0$. I come to that later.

But before I want to mention that one can learn more about the shape of the optimal potential from Eq. (3.7), namely, that it must be antisymmetric with respect to $x = 1/2$.

Optimal potentials are antisymmetric with respect to $x = 1/2$

The functional in Eq. (3.7) is invariant under the transformation $U(x)$ to $-U(1-x)$, that is, $\tau\left[U(x)\right] = \tau\left[-U(1-x)\right]$. With this invariance, my collaborators derived the following inequality

$$\tau\left[U(x)\right] \geqslant \tau\left[U_{\text{asym.}}(x)\right]. \tag{3.11}$$

where $U_{\text{asym.}}(x) \equiv (U(x) - U(1-x))/2$ denotes the antisymmetrised version of $U(x)$.

In other words, the above inequality expresses the fact that the antisymmetrized version of any given potential will yield shorter mean escape times.

This can be seen by applying Cauchy-Schwarz' inequality $\langle f, g\rangle^2 = \left(\int_0^1 \mathrm{d}x\, f(x)g(x)\right)^2 \leqslant \int_0^1 \mathrm{d}x\, f^2(x)\int_0^1 \mathrm{d}x\, g^2(x)$ to the equation below

$$\tau^2 = \tau\left[U(x)\right]\tau\left[-U(1-x)\right] = \int_0^1 \mathrm{d}x\, e^{-U(x)}\int_x^1 \mathrm{d}y\, e^{U(y)} \cdot \int_0^1 \mathrm{d}x\, e^{U(1-x)}\int_x^1 \mathrm{d}y\, e^{-U(1-y)}. \tag{3.12}$$

More specifically, it can be applied with the replacements $f^2(x) = e^{-U(x)} \int_x^1 \mathrm{d}y\, e^{U(y)}$ and $g^2(x) = e^{U(1-x)} \int_x^1 \mathrm{d}y\, e^{-U(1-y)}$, which leads to

$$\tau^2\left[U(x)\right] \geqslant \left(\int_0^1 \mathrm{d}x\, e^{-\frac{U(x)-U(1-x)}{2}} \sqrt{\int_x^1 \mathrm{d}y\, e^{U(y)} \cdot \int_x^1 \mathrm{d}y\, e^{-U(1-y)}} \right)^2 . \tag{3.13}$$

Another application of the inequality with $f(y) = e^{U(y)}$ and $g(y) = e^{-U(1-y)}$ then yields the desired inequality

$$\begin{aligned} \tau^2\left[U(x)\right] &\geqslant \left(\int_0^1 \mathrm{d}x\, e^{-\frac{U(x)-U(1-x)}{2}} \int_x^1 \mathrm{d}y\, e^{\frac{U(y)-U(1-y)}{2}} \right)^2 \\ &= \tau^2\left[\frac{U(x)-U(1-x)}{2} \right] = \tau^2\left[U_{\text{asym.}}\right]. \end{aligned} \tag{3.14}$$

Two Regularisations: Bounds on U and spatial discretisation

The two constraints that I discuss in the following, namely (A) bounds on U and (B) regular space-discretisation, are both compatible with antisymmetry.

Imposing bounds on the potential constitutes the simplest form of regularisation. For the sake of simplicity, my collaborators restricted the analysis to constant bounds $U_{\min} \leqslant U(x) \leqslant U_{\max}$, and refer to this constraint as A. They furthermore split the interval $[0,1]$ into three regions 1,2, and 3. In region 1, $U(x) = U_{\min}$, while in region 3, they require $U(x) = U_{\max}$. The potential is only free to change in region 2. From Eq. (3.10) it can be seen that it must be linear there. The position x^* and y^* of the intersection between the two plateaus and the linear part can be calculated from Eqs. (3.7) and (3.9), as well as the associated optimal mean exit-path time

$$\tau^A_{\text{opt}} = x^* = 1 - y^* = \frac{1}{2 + U_{\max} - U_{\min}}. \tag{3.15}$$

For symmetric bounds $U_{\min} = -U_{\max}$, the optimal potential is antisymmetric with respect to $x = 1/2$, as predicted in Sec. 3.2.1. The profile of the optimal potential under constraint A is shown in Fig. 3.9**a**. However, as the figure shows, the infinite slope at $x = 0$ and $x = 1$ renders this potential profile unrealistic and experimentally unattainable. This lead my collaborats to conclude that constraint A is flawed. They therefore turned to a different constraint B, which requires $U(x)$ to be piecewise-linear.

Such a potential profile is defined by the values U_i at equidistantly distributed points $x_i = i/n$ for $i = 0, \ldots, n$, here referred-to as n-support function. As before, the potential is chosen to be energy-neutral, that is, $U_0 = U_n = 0$. In contrast to constraint A, constraint

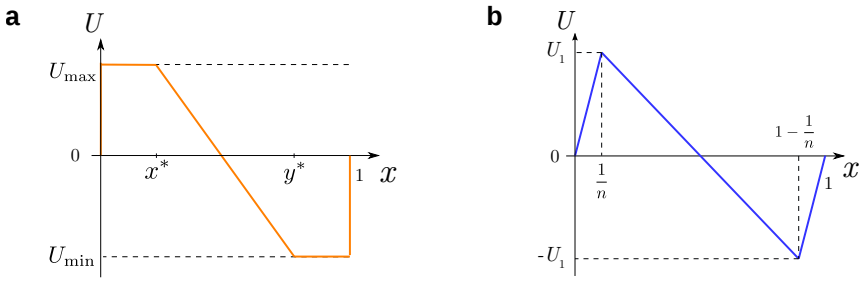

Fig. 3.9 **a** Plot of the optimal potential profile A with symmetric bounds $U_{\min}$ and $U_{\max}$. **b** N-shaped approximation of the n-support discretised solution under regularisation B with only one variable parameter U_1 (optimal potential B'). The plots were created by my collaborators Dr. M Chupeau and Prof. E Trizac.

B does not impose any restriction on the value of U_i. As I mention in the introductory example, high barriers can only speed up escapes when the width of the initial minimum vanishes. But because of the discretisation, the well width is bounded from below by $1/n$ such that a bound on the potential becomes obsolete. In order to compute the associated optimal potential profile B, my colleagues carried out simulated annealing (details of the simulation can be found in the publication [161]). In accord with the result in Sec. 3.2.1, the optimal profile is again antisymmetric with respect to $x = 1/2$ and reminiscent of an N-shape with an overshoot and an undershoot on both sides of the intermediary slide. The overshoot prevents the particle from falling back into the initial well and having to re-escape. A simple approximation of this optimal potential profile is given by an N-shaped function (denoted here as B') which is only parametrised by the potential barrier height U_1, as shown in Fig. 3.9**b**. A minimisation of Eq. (3.7) for this potential profile yields a $\ln n$-scaling of the corresponding mean-exit-path time

$$\tau_n^{B'} = \frac{1}{2\ln n} + O\left(\frac{1}{\ln n}\right). \tag{3.16}$$

Optimal potential B' captures the scaling of the mean exit-path time resulting from the sliding time τ_{slide}. More important, however, is the fact that optimal potentials B and B' do not exhibit pathologies such as infinite curvatures and should therefore be experimentally realisable in contrast to optimal potential A. In my experiments, I concentrate on the simpler version of the n-support-optimised potential, optimal potential B'.

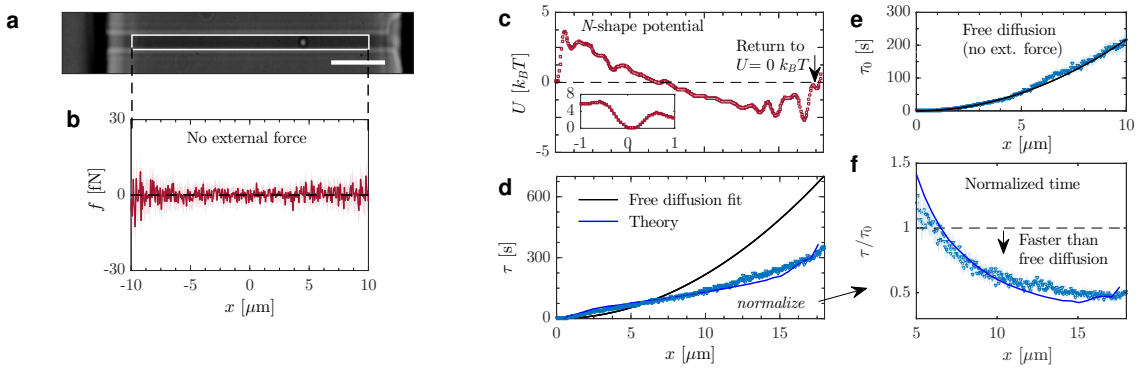

Fig. 3.10 **a** Picture of microfluidic channel used containing a single particle. The region of interest is highlighted in white. The scale bar corresponds to 5 μm. **b** Forces along the channel inferred for the zero-potential case. The error-envelope is on the order of the marker size. I here plot the force rather than the potential to highlight the accuracy of my force estimator. **c** N-shaped potential created by the HOT. The inset shows the asymmetric barrier used to approximate the initial reflecting boundary. **d** Measured first-passage times at position x for the same potential compared with the free-diffusion fit. **e** First-passage time measured symmetrically from the centre of an interval in the absence of optical forces. **f** Measured first-passage time at position x, normalized by the corresponding free diffusion time $x^2/(2D)$.

3.2.2 Experimental results

In order to test whether experimental potentials can be tailored to deliver the predicted speed-ups, I leveraged the wave-shaping abilities of my holographic optical tweezer (HOT) to create the potential shape associated with the optimal potential B' (see Fig. 3.9**b**) in the focal plane of a microscope. In addition, I again used a microfluidic device to confine movements of colloidal particles, to a quasi-one-dimensional line, eliminating entropic forces and variations in hydrodynamic drag [166, 167] (see next section). As discussed in Sec. 1.1.1, motion of colloidal particles is well within the overdamped regime, such that our theory applies. All experiments were carried out using my automated "drag-and-drop" routine (see Sec. 2.2.4).

As a first step, I measured first-passage times τ_0 of a colloid released in the centre of a channel, shown in Fig. 3.10**a**, without the influence of laser forces (see panel **b**). As the data in Fig. 3.10**e** shows, these times adhere closely to theory and scale quadratically with distance. From this data set, I infer a diffusion coefficient of $D = 0.23\,\mu\text{m}^2/\text{s}$, by fitting a square function $\tau(x) = ax^2$ to the measured mean exit-path times.

Next, I create the optimal potential shape B' as shown in Fig. 3.10**c**. The holographic parameters necessary to form the right balance of intensity-gradient and phase-gradient forces were found by educated guesswork. Specifically, the N-shaped potential was created

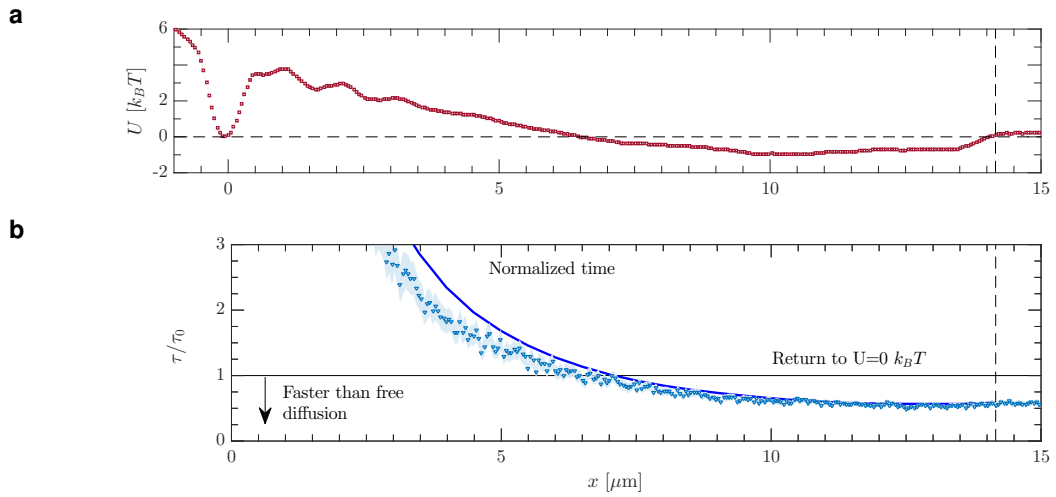

Fig. 3.11 Second potential, which reliably shows faster-than-free diffusion escapes. **a** Inferred potential profile. **b** Mean-first passage time along x.

by a combination of a single point trap providing the initial potential well and three line traps with phase-gradients and lengths as specified in Tbl. 3.1. The reflecting boundary condition is approximated by an asymmetry in the initial well depth (see inset in Fig. 3.10**c**). The experimentally measured (triangle symbol) and the theoretically predicted (blue line) mean exit-path times are shown in panel **d**. At the point where the potential returns to its initial value, the mean exit-path time reaches a value of 336 ± 19 s. For comparision, the corresponding free-diffusion value is 684 s, so that I obtain a speed-up factor of exactly 2.

The speed-up also appears in potentials that do not really resemble an "N". In Fig. 3.11, I give another example of a potential with fewer oscillations towards the end, near $x = 14\,\mu$m. The figure shows that the rate-boost effect is robust against slight variations of the potential. However, the speed-up here is lower than the one achieved in Fig. 3.10. In Fig. 3.11**b**, I measure a mean exit-path time $\tau = 251.2 \pm 20.4$ s at the x-position demarked by the dashed vertical line. Compared to a free diffusion 432.2 s, this yields a speed-up factor of only around 1.72.

In the experiments corresponding to Figs. 3.10 and 3.11, I create the initial minimum using an asymmetric potential well. This turns out to be the right approach. Previous to this, I attempted to use the fact that Kramers escape times are fundamentally similar to exit-path times, if the potential is symmetric with respect to the initial position. However, creating precisely symmetric potentials is naturally difficult. Moreover, this approach requires twice the interval length of the single-exit approach. The resultant symmetric potential landscapes therefore barely fit into the 20 μm channels.

In addition, I also explored physical realisations of reflecting boundaries, i.e. dead-end channels. However, high hydrodynamic drag close to the dead-end resulted in lower overal speed-up. A higher probability of the colloid getting stuck at the dead-end wall against which it must be initially pressed rendered this approach impractical.

3.2.3 Unsuccessful experiments, potential causes, and mitigating measures

The process of guessing holographic parameters is slow and involved since almost all parameters have some degree of cross-talk. Changing the intensity of one optical trap will likely affect all other intensities and shapes. It is illuminating to review a few examples of experiments where predicted and measured first-passage times do not agree well. In Fig. 3.12, I show a few of the potential landscapes that I created during the optimisation process. I only plot the relevant part of the x-axis, from $x = 0$ until the point $x = L$ where the potential returns to $U = 0\, k_B T$. Every parameter study typically lasts several days since the potential can only be reliably estimated once sufficient data is collected.

As data in Fig. 3.12**b** shows, the theoretical prediction does not always reproduce the measured first-passage time very well. Given the large amount of data collected, it can be very difficult to ascertain the precise cause for such a mismatch between theory and experiment. In order to facilate future research on similar scenarios, I state possible reasons below:

1. Transient binding of the colloid to the PDMS walls of the channel or the glass surface: Transient binding can be identified by visually inspection of individual trajectories $x(t)$. During such events, the trajectory appears constant for brief periods of time (few milliseconds to seconds).

2. Time-dependent potential landscape: The potential inference used here is unable to measure variations in the potential landscape. Such variations, however, may well arise due to the following causes:

   (a) *Unstable settings of the autofocus routine.* A plot of the z-position of the piezo over the course of the experiment may contain clues as to whether this is a problem or not. Especially oscillatory behaviour of the z-position is a strong indicator of an unstable autofocus.

   (b) *Optical drift.* Fundamentally, the operation of the HOT setup rests on the assumption that the virtual and physical position of optical traps coincide. Due to thermal variation, the physical position of a laser trap may wander over

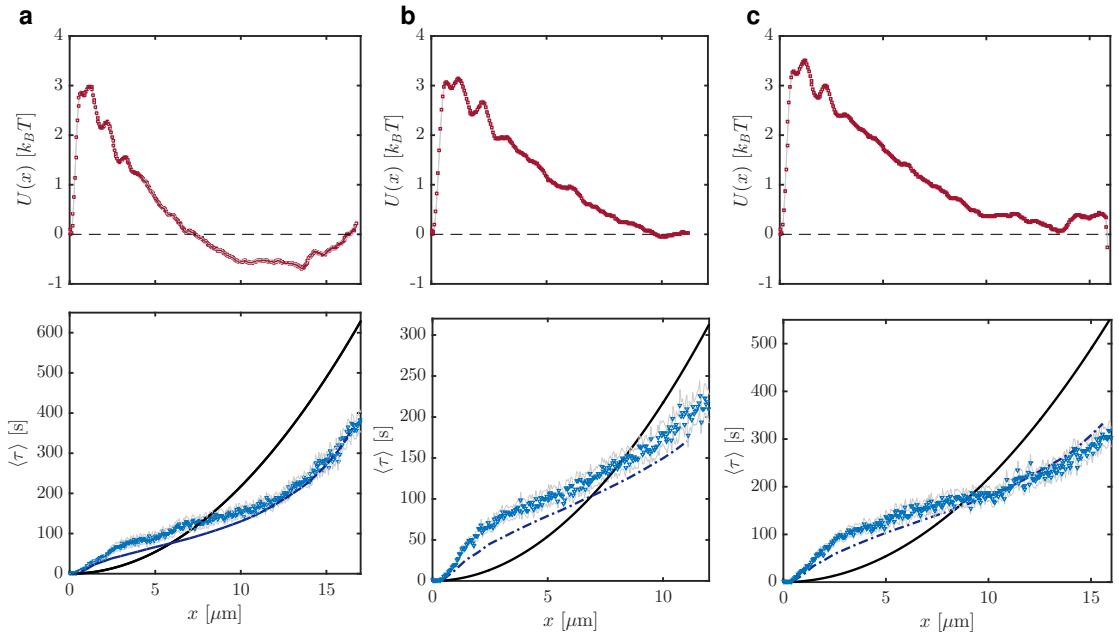

Fig. 3.12 Study of the effects of various central trap intensities. While the upper row shows the inferred potential landscape, the row on the bottom show predicted (dash-dot) and measured (triangle) first-passage times. The solid black line in the plots on the bottom is a plot of the quadratic free diffusion time. The thin line connecting the potential energy data points in the upper row represents the standard error of the mean, which is here visibly smaller than the red markers. The value of the central trap intensity parameter I varies across the panels: **a** $I = 0.205$, **b** $I = 0.215$, **c** $I = 0.225$. See Tbls. 3.1 and 3.2 for comparison.

time. The hologram software, however, has no intrinsic means to detect such a drift. In order to mitigate this problem, I implemented a deviation detection, which periodically compares the average position of a trapped colloid to the virtual trap position and corrects the latter accordingly. In case of suspected drift, this routine should be activated or the virtual-to-physical space mapping should be updated manually.

(c) *Variations in particle size.* Optical trapping forces are dependent on particle size. While the autofocus routine usually ensures re-use of the same colloid, it may occasionally loose track of a colloid and instead grab a different one. This is especially likely during autofocus adjustments caused by external shocks such as slammed doors. In most colloidal suspensions used in this thesis, the size dispersion is small enough for this not to be measurable problem. However, the routine can occasionally grab small floating objects such as debris. If multiple populations of particle sizes are visible, the experiment

should be immediately aborted and the microfluidic chip and solution should be exchanged. A trajectory-to-trajectory comparision of diffusion coefficients may contain clues about the sizes of the particle that have been used. Another variable of interest in such a case is the intensity of the imaged particles.

(d) *Transient hydrodynamic flows.* Pressure-driven flows through the channel are one of the main causes of experiment-theory deviations. Flows that are permanent will skew the inferred potential landscape and hence not be source of deviation. However, the absence of transient variations in flow intensity cannot be guaranteed. If small, such variations are almost impossible to detect. In general, prior to an experiment one should check whether colloidal in-channel motion is biased. If so, the microfluidic chip should be exchanged.

3.2.4 Conclusion

Building on theory developed by my coauthors, I explored escape rates over fine-tuned barriers and showed that such rates can exceed rates obtained for zero-potentials. The increase in rate was achieved without altering the free-energy balance of the system, the colloid ends on the same free-energy as it starts. While the optimisation of holographic parameters is slow, iterative, and requires educated guesses, I eventually managed to demonstrate a doubling of escape rates in N-shaped potentials. This particular shape was predicted to be optimal in terms of escape rates under constraint B, as discussed in the text. I omit here the part of the theoretical work of my colleagues which shows that the speed-up effect extends into the inertial regime. However, I regard this addition as significant, since it demonstrates that the speed-up effect is not a pathology of the otherwise sometimes ill-behaved overdamped regime. Furthermore, I also omit the theoretical discussion of the effect in higher dimensions, which shows that the rate-boost remains relevant as the dimensionality increases. Both parts feature in our publication [161].

To conclude, I believe that our result will have a profound impact on studies where escape rates matter, for instance, in the aforementioned studies of protein folding.

3.2.5 Experimental methods

The same colloids and the same buffer solution was used as in the previous section. However, the holographic pattern used were considerably more involved. Table 3.1 contains the parameters used to create the N-shaped pattern shown in Fig. 3.10**c**, while Tbl. 3.2 summarises the parameters used to create the potential in Fig. 3.11**a**. As both table show,

Type	I	x_c	L	p
Line	0.2	$-0.91\,\mu$m	$2\,\mu$m	1
Point	0.09	$0\,\mu$m	N.A.	N.A.
Line	1.1	$8.05\,\mu$m	$16\,\mu$m	0.1
Line	0.6	$17.45\,\mu$m	$3\,\mu$m	-1

Table 3.1 RedTweezers parameters used to create the potential $U(x)$ in Fig. 3.10**c**.

Type	I	x_c	L	p
Line	0.2	$-31\,\mu$m	$2\,\mu$m	1
Point	0.205	$-32\,\mu$m	N.A.	N.A.
Line	1.1	$-39\,\mu$m	$14\,\mu$m	-0.06
Line	0.25	$-50.803\,\mu$m	$6\,\mu$m	0.35
Point	1	$-30.28\,\mu$m	N.A.	N.A.

Table 3.2 RedTweezers parameters used to create the potential $U(x)$ in Fig. 3.11**a**.

In both tables, I denotes the parameter controlling the relative intensity, x_c represents the centre of the line or point trap, L denotes the length of line traps, and p represents the phase-gradient parameter introduced in Sec. 3.1.6. Trap centre positions x_c are given in absolute distance to the virtual position of the zeroth order.

almost all parameters required fine-tuning, often down to two subdecimal digits. Only one parameter can be optimised at a time due to their non-linear interdependence.

3.3 Transition-path times in corrugated channels

While the microfluidic channels in the previous sections in this chapter have played no role other than effectively reducing the dimensionality of the colloidal random walk, I here consider the effect of microfluidic channels with spatially-varying widths, which result in a two-dimensional problem.

Brownian motion inside a tube of varying diameter is a classic and much-studied problem [168, 166, 169–175]. According to Zwanzig, even Adolf Fick, one of the founding fathers of modern diffusion physics, had considered this problem [166]. In any case, judging by the number of publications, the problem appears to be of interest again. This is probably so for three reasons: first, the varying diameter offers a perspective on the effects of entropic forces and secondly, the coarse-grained diffusion coefficient appears to scale in an interesting way with the shape of the channel. Thirdly, the discovery of 'entropic splitting' reinvigorated interest in these systems [174, 175, 145]. However, almost all modern studies are of theoretical nature, with little experimental input. Crucially, few studies incorporate or discuss spatially-varying drag coefficients $\gamma(x)$ which will likely play a huge role in microfluidic realisations of these systems due to the rich and surprising hydrodynamics of confinement. Such hydrodynamic effects have been studied in detail by our group [44, 45].

The results presented in this section are based on experiments that I carried out with my colleague Dr. Karolis Misiunas. He kindly provided me with a microfluidic silica mold with a number of periodically shaped channels on. As I mentioned in the beginning of this chapter, this line of work never got published since we stopped working on it when a competing group unexpectedly published ahead of us [135].

3.3.1 Fick-Jacobs theory

In the first section (Sec. 1.1.1), I write about 'integrating out' degrees of freedom. Here, I marginalise over the variable y, which I define as the lateral position, perpendicular and with respect to the channel axis ($y = 0$). Let $U(x, y)$ be the potential a Brownian particle is subject to in the channel, then

$$e^{-\frac{U(x)}{k_B T}} = \int_{-w(x)/2}^{w(x)/2} \mathrm{d}y\, e^{-\frac{U(x,y)}{k_B T}} \tag{3.17}$$

defines the marginalised version. Crucially, $U(x)$ will pick up an entropic component if the channel width $w(x)$ varies over x. Our channels are two-dimensional, only the width

changes, not the height to the ceiling. For a flat, hard-sphere potential $U(x,y)=0$, the above equation amounts to $U(x)=-k_BT\ln(w(x))$, a classic entropic potential.

In the absence of forces, one can write down an approximate FPE called Fick-Jacobs equation,

$$\frac{\partial\rho}{\partial t}(x,t)=D\frac{\partial}{\partial x}\left(A(x)\frac{\partial}{\partial x}\left(\frac{\rho(x,t)}{A(x)}\right)\right), \tag{3.18}$$

where $A(x)$ denotes the cross-sectional area at x.

Robert Zwanzig published an article in 1992 explaining that one can generalise Fick-Jacob's equation to cases with a potential using Eq. (3.17) and, at the same time, improve on its accuracy by introducing a spatially-varying diffusion coefficient [166]

$$\frac{\partial\rho}{\partial t}(x,t)=\frac{\partial}{\partial x}\left(D(x)e^{-\frac{A(x)}{k_BT}}\frac{\partial}{\partial x}\left(e^{\frac{A(x)}{k_BT}}\rho(x,t)\right)\right) \tag{3.19}$$

where $D(x)/D_0\approx 1-\frac{1}{3}w'(x)^2$. D_0 here denotes the microscopic diffusion coefficient, that fulfils some fluctuation-dissipation theorem such as Eq. (1.15). If one wants to actually compare theory and microfluidic experiments, one has to incorporate the spatially-varying drag $\gamma(x,y)$ as well [135]. The position-dependency of D above is just one of several approximations that have been put forward in the literature [176]. Testing the accuracy of Eq. (3.19) and comparing the various entropic diffusion corrections constituted the original aim of this project.

3.3.2 Experimental results

Using my automation routine, I acquired several data sets each comprising around 300 uninterrupted trajectories, up to 10 minutes in length, of equilibrium Brownian motion in channels of varying diameter. In Fig. 3.13, I plot mean transition-path times across two different distances (black box in the upper image). The box is continuously moved across the field of view, such that I can plot two continuous curves: mean transition-path times left-to-right (blue) and vice-versa (red).

Interestingly, the behaviour of the times appears to crucially depend on the length of the box: For larger distances (panel Fig. 3.13**a**), the expected entropic speed-up seems to dominate. This situation is comparable to the right minimum of transition-path times in Fig. 3.2**d**, which occurs over a maximum of the potential in Fig. 3.2**a**. It is important to remember that the transition-path times considered in this thesis are conditional first-passage times, which are conditioned on reaching their destination in one go. Without

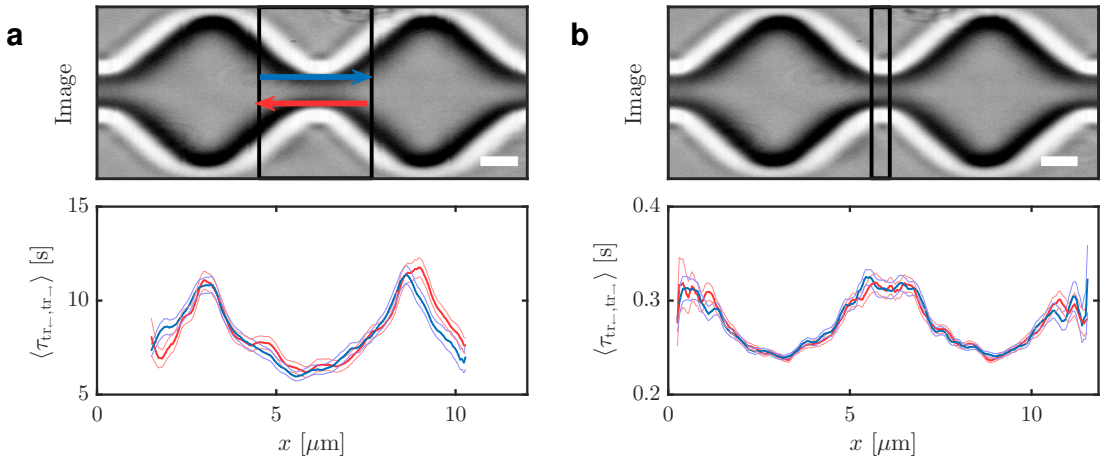

Fig. 3.13 Transition-path times across distances (black box) in channels of varying diameter. The x variable in the lower row corresponds to the horizontal position of the centre of the black rectangle shown in the upper row. The rectangle is dragged along the channel axis, which results in a continuous curve of transition-path times, similar to the curve in Fig. 3.2**d**. The scale bar in the images corresponds to 1 μm. **a** Transition-path times across larger distances are dominated by entropic effects. Across the constriction, transition-path times are at a minimum. **b** Across shorter distances, hydrodynamic effects dominate transition-path times. The increased drag in the constriction slows transitions down. This leads to an inversion of the shape of the curve relative to panel **a**.

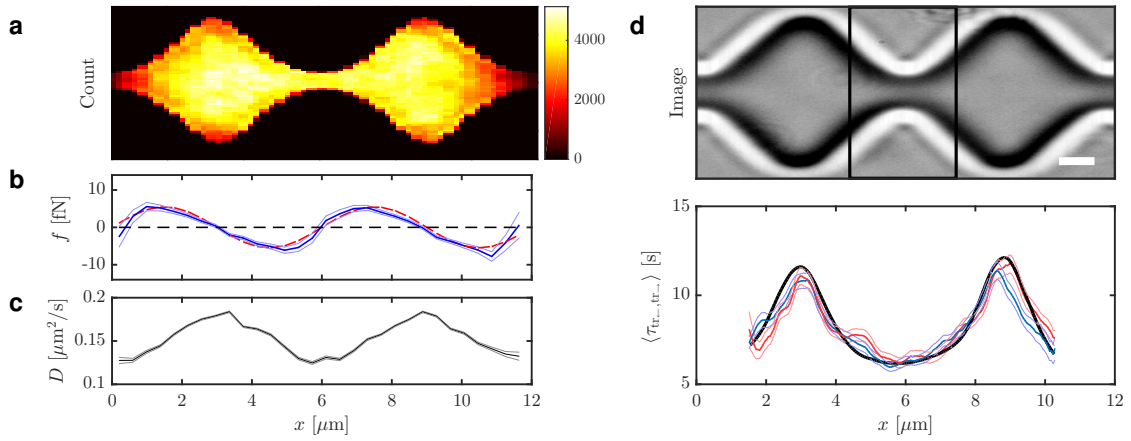

Fig. 3.14 Comparison with theoretical mean transition-path times. **a** Histogram of localisation statistics in the channel. **b** Force estimate from equilibrium trajectories (blue). Only x-position is considered for estimation. The red-dashed curve shows a sinusoidal fit. **c** Estimated microscopic diffusion coefficients. The variation is not due to entropic effects, it is caused by varying hydrodynamic drag. **d** Reprint of data from Fig. 3.13**a** with theoretical predicion of mean transition-path time (black).

this condition, trajectories traversing a potential maximum or entropic constriction would of course be impeded and the corresponding transition would appear slowed-down rather than sped-up. By contrast, over short distances (panel Fig. 3.13**b**), a different effect seems to dominate direct transitions: Here, the increased drag in the constriction slows the dynamics down.

The entropic forces caused by the confinement are directly measurable. If I apply my force-estimator Eq. (3.4) on the x-position of the trajectories, I obtain the curve shown in Fig. 3.14**b**. An approximately sinusoidal force is what one would expect from Eq. (3.17), since the width in this channel behaves roughly as $w(x) \approx a + b\sin(kx)$ with $a = 5\mu$m , $b = 4\mu$m, and $k = \pi/(5\mu\text{m})$ (I note that these are effective widths for the colloidal centre point). I thus have $w_{\min} = 1\mu$m and $w_{\max} = 9\mu$m. For a hard-sphere-like potential $U(x,y) = 0$, I find $U(x) = -k_B T \log(w(x))$ and thus $f(x) = -\partial_x U(x) = k_B T \frac{w'(x)}{w(x)} = k_B T \frac{kb\cos(kx)}{a+b\sin(kx)}$, which when plotted resembles a standard sine-curve $\sin(x)$ with triangular bulges. The degree of deviation from a standard sine-curve will grow with the ratio a/b. The careful observer will find a slight asymmetry in the bulges in Fig. 3.14**b**.

The diffusion can be estimated too, which yields the curve in panel Fig. 3.13**c**. The variation here shows again the position-dependence of hydrodynamic drag in confinement discussed in Fig. 3.13**b**.

Using Eq. (1.52), I can now compute the mean transition-path time across the black box in the microscope image in Fig. 3.14**d**. The result is shown in the plot below in

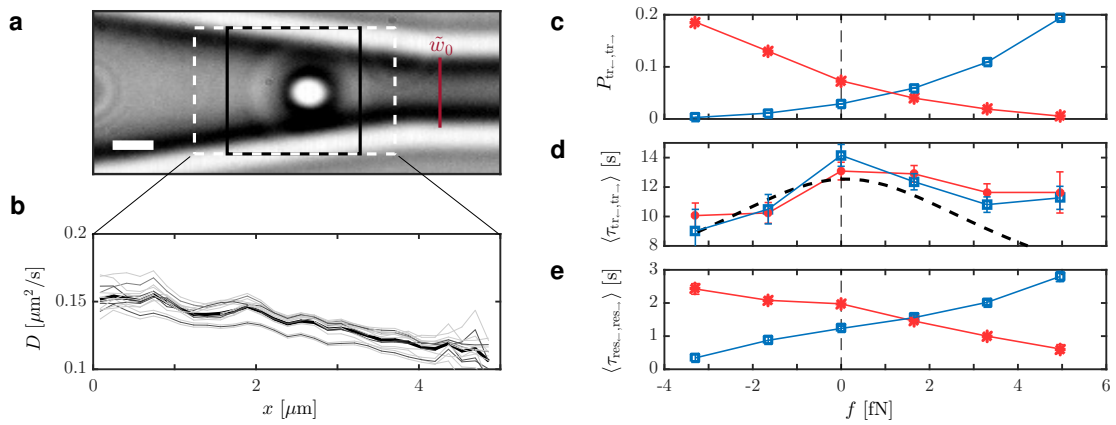

Fig. 3.15 Colloidal dynamics in a funnel filled with a line trap. **a** Image of the funnel. The scale bar represents 1 μm. **b** Diffusion coefficients (grey) with standard error of the mean envelope for all phase-gradients used. The mean diffusion coefficient is plotted in thick black. **c** Transition probabilities across the black box as a function of the applied force. **d** Mean transition-path times in both directions across the back box. **e** Mean residence time of all trajectories leaving to the left (red) and to the right (blue). This includes returning trajectories.

black. In this plot, I assumed a constant diffusion coefficient of $D = 0.155\,\mu\mathrm{m}^2/s$, but otherwise used the force-fit $f(x) = A\sin(kx+\varphi_0)+f_0$ shown in Fig. 3.14**b**. Again, I find that transition-path time theory is applicable to experiments in a self-consistent way.

In the following, I turn to a different channel shape: the funnel shown in Fig. 3.15**a**. The idea of the experiment described here is to counter entropic with actual forces. To this end, I fill the funnel with a line trap and repeat essentially the experiment in Fig. 1.4**d**,**e**. The entropic forces exerted by funnel which gets wider at a constant rate λ are given by

$$f_{\mathrm{entrop}}(x) = -\partial_x U(x) = k_B T w'(x)/w(x) = k_B T\lambda/(w_0+\lambda x). \tag{3.20}$$

In panel **c**, the transition probabilities over various phase-gradient forces are shown. The offset to the right of the force that precisely counters the entropic effect is clearly visible. The minimal width in this system is $\tilde{w}_0 = 2\,\mu$m, where the tilde indicates that I have not yet accounted for the width of the colloid, which is $d = 1\,\mu$m in this experiment. The effective width that the centre of the colloid experiences is thus $w_0 = 1\,\mu$m. The rate of change of the width can be approximated to $\lambda \approx 0.3$, by measuring the width at various points. For instance, the left side of the white-dashed region-of-interest in Fig. 3.15**a** has a width of $\tilde{w}_1 = 3.6\,\mu$m, so the effective width is $w_1 = 2.6\,\mu$m. The x-width of the this region is $w_x \approx 5\mu$m. Eq. (3.20) then yields $f_{\mathrm{entrop}}(x = 5\mu\mathrm{m}) \approx 0.5$ fN and $f_{\mathrm{entrop}}(x = 0) \approx 1.2$ fN, roughly in the regime of the experimentally determined break-even force $f \approx 1.15$ fN.

Importantly, the phase-gradient forces for this experiment were calibrated within the straight part of the channel, in the absence of entropic influences.

The mean back-and-forth transition-path times measured across the black box in panel **a**, agree in both directions for all phase-gradient parameters. Interestingly, the theoretical prediction Eq. (1.52) fails to capture the slight asymmetry of the experimental transition-path times with respect to zero force.

Before I conclude, I want to briefly discuss the residence time plotted in panel **e**. The residence time is defined as the weighted average of the average return time and the translocation time, $\langle \tau_{\mathrm{res}} \rangle = P_{\mathrm{ret}} \langle \tau_{\mathrm{ret}} \rangle + P_{\mathrm{tr}} \langle \tau_{\mathrm{tr}} \rangle$ [90]. This time is not direction independent and, crucially, are typically measured in systems in which return and translocation events cannot be distinguished, such as resistive-pulse sensing in nanopores [149]. Residence times might therefore offer a way to infer the orientation of asymmetrically-shaped biological channels in electrophysiology experiments.

3.3.3 Conclusion

In this section, I investigate the kinetics of entropically-driven Brownian motion in confinement. My results repeatedly underline the importance of hydrodynamic friction as a function of distance to the confining PDMS enclosure. In accord with theory, I find that a spatially-varying width leads to the emergence of an entropic force pointed along the channel as shown in Fig. 3.14. Interestingly, the speed-up of transition paths caused by this force is partially countered by increased friction near constrictions. Even though the two effects scale differently with the distance to the constriction, there must be a length scale over which they cancel to some extent, which should result in a flatter transition-path time profile (see Figs. 3.13).

Over the last decade, a number of publications have discussed whether a change in drag causes measurable, albeit spurious forces [177–180]. According to theory [179], such forces would behave as $f_{\mathrm{spurious}} = -\gamma(x)\partial_x \left(k_B T/\gamma(x)\right)$. In the case of the funnel geometry in Fig. 3.15, this force can be approximated from the slope of the diffusion coefficients in panel **b**. I find a value of $f_{\mathrm{spurious}} \approx 0.35$ fN, which might be just below what I can reliably detect, given the presence of confounding factors such as entropic forces and (often present) latent hydrodynamic flows. However, measuring spurious forces in more abruptly changing geometries such as channel entrances could constitute an interesting project in the future.

In the following, I briefly discuss another aspect of channel transport that is of importance in biological systems: selectivity.

3.4 Brownian dynamics simulations of channel transport

Previous to this thesis, Pagliara *et at* used the predecessor of the HOT setup to study optimal parameters and shapes of channel-filling potential landscapes [156]. The quantity that is being optimised in their study is the flux of colloids through the channel, irrespective of the direction. However, in biological systems it is often not just the total flux through a channel, e.g. a membrane channel protein, that is of interest, but relative fluxes of different species of, say, ions or proteins [142]. To assist and inform the design of my experiments, I set up three-dimensional Brownian dynamics simulations featuring a channel geometry and line-shaped potentials.

Here, I discuss some of the results of these simulations. In particular, in Sec. 3.4.2 I describe a recipe to spot violations of detailed balance in Brownian dynamics simulations that may for instance originate from flawed implementations of interaction potentials.

3.4.1 White-list selectivity via channel-entrance binding

The aforementioned HOT study found that line traps which protrude from the channel into the bulk increase the transition rate of colloids drastically [156]. The question I wish to investigate here is whether entrance binding can also be used to make channels selective for affine species. If a channel attracts a sufficient amount of the selected-for species, these particles might then block particles of other species from translocating. The idea is thus to compare the fluxes of two species, one of which, the affine species, feels the potential to a greater extent than the other (non-affine species).

Details of the simulation

In Fig. 3.16**c**, I plot the number of attempts per hour. The dimensions in the $x \times y$-plane are shown in panel **a**. The $x \times z$-plane perspective is exactly similar. The grey region in panel **a** marks the wall, pierced by the channel. The boundary conditions in x are periodic and reflecting in the y- and z-direction.

At $t = 0$, 10 particles of two species, A and B, are initialized at random positions inside the box, excluding the channel. Species A is attracted by the trap shown in panel **a**, while species B is not. However, instances of the two species do interact with each other as hard-spheres with radius $r = 250$ nm. All particles have a diffusion coefficient of $D = 0.8\ \mu\text{m}^2/s$. Position and velocity of each particle are stepped using an Euler-step of the (underdamped) Langevin Equation, Eqs. (1.3) and (1.4) for purely viscous kernels $\hat{\alpha} = i\omega k_B T/D$. The timestep Δt in all simulations is set to $\Delta t = 0.2\,\mu$s, the total simulation amounted to $T = 400$ s. However, I write results to file only in every 5000^{th} step. The

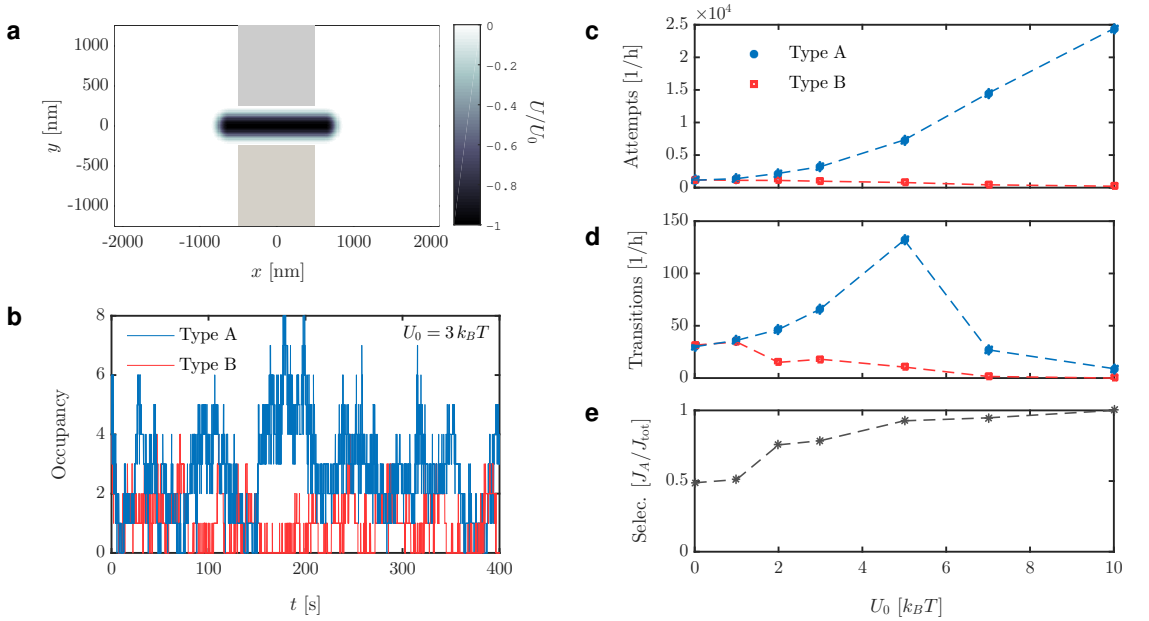

Fig. 3.16 Brownian dynamics simulations overview. **a** Geometry of the wall, channel and line-trap in the $x \times y$-plane. **b** Example trace of the occupancy of the channel by instances of the two species. **c** Number of attempts as defined as number of times a particle entered the channel region. **d** Number of translocations of species A and B as a function of potential depth U_0. **e** Selectivity, defined as the number of translocations of species A relative to the total number of translocations, over potential depth. At $U_0 = 0$, the channel is unselective as expected.

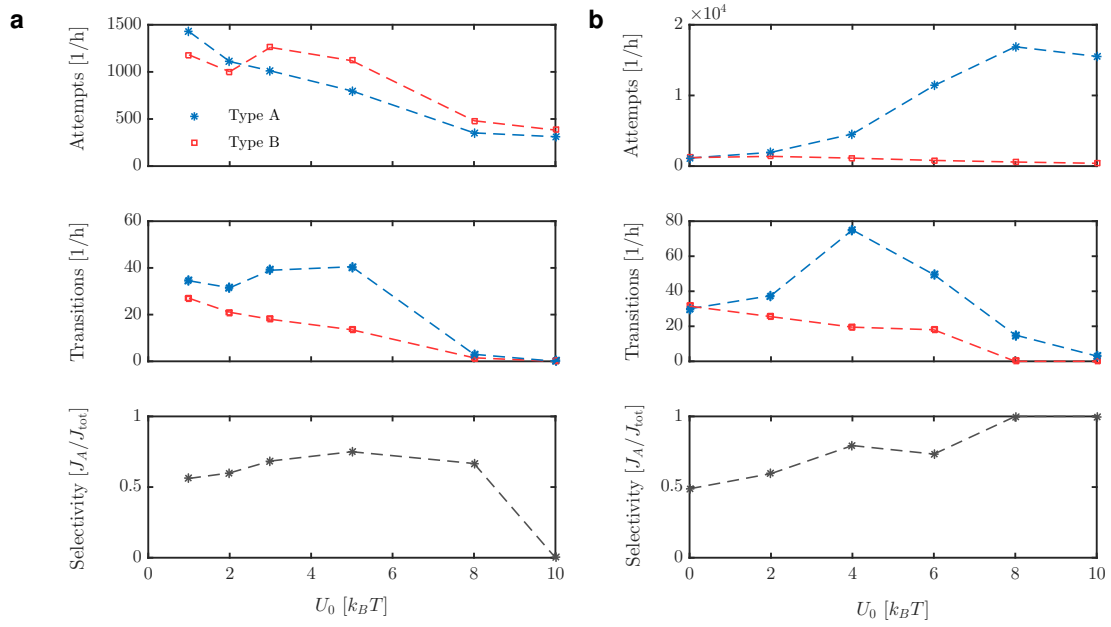

Fig. 3.17 Further results of the simulations. **a** Competition for line traps that do not protrude from the channel between affine (A) and non-affine species (B). Selectivity is markedly reduced compared to Fig. 3.16. **b** Similar simulations to Fig. 3.16, but species B is made partially affine, $U_{\mathrm{B}}(\vec{x}) = 0.2U_{\mathrm{A}}(\vec{x})$.

program is parallelized with up to 12 instances running concomitantly. The simulation was implemented in the Brownmove package [181, 182], which supports the channel geometry shown in Fig. 3.16**a**. I customised the package and added the potential, affinity and parallelisation.

Results

In panel **b**, I show a typical occupancy time series over all 400 seconds. The preference for type A is clearly visible. Pagliara *et al* report an optimum depth of the binding potential between $2-4\,k_B\,T$. The translocations in panel **d** clearly indicate a maximum around $4\,k_B\,T$ with higher potentials leading to a jamming and thus reduced fluxes. The attempt rate appears to be monotonically increasing over the interval of potential depths considered here. The selectivity for type A, defined as the translocations of A J_{A} divided by the translocations of all particles, J_{tot}, increases from no preference (selectivity = 0.5) to fully-selective (selectivity = 1) as shown in panel **e**. These results are interesting when contrasted with the results shown in Fig. 3.17**a**, where I repeat the same simulations but with line-traps that only fill the interior of the channel.

The selectivity, again shown in bottom panel, is markedly reduced. Overall, translocation numbers are subdued and attempt rates are magnitudes lower than in channels with entrance-binding. The reduction in translocation can be explained with the difficulty of exiting the attractive region inside the confining channel. It is generally more probable to escape from a potential within a given period of time, the larger its exit-surface area. In panel **b**, I show results where species B has an affinity of 0.2 for the trap, i.e. feels a potential $U_{\mathrm{B}}(\vec{r}) = 0.2U_{\mathrm{A}}(\vec{r})$. While the channel largely remains selective, the number of translocations peaks around $U_0 = 4\,k_B T$ before decreasing to zero. This could be due to jamming, since more particles compete for the trap. For low translocation numbers, selectivity is not meaningful in a statistical sense, since both numerator J_{A} and denominator J_{tot} might be close to zero.

Conclusion

Taking only attempt and translocation numbers into account, I conclude that the simulations corroborate Pagliare *et al*, who found that entrance binding strongly increases channel translocations. I here show, that entrance binding can also create a white-list-type selectivity, in which a positive selection for affine species boosts their translocation numbers sufficiently such that other species get crowded out of the channel area. This effect, however, requires a stark difference in affinity between the species as the last simulation indicates in Fig. 3.17**b**.

My original intent was to reproduce the simulations using the HOT setup and let two types of colloidal particles compete for access to microchannels filled with line-shaped traps. Particle species consisting of different materials with different refractive indices would feel the optical traps to a different extent, such that one species would be more affine for these traps than the other. I planned to differentiate the two species using fluorescence (hence the epi-fluorescence LED in the setup, see Fig. 2.6). However, the green fluorescence signal of 500nm melamine particles (microParticles GmbH) used for this purpose, turned out to be too weak and, crucially, appears in a different optical plane than the out-of-focus plane typically used for imaging colloids.

3.4.2 A method to locate sources of non-equilibrium dynamics

Progress in this project was hampered by a *bug* in the Brownmove package. The bug consists in a broken symmetry of wall-repulsion potentials near the sharp corners encountered at the mouth of the channel. The repulsive Lennard-Jones potential used to ensure that no particle can diffuse into obstacles was implemented in such a way that

detailed balance was violated in transitions near the channel entrances. More precisely, the force vector in response to transitions into the wall close to the corner from either the outside or inside of the channel pointed into the channel with a slightly higher frequency than back into the bulk. The bug only occurred for ball-objects with finite volume, where the geometry of the object surface in relation to the wall or corner needs to be considered. In the following, I briefly describe the method I used to locate the bug, since it might be of help for future projects. The method is inspired by a former project I worked on, which attempted to establish non-equilibrium indicators for biological matter [155, 61]. The idea is to look for points in phase-space where detailed balance breaks down. In a finite data set, detailed balance will, of course, never be perfectly maintained. However, as I indicated in Sec. 1.1.3, a violation of detailed balance in steady-state will cause circular probability currents. Such patterns are unlikely to develop by chance; the presence of circular currents is thus strong indication of non-equilibrium activity. Secondly, one can try to estimate error bounds for the direction and length of probability currents in discretised spaces as explained in Battle *et al* [155].

In Fig. 3.18, I plot discretised probability fluxes over the $x \times y$-plane of the simulation. The boundary conditions here were set to be reflective everywhere.

As the Fig. 3.18 clearly shows, the flawed stochastic dynamics near the channel entrances break detailed balance and cause local circular probability currents. The external potential in panel **a**, which protrudes from the channel only on one side of the barrier, rectifies these currents into a current through the channel. This results in a steady-state imbalance in particle numbers between the left and right side of the channel as shown in panel **c**.

3.5 Conclusion

In this chapter, I explore several different flavours of first-passage times on the microscale using my automated holographic optical tweezers setup. In particular, I demonstrate exit- and transition-path time symmetries over confined pathways and show a breakdown of this symmetry once the system is forced to diversify transition pathways. Moreover, I demonstrate that escape rates across fine-tuned energy-neutral barriers exceed Kramers rates by up to a factor of 2. I find excellent agreement of experimentally measured transition-path times in channels of varying width with predictions based on marginalised, entropic potentials and demonstrate that entropic forces in such channels can be counterbalanced by directed optical scattering forces. Finally, I show in Brownian dynamics simulations that channel-entrance binding can result in a white-list type selectivity for

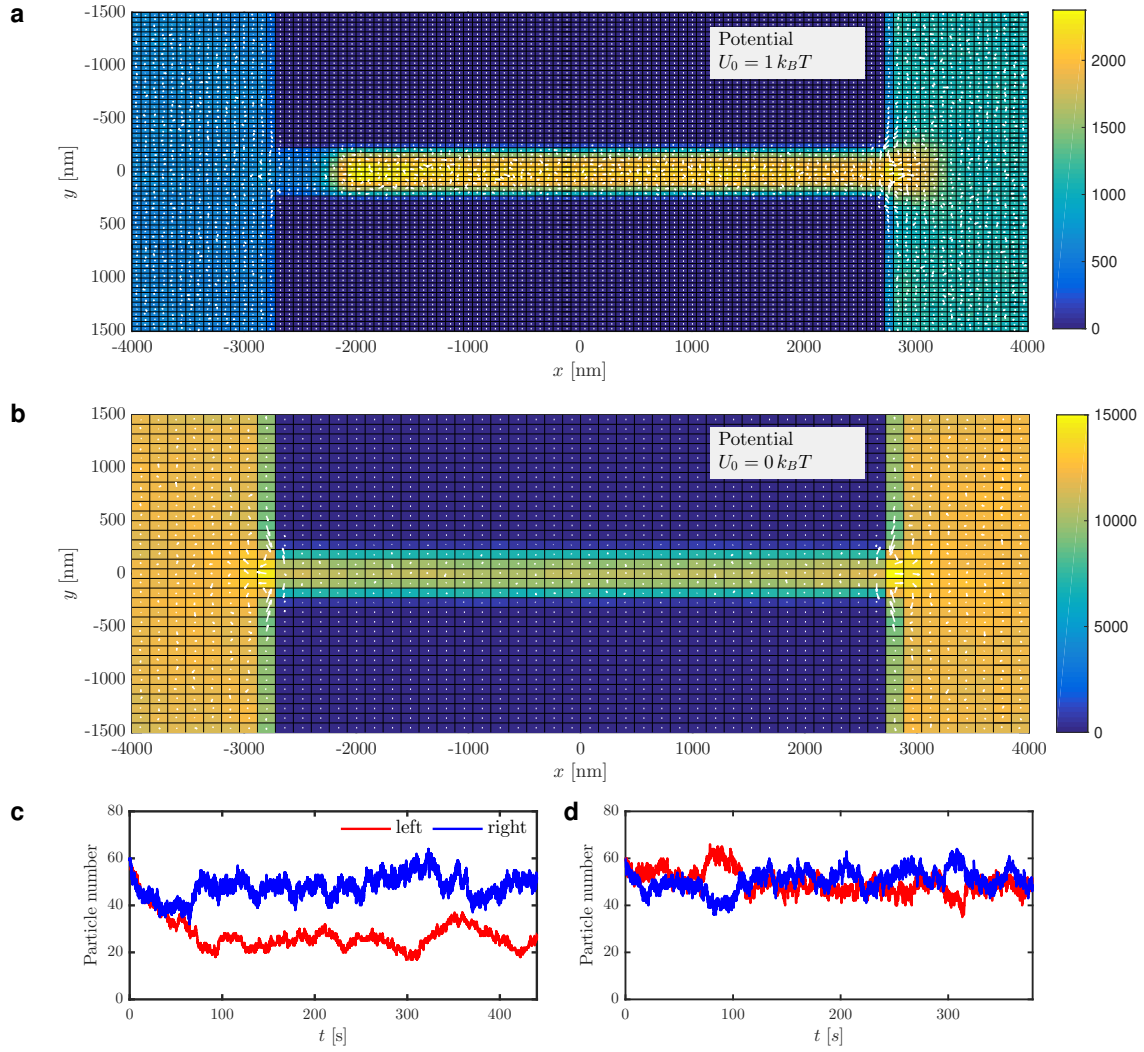

Fig. 3.18 Broken detailed balance in Brownian dynamics as bug indicator. **a** An asymmetric line-trap which only extends on one side of the simulation rectifies the flawed entrance dynamics into a probability current. **b** The effect can also be observed, albeit on a smaller scale, in channels without any external potential. The grid size here is increased to capture the effect. **c** Particle numbers on the left and right side of the channel for the scenario in panel **a**. **d** Particle numbers without the rectifying influence of the line-trap.

Brownian particles which feel the attracting potential to a greater extent than competing species. Whenever possible, I directly compare my experimental results with first-passage-time predictions based on overdamped Fokker-Planck theory. Overal, the results in this chapter turn out to be in excellement agreement with theory, often without any free parameters.

In the future, combining the theoretical and experimental methods and approaches used in Secs. 3.1 and 3.2 might result in fruitful projects. For instance, one could compute the optimal potential landscape filling a microfludic channel by applying variational calculus to Eq. (1.51). Of course, the crucial question here is what optimality in this context means. In the case of Eq. (1.51), one would optimise the speed of transitions across a channel while requiring a certain potential difference ΔU between both ends 3.2. The profile of the resulting optimal potential could be implemented using a combination of optical and entropic potentials, since the results obtained in Sec. 3.3 support Fick-Jacobs-based theory.

As both my simulation results in Sec. 3.4 and previous results from our group show [156], the accessible surface of a channel-solute-interaction potential to the bulk is of significant importance. Differences in surface area between the channel entrances can even rectify localised non-equilibrium effects, such as the flawed interaction potential mentioned in Sec. 3.4.2, into a permanent particle flux through the channel (see Fig. 3.18**c**, **d**). Potential-geometry-related effects, such as 'entrance binding', could also be accounted for in a numerical optimisation of channel transition times, which should lead to more widely applicable conclusions.

Finally, the experimental difficulties described in Sec. 3.2 indicate a need for better algorithmic tools to create precise light intensity landscapes on the micronscale. To address this need, I describe a proof-of-principle realisation of what is, to my knowledge, a novel approach to the optical inverse problem encountered in holographic tweezing in the next chapter.

Chapter 4

Inverse digital holography using conditional generative models

4.1 Ill-posed inverse problems

The term 'inverse problem' can mean different things, depending on the context. In this thesis, it refers to finding the inverse $f^{-1} : y \to x$ to a function $f : x \to y$, or an approximation thereof.

Of course, using this definition, the inverse of the inverse problem is the original problem of finding the 'forward' mapping $f : x \to y$. For well-behaved invertible functions, finding the inverse from a given data set is as difficult as finding the forward function: One simply exchanges the order in the inference engine. However, many mappings and relationships in experimental situations are not invertible in the sense of mathematical functions. Under such circumstances, mathematicians speak of 'ill-posed inverse problems'. A problem is ill-posed if it is not well-posed. In the context of finding the (an) inverse mapping, well-posedness requires, according to J. Hadamard, (1) the existence of a solution, (2) uniqueness of that solution, and (3) that the solution has a continuous dependency on the data (stability) [183].

In Sec. 2.2.5, I introduce digital holography as an inverse problem. In this section, I explore if some of the novel machine learning techniques that have been publicised in the last few years might help solve the inverse problem that appears in digital holography: *Given a desired intensity distribution* $\mathbf{I}$ *in the image plane, find one or all holograms* $\mathbf{h}$ *that get closest to reproducing* $\mathbf{I}$.

This inverse problem is certainly ill-posed. Since, as I show in Sec. 2.2.5, at least the second condition (uniqueness) is violated. The first condition (existence) is fulfilled whenever the algorithm is asked to recreate a previous intensity distribution (disregarding noise). This is the natural situation that occurs when an intensity is drawn from a data set, where the existence of hologram that lead to that distribution is guarantueed. The third condition (stability), is more involved. For single Fourier components, Fig. 2.9**b** suggests a smooth, albeit non-monotone dependency of the laser intensity on Fourier amplitudes. The situation is less clear for mixtures of components. Barring possible intensity-thresholds of the camera sensor, there appears to be no physical reason to expect discontinuities.

While a number of deep learning frameworks are publicly available, such as tensorflow or pytorch, I chose to publish yet another full-fledged, modular deep-learning library, written in C++. The library can be loaded into LabView as external DLL and is available on my GitHub page [184]. The idea was to train the models discussed in this section directly on the setup, which would have enabled the learning algorithm to interact with the problem and, for instance, request the light pattern that corresponds to a newly created hologram during training. Even though the library is fully-functional and supports almost all of the normalisations or gradient-descent optimisers required here, the advantages of online learning turned out to be smaller than expected. Quantity of training data and speed of training are more important factors. I therefore followed the typical machine-learning workflow and trained all models on fixed data sets and used GPU-accelerated tensorflow-based code for swift training.

Finally, I want to point to a connection between thermodynamics and deep learning. In the latter, the goal usually is to minimise some sort of error function by means of stochastic gradient descent. As has recently been pointed out, this problem can be mapped to a Langevin-gradient descent of a free-energy function as encountered in problems such as protein folding [185]. One of the key lessons of this mapping is that deeper networks seem to generate error-landscapes with minima that feature larger basins and hence larger entropies. This appears to be one of the factors of the success of deep neural networks when compared to large, but shallow networks. More generally, this result indicates that lessons from protein folding may indeed help explain the success of deep learning.

But before I introduce the problem in this chapter, I proceed with a concise introduction to neural-network-based regression in order to put my thinking into context and introduce relevant notation.

4.1.1 Introduction to neural-network-based regression

In this section, I introduce the concepts underlying neural-network-based regression ("curve fitting") and, importantly, define the notation used in this chapter.

A regression example using polynomials. In a typical regression scenario, one is given a data set of input-output pairs of an unknown functional relationship

$$\mathscr{D} = \{(x_1, y_1), (x_2, y_2), (x_3, y_4), \ldots, (x_N, y_N)\}, \tag{4.1}$$

with x, y respectively representing input and output values.

Conceptually, neural-network-based regression is very similar to regression using other, perhaps more familiar function families, such as polynomials. I shall therefore begin this introduction using polynomials. In its most simple form, polynomial regression aims at finding an optimal set of coefficients $\theta^* = \{p_0, \ldots, p_n\}$ for a polynomial of the form

$$\hat{y}_{\theta^*}(x) = p_0 + p_1 x + \cdots + p_n x^n \tag{4.2}$$

which minimises the mean-squared-error

$$\mathscr{E}(\theta, \mathscr{D}) = \sum_{i=1}^{N} \left(\hat{y}_\theta(x_i) - y_i\right)^2. \tag{4.3}$$

Mathematically speaking, the aim is to perform $\theta^* = \text{argmin}_\theta \mathscr{E}(\theta, \mathscr{D})$. Importantly, the function in Eq. (4.2) has the desirable property of a linear dependence of the output value to the parameters θ. Indeed, when the parameters are rewritten as a vector, $\vec{\theta}^T = (p_0, p_1, \ldots, p_n)$, Eq. (4.2) can be written as $y_{\vec{\theta}}(x) = \vec{\theta}^T \vec{\phi}(x)$ where $\vec{\phi}^T(x) = (x^0, x^1, \ldots, x^n)$. The key advantage of *linear* regression is that the optimal parameter vector $\vec{\theta}^*$ can be written down immediately. Following Bishop's notation [186], I write $\vec{t}$ as the vector containing all target values $\vec{t} = (y_0, y_1, \ldots, y_N)$ and $\Phi_{k,l} = \phi_k(x_l)$ as a matrix where each element consists of the kth-basis function (here $\phi_k = x^k$) evaluated on the lth-input from the training set, x_l. The optimal choice of parameters for a given polynomial degree n is then

$$\vec{\theta}^* = \left(\Phi^T \Phi\right)^{-1} \Phi^T \vec{t}. \tag{4.4}$$

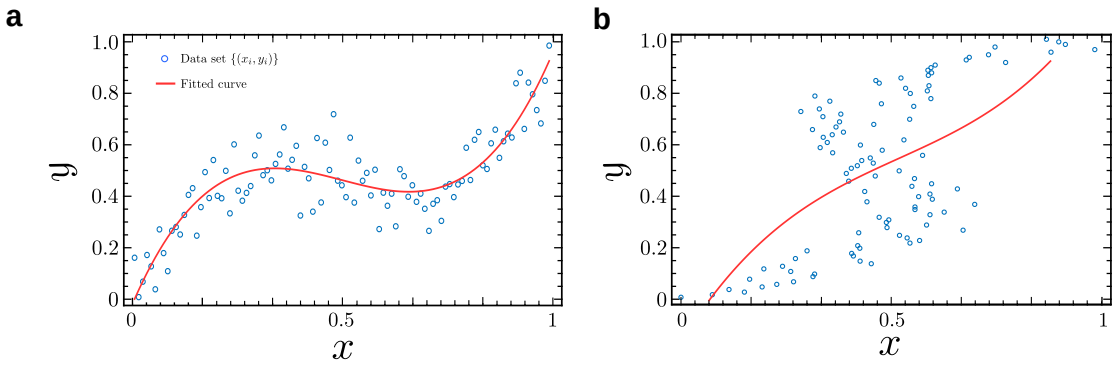

Fig. 4.1 Polynomial regression of a data set of noisy input-output pairs of a function. **a** Forward regression of the x-to-y dependence. **b** Inverting the order (y-to-x) of the regression leads to severe problems if the underlying function is not invertible: polynomials, like any other function, cannot assume multiple values at once.

In Fig. 4.1**a**, I show an example of polynomial regression with $n = 3$ (cubic) for $N = 100$ data points. The true function that I wish to approximate here is $y(x) = x + 0.3\sin(2\pi x)$, which is perturbed by Gaussian noise $\xi \sim \mathcal{N}(0, \sigma^2 = 0.01)$ on the output end.

In contrast to the successful regression in Fig. 4.1**a**, I show what happens when I switch the role of x and y in each data point in Fig. 4.1**b** [186]. Like all mathematical functions, polynomials cannot assign more than one value to an input. Any attempt to fit polynomials to distributions of x, y-pairs, which are multimodal in the y-direction will result in a potentially severe approximation error, since the fit will try to average the different y-modes. This is a simple example of an ill-posed inverse problem.

Neural-network-based regression is non-linear. Neural-network-based regression is analogous to the example of polynomial regression shown above in the sense that it too seeks to minimise a given error function, $\mathscr{E}(\theta, \mathscr{D})$. However, an important difference is that the dependency between output y and parameters θ is non-linear, which makes it impossible to directly find the optimal parameter set. Instead, one resorts to an iterative procedure which gradually descends the loss function in parameter space, an approach that is aptly named 'gradient descent'. In a popular variant of this approach, one computes the gradient $\vec{\nabla}_\theta \mathscr{E}(\theta, \mathscr{B})$ only for a small '*batch*' $\mathscr{B} \subset \mathscr{D}$ of data points instead of over all data points in $\mathscr{D}$. In each step n, a different batch $\mathscr{B}_n$ is chosen. These batch gradients are then successively applied to the parameters as shown below in Eq. (4.5),

$$\theta^{n+1} = \theta^n - \eta \vec{\nabla}_\theta \mathscr{E}(\theta^n, \mathscr{B}_n) \tag{4.5}$$

where η is referred-to as learning rate. Since this succession of gradients will randomly switch direction based on the data contained in the respective batch, this approach is known as 'stochastic gradient descent' and has become a cornerstone of modern neural-network-based learning [186].

For the sake of simplicity, I will consider here simple neural-networks which only consist of fully-connected layers. A fully-connected layer performs a matrix multiplication of its inputs with a 'weight' matrix, followed by an addition of a bias term and a non-linearity. The operation of a single p-q-layer can be summarised as follows

$$\begin{pmatrix} \hat{y}_1 \\ \vdots \\ \hat{y}_p \end{pmatrix} = h\left(\begin{pmatrix} w_{11} & w_{12} & \dots & w_{1q} \\ \vdots & \vdots & \ddots & \vdots \\ w_{p1} & w_{p2} & \dots & w_{pq} \end{pmatrix}\begin{pmatrix} x_1 \\ \vdots \\ x_q \end{pmatrix} + \begin{pmatrix} b_1 \\ \vdots \\ b_p \end{pmatrix}\right) \tag{4.6}$$

with $h(x)$ denoting some non-linear function, which is applied element-wise. A few non-linear functions are particularly popular for deep learning such as $\tanh(x)$, the sigmoid $\sigma(x) = (1+e^{-x})^{-1}$, and the rectified-linear function $\mathrm{ReLu}(x) = \max(x,0)$ [186]. The aim is thus to find optimal values for a set of parameters $\theta = \{w_{11}, w_{12}, \dots, w_{pq}, b_1, \dots, b_p\}$.

In Fig. 4.2, I give an example of neural-network-based regression using the same sample function as used in Fig. 4.1. The succession of increasingly visible red lines in Fig. 4.2**b**, **c** shows the convergence of the output of the network sketched in Fig. 4.2**a** to the noise-less ground truth, $y(x) = x + 0.3\sin(2\pi x)$. Like polynomial regression, neural-network-based regression with a l_2-loss averages and interpolates noisy data. However, due to the sequential nature of stochastic gradient descent, neural-network training requires large datasets. In this case, convergence to a loosely N-shaped curve requires on the order of $100-1000$ training samples and several thousand batched gradient descent steps at a learning rate of $\eta = 10^{-4}$. The required ratio of training samples to adjustable parameters is a question of ongoing research: Large nets with many adjustable parameters have the risk of memorising training data which can lead to a deteriorating ability to generalise [186].

Implementing gradient descent using backpropagation

In order to implement stochastic gradient descent for the operation described in Eq. (4.6), one has to find a way to calculate the gradient in Eq. (4.5) with respect to each parameter in θ. Of course, modern deep networks consist of more than one layer. Fortunately, the rules derived for gradient calculation for a single layer can be applied iteratively to each layer in a deeper network. All one has to do is to specify how gradients flow from one layer to the

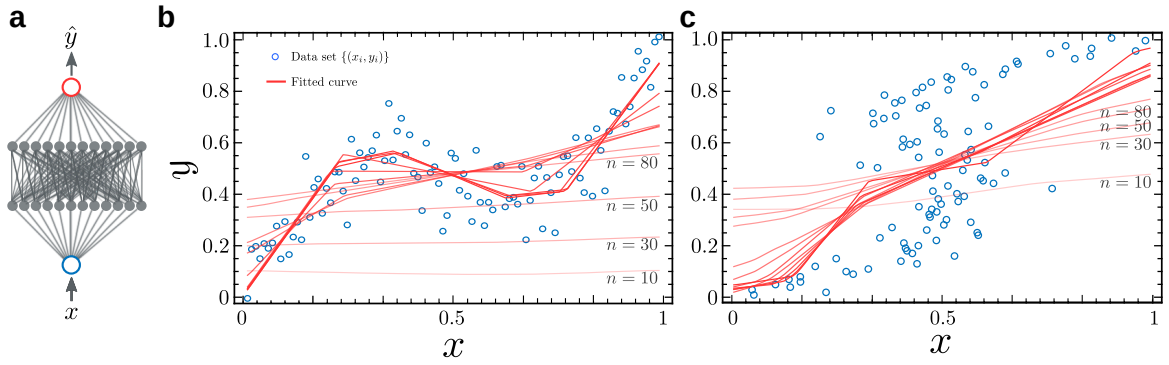

Fig. 4.2 Regression using neural networks. **a** Sketch of a 3-weight-layer neural network with 12 hidden units in each hidden layer and one in- and output. **b** Convergence of the neural-network-based regression over training steps. The intensity of the red colour encodes the stage of training. **c** Same plot as in panel **b**, but with inverted roles of x and y.

next. Successive layers are viewed as functions that are iteratively applied to the ouput of the previous function. The gradient of a deeper layer can therefore be calculated using the chain rule of differentiation, which leads to a set of recursive equations that allow for an error propagation through the layers. This technique is known as backpropagation [186].

The rules of backpropagation can be derived by calculating the gradient of an error function, e.g. $\mathscr{E}(\{x,y\};\theta) = \sum_{n=1}^{N} ||\hat{y}(x_n) - y_n||_{l_2}^2 \equiv \sum_{n=1}^{N} E_n$ with respect to each parameter, i.e. $\partial E_n/\partial w_{ij}^l, \partial E_n/\partial b_i^l$ where l marks the layer number. Again, following Bishop [186], I define abbreviations for the layer output $a_i^l = \sum_j w_{ij}^l z_j^{l-1} + b_i^l$ with z_i^l denoting the activation, $z_i^l = h(a_i^l)$.

An application of the chain rule to $\partial E_n/\partial w_{ij}^l$ yields $\partial E_n/\partial w_{ij}^l = \partial E_n/\partial a_j^l \partial a_j^l/\partial w_{ij}^l$. The definition of a_j^l can then be used to simplify the second factor, $a_j^l/\partial w_{ij}^l = z_i^l$, while the first factor itself is abbreviated as $\partial E_n/\partial a_j^l = \delta_j^{l+1}$, such that one arrives at

$$\frac{\partial E}{\partial w_{ij}^l} = \delta_j^{l+1} z_i^l. \tag{4.7}$$

Even though the layer indexing may appear strange at first, the underlying logic is that the to-be-updated weight matrix w_{ij}^l will receive a gradient which consists of the product of activations coming from below, z_i^l and δ's coming from above, i.e. δ_j^{l+1}. Crucially, the δ's can be passed from layer to layer using again the chain rule: $\delta_i^{l+1} = \partial E_n/\partial a_i^l = \sum_m \partial E_n/\partial a_m^{l+1} \partial a_m^{l+1}/\partial a_i^l$, where the index m runs only over upstream nodes which are connected to node i. This sum can be simplified using the definition of δ_i^l, which leads

to [186]

$$\delta_i^{l+1} = h'(a_i^l) \sum_m w_{mi}^{l+1} \delta_m^{l+2}. \tag{4.8}$$

A C++-implementation of this logic can be found in the *FullyConnectedLayer*-class in my custom-written deep learning library [184]. The last remaining brick is the identification of the direct derivative of the error with respect to the network output with the most-upstream δ, i.e. $\delta_i^{L+1} = h'(a_i^L)(\hat{y}_i - y_i)$, where L is the total number of layers.

A single step in neural-network training has therefore three phases: (1) Forward pass (compute and save activations a_i^l), (2) error computation and backpropagation (compute and save δ_i^l), and (3) computation and application of the gradient, $w_{i,j}^l = w_{i,j}^l - \eta \delta_j^{l+1} z_i^l$ and $b_i^l = b_i^l - \eta \delta_j^{l+1}$. In matrix notation, the last step is written as $\mathbf{w}^l = \mathbf{w}^l - \eta \delta^{l+1} (\mathbf{z}^l)^T$.

Although the regression examples presented here are fairly simple, they do not differ in terms of the workflow from the more complex types of regression undertaken in this chapter. An important difference is that the models trained in this chapter seek to regress distributions, rather than functions.

4.2 Problem definition & inference models

In recent years, deep-neural-network-based machine learning has revolutionized a number of inference problems, such as image comprehension [187–193] or natural language processing [194–196]. During training, for instance on classification tasks, the network shapes a flow of increasingly abstract representations, from layer to layer, suppressing unnecessary features, while relevant ones are amplified [197]. As shown in the preceding section, under a supervised scheme, training refers to a minimisation of a loss function $\mathscr{E}$, e.g. $\mathscr{E}(\hat{y}, y) = ||\hat{y} - y||_{l_2}^2$ over a labelled data set, such that the neural network output $\hat{y}(x)$ approaches the target output y given the input x. The crucial question is then whether the network will accurately predict the corresponding target value of previously unseen inputs and distill features that are relevant for generalisation [186].

This, however, necessitates the existence of some underlying generalising structure in the data. Such a structure can be conceptualised as a mapping $f : x \to y$, for instance, between images of handwritten digits and their numeric label. This mapping is usually only implicitly present in the data and needs to be inferred. Crucially, such an approach is limited to modelling data relations that can be expressed by functions in the mathematical sense, i.e. by one-to-one- or many-to-one-type relations. Many relevant real-world

problems, such as inverse holography, however, are ill-posed inverse problems, amenable only to certain algorithms including generative models [198, 199].

Before I introduce generative models, I briefly define the physical problem that I wish to solve in this chapter. Furthermore, I describe the choice of features made here as well as the process of data acquisition, which is shown in Fig. 4.3**a** and **b**.

4.2.1 Dimensionality reduction by Fourier-feature preselection

Only a small number of hologram patterns actually result in light being recorded on the camera. Most patterns correspond to spatial frequencies that redirect the laser beam somewhere else into the setup. To reduce the number of "empty" patterns and overall complexity of the problem, I restricted the holograms **h** to only include frequencies from a preselected 8×8-matrix in Fourier-space, $f_{i,j}$ as shown in Fig. 4.3**a**. These 64 Fourier components are selected such that when one displays any of these frequencies $f_{k,l}$ in isolation, i.e. $f_{i,j} = f_0 \delta_{i-k,j-l}$, the SLM directs light to a region of interest (ROI) on the camera, with pixel dimensions 100×100. By restricting these 64 Fourier components to the real plane, I eliminate any relative phase difference between the Fourier components (see Fig. 2.9**a**). Global phase offsets, such as the example shown in Fig. 2.8, are prevented from occurring in this configuration too.

My workflow of acquiring data sets is sketched in Fig. 4.3**b**: I begin by setting all corresponding matrix coefficients to zero. Then, I draw 64 random numbers from a uniform distribution and reconsider those coefficients where the respective random number surpasses a predefined sparseness-threshold C. The so-selected coefficients are then redrawn from the distribution in order to ensure that Fourier coefficients can take up all possible values $f_{i,j} \in [0, 100]$. This is done in order to avoid overloading the holograms, which would lead to low light intensities. The sparseness is thus a data-set parameter (typically $C \in [90, 97]$). By restricting the values to the interval $[0, 100]$, I eliminate even the bi-modality shown in Fig. 2.9**b** at least for Fourier matrices with only a single non-zero component. However, pixel values in holograms corresponding to Fourier matrices with several non-zero components can still exceed the 100-mark, since their hologram pattern superimpose. Its highly non-linear nature in conjunction with imaging noise, such as shot noise, and the remaining high number of variables (10k + 64) still make this problem enough of a challenge.

Once the 8×8-matrix is obtained, it is embedded into a 200×200-matrix of zeros, on which I perform a Fast Fourier Transformation (FFT) and subsequently take the real part (see Fig. 4.3**a** and **b**). The resultant 200×200-hologram **h** is then embedded centrally in

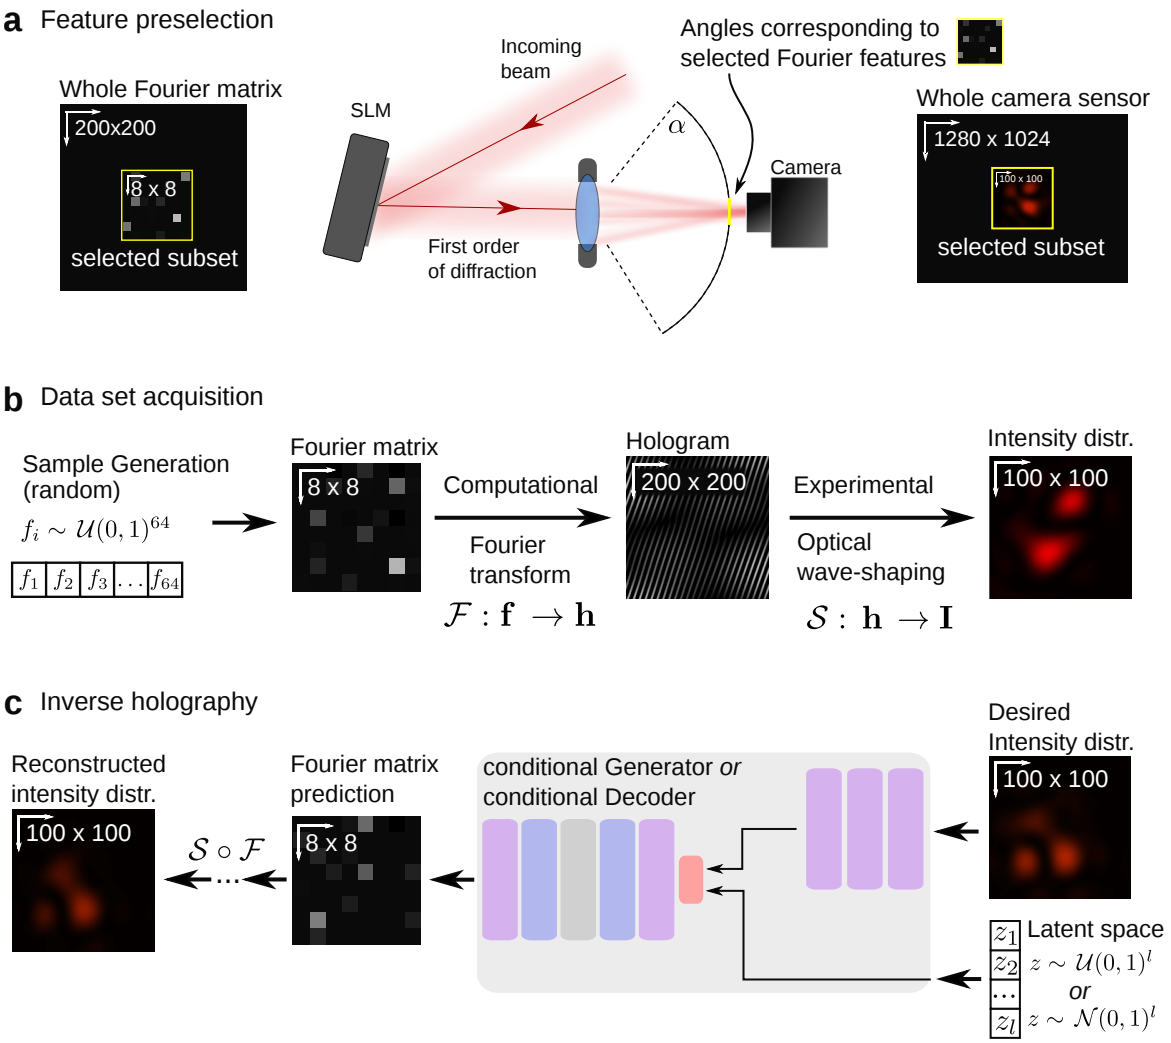

Fig. 4.3 **a** I consider the intensity within a 100x100 pixel sized region-of-interest (ROI) on the camera. Only a small number of spatial frequencies send rays towards that ROI, so that one can reduce the dimensionality in Fourier space considerably to an 8x8 matrix. All other Fourier components are always set to zero. **b** Workflow of sample acquisition approach: Generate real-valued random 8×8 Fourier matrices $\mathbf{f}$, embed them centrally into the overall 200x200-Fourier matrix, calculate the corresponding 200x200-hologram $\mathbf{h}$, display it centrally in the 800x600-sized SLM, and measure the ensuing 100x100-intensity distribution $\mathbf{I}$. **c** The conditional generative model is then trained on the mapping $(S \circ F)^{-1} : (\mathbf{I}, z) \rightarrow \mathbf{f}$, omitting the intermediate hologram $\mathbf{h}$. In principle, conditional generative models could be used to learn $S^{-1} : (\mathbf{I}, z) \rightarrow \mathbf{h}$, which is more computationally demanding.

the 800×600-pixel matrix of the SLM. Before feeding the intensity and Fourier matrices into the models, I normalise all values such that $I_{i,j} \in [0,1]$ and $f_{i,j} \in [0,1]$.

4.2.2 Conditional generative models

The approach that I describe here consists in training conditional generative models in order to (re)construct 8×8-Fourier matrices $\hat{\mathbf{f}}$ from a given intensity distribution $\mathbf{I}$. Conceptually, the approach is not entirely novel; generative models in inverse problems have been applied to molecular design [200] or computer tomography [201].

Conditional Generative-Adversarial Networks

Arguably, the best-known generative model are 'generative-adversarial networks' (GAN). Ever since their discovery, GAN have received attention due to their expressive abilities to create realistic, but purely synthetic pictures [202, 203]. However, from a mathematical viewpoint, GAN are an approach to learning a probability distribution $\rho(x)$ in a generative way [198, 204].

More precisely, a GAN is a scheme aimed at training a generative model such that it is able to morph "latent space" vectors $z = \{z_1, \ldots, z_l\}$ drawn from a standard probability distribution into samples $\hat{x} = \{\hat{x}_1, \ldots, \hat{x}_m\}$ that appear to be drawn from the data set. This ability of GAN to learn distributions, can be used to solve one-to-many-type inverse problems $f^{-1} : y \to \{x_k\}_{k \in \mathscr{K}(y)}$ where $\mathscr{K}(y)$ is the subspace over which $f(x) = y$. To this end, the GAN must be conditioned on the output variable y (here the intensity $\mathbf{I}$) to approximate the conditional distribution $\rho(x|y)$ [205]. Once trained, one can generate multiple estimates $\hat{x}$ (here Fourier matrices $\hat{\mathbf{f}}$) by redrawing latent space vectors $z \in Z$ for a given y and test for the optimum.

Since GAN are notoriously difficult to train, a number of normalisation schemes have been proposed to prevent an otherwise likely loss of multi-modality of the learned distribution $\rho_\theta(\hat{\mathbf{f}}|z,\mathbf{I})$, known as mode collapse [206–209]. This approximative distribution is controlled by the generator G_θ. It is only defined by samples drawn from the generator, i.e. $\hat{\mathbf{f}} = G_\theta(z,\mathbf{I})$. I trained all conditional GAN-models by applying gradient descent to the following loss function [198]

$$\begin{aligned}\mathscr{E}_{\text{cGAN}}(\{\mathbf{f},\mathbf{I}\};\phi,\theta) = \mathbb{E}_{\mathbf{I}\sim\rho(\mathbf{I})}\left[\mathbb{E}_{\mathbf{f}\sim\rho(\cdot|\mathbf{I})}\left[v\left(D_\phi(\mathbf{f},\mathbf{I})\right)\right] + \mathbb{E}_{z\sim\mathscr{U}^l}\left[w\left(D_\phi(G_\theta(z,\mathbf{I}),\mathbf{I})\right)\right]\right] \\ + \beta\mathscr{E}_{\text{rec}}\left(\{\mathbf{f},I\};\theta\right)\end{aligned} \tag{4.9}$$

with respect to the parameters controlling the discriminator D_ϕ and generator G_θ, ϕ and θ. As discussed in the introduction in Sec. 4.1.1, the parameter sets, ϕ and θ, contain all trainable weights and bias terms of the respective underlying neural network. As Nowozin *et al* [204] show, Eq. (4.9) leads to an approximate minimisation of the symmetric Jensen-Shannon divergence $J_{\mathrm{SH}}(\rho||\rho_\theta)$ between the data distribution ρ and its approximation, ρ_θ.

The functions $v(\cdot)$ and $w(\cdot)$ in Eq. (4.9) depend on the type of GAN: in standard GAN, for instance, they are given by $v(x) = \log(\sigma(x))$ and $w(x) = \log(1-\sigma(x))$ with $\sigma(x) = (1+e^{-x})^{-1}$. In practice, $w(x)$ is often changed to $w(x) = -\log(x)$ [198]. Importantly, using this notation, the discriminator network does not apply non-linear functions to its output nodes. Expectations in Eq. (4.9) are estimated by running averages over small batches of data (mini-batching). The last term in Eq. (4.9) denotes a reconstruction loss, implemented throughout this chapter as $\mathscr{E}_{\mathrm{rec}}(\{\mathbf{f},\mathbf{I}\};\theta) = L^{-1}\mathbb{E}_{\mathbf{I},\mathbf{f},z}\left[||\mathbf{f}-\hat{\mathbf{f}}(z,\mathbf{I})||_{l_2}^2\right]$ with $L = 64$, the dimensions of $\mathbf{f}$. The hyperparameter β controls its relative importance in relation to the rest of the GAN-loss. The flow of gradients for cGAN models is sketched in Fig. A.5**a**. The cGAN model and its training workflow are described in further detail in the appendix Secs. A.2.1 and A.2.2.

Wasserstein GAN (WGAN) are a particular type of GAN, which I use in this study. In WGAN, one seeks to minimise the Wasserstein distance between the distribution that the GAN implicitly learns and the target distribution [210]. For WGAN, the functions in Eq. (4.9) are simpler than in standard GAN and read $v(x) = x$ and $w(x) = -x$. However, the discriminator-network in WGAN is required to fulfil a K-Lipschitz condition, $|D(\mathbf{f}_1) - D(\mathbf{f}_2)| \leqslant K||\mathbf{f}_1 - \mathbf{f}_2||_{l_2}$ for a scalar K, which was originally achieved by weight-clipping. Other methods have since been put forward to ensure the Lipschitz condition and render training of WGAN more stable, such as gradient-penalty [211]. Furthermore, some of the benefits of Lipschitz-constrained discriminators have been shown to extend to standard GAN as is exemplified by spectrally-normalised GAN, which exhibit increased expressiveness and reliable training progress [208].

Conditional VAE model

Importantly, GAN do not attempt to structure their latent space Z, which may complicate finding a suitable z-vector for minimal loss (see example in Sec. 4.2.2). By contrast, Variational Autoencoder (VAE) [212] enforce a particular distribution over $z \in Z$; typically a standard Gaussian is chosen. Similar to GAN, VAE can be trained in a conditional way to approximate $\rho(y|x)$ while preserving their regularised latent-space structure. VAE consist of two networks, a decoder and an encoder. The latter is trained to transform input and

conditioned-on variables x, y into a standard-normally distributed latent space vector $z \sim \mathcal{N}(0, \mathbf{1})$. The decoder then generates a sample $\hat{x}$ from a newly drawn latent space vector z' and conditional variable y. The expected output of the trained model is thus solely determined by the conditional variable y.

In addition to the three GAN models (cGAN, cWGAN, and cGAN with reconstruction loss), I trained a conditional VAE (cVAE) on my data set. As mentioned before, VAE consist of two networks. The primary network is the decoder, which controls the probability of Fourier candidates $\hat{\mathbf{f}}$, $p_\theta(\hat{\mathbf{f}}|z, \mathbf{I})$ as a function of latent space vectors z and desired intensity $\mathbf{I}$. During training, the latent space vector is drawn from a distribution $q_\phi(z|\mathbf{f})$ which is controlled by the second, the encoder network. Once trained, the encoder can be removed; all latent space vectors are then drawn from $z \sim \mathcal{N}(0, \mathbf{1})$. The loss function of a cVAE is given by

$$\mathcal{E}_{\text{cVAE}}(\{\mathbf{f}, \mathbf{I}\}; \phi, \theta) = \mathbb{E}_{\mathbf{I}\sim\rho(\mathbf{I})}\Big[\beta \mathbb{E}_{z\sim q_\phi(\cdot|\mathbf{f})}\left[\log p_\theta(\mathbf{f}|z, \mathbf{I})\right] - \mathbb{E}_{\mathbf{f}\sim\rho(\cdot|\mathbf{I})}\left[D_{\text{KL}}\left(q_\phi(z|\mathbf{f})||\mathcal{N}(0, \mathbf{1})\right)\right]\Big]. \tag{4.10}$$

Broadly speaking, the idea behind Eq. (4.10) is to maximise (a lower bound to) the probability of the distribution controlled by the decoder $p_\theta(\hat{\mathbf{f}}|\mathbf{I})$ of matching the actual, but unknown distribution $\rho(\mathbf{f}|\mathbf{I})$. The first term corresponds to the aforementioned reconstruction error in Fourier space, $\mathcal{E}_{\text{rec}}(\{\mathbf{f}, \mathbf{I}\}; \theta)$, with β again controlling its relative importance. In all experiments with standard cVAEs, I set $\beta = 1$. The second term in Eq. (4.10) enforces a particular distribution of the latent space variable $z \in Z$, its effect can thus be conceptualised as a latent-space regularisation. A sketch of the gradient flow can be found in the appendix in Fig. A.5**b**. VAE are typically used for data compression: the dimensionality of the latent space, for instance, in this study could be made smaller than 8×8. However, I am not interested in this feature of VAE: my encoder network is only tasked with ensuring that the latent space variables z adhere to a standard normal distribution. Finally, I should add that I tacitly assume here that normally distributed latent space vectors work well for the problem at hand. The optimal latent space distribution is, admittedly, an open question. More details regarding the architecture and the training workflow can be found in the appendix Secs. A.2.3 and A.2.4.

Descending the forward loss (cVAE with a forward loss)

As I show in the introductory regression examples, an important property of regression techniques is the ability to interpolate the function of interest in the presence of noise or measurement error on the output end. However, what I attempt in this chapter is an interpolation in the opposite direction: I need to find input values for given output

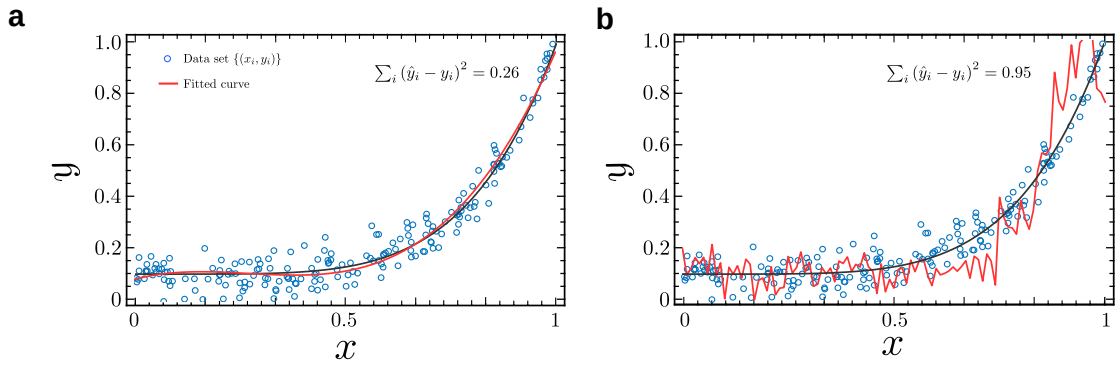

Fig. 4.4 Forward/backward polynomial interpolation. The function to interpolate is $s(x) = 0.9x^5 + 0.1 + \xi$ where $\xi \sim \mathcal{N}(0, \sigma = 0.1)$. A noiseless version of this function is plotted in black. The interpolation polynomial (red) is restricted to third degree. The total interpolation failure is printed in both panels. **a** Result when (x_i, y_i)-pairs are interpolated in y-space. **b** Result when (x_i, y_i)-pairs are interpolated in x-space, followed by a subsequent application of $s(x)$ (including noise). The sample size is 200 in both panels.

values using (effectively) a regression technique. In Fig. 4.4 I give an example of the kind of distortion that this inverted interpolation may cause. Importantly, the function $s(x)$ considered in Fig. 4.4 is invertible (see caption). The specific error that results from an inverted interpolation discussed here should not be confused with the fundamental failure of inverted interpolation for non-invertible functions shown in Figs. 4.1**b**, and 4.2**b**. In the example in Sec. 4.2.2, I investigate this issue further.

Attempting to solve an inverse problem by seeking to reconstruct function inputs could therefore be a problematic strategy: Descending the Fourier reconstruction loss only bends the models towards reproducing Fourier components given the intensity, which is the wrong error to minimise here. However, any attempt at directly minimising the forward error, $||\widehat{\mathbf{I}} - \mathbf{I}||_{l_2}^2$, requires a differentiable approximation $U : \mathbf{f} \rightarrow \mathbf{I}$. Such an approximation can be provided, of course, using neural-networks, particularly of the convolutional type. Indeed, training a forward model on this problem proved to be so simple, that I omit this result here. However, connecting the forward model U with the Fourier output of my generative models, $\hat{\mathbf{f}}$, turns out to be involved: Roughly speaking, the forward network needs to only relate actual Fourier inputs $\mathbf{f}$, not some incorrect matrix $\mathbf{f}'$, to the observed intensity $\mathbf{I}$. Otherwise it might nudge the generative model into producing Fourier matrices that satisfy the forward network, but not the actual forward operation, $S \circ F : \mathbf{f} \rightarrow \mathbf{I}$. This can be remedied, to some degree, by retaining the reconstruction-loss term in the Fourier plane, $||\hat{\mathbf{f}} - \mathbf{f}||_{l_2}^2$, however adding this loss term may result in trade-offs with the forward loss.

A strategy, which appears to work (see Sec. 4.3.1) looks as follows: The forward network $U_\xi : f_{i,j} \to I_{k,l}$ with parameters ξ is pre-trained on the data set using a forward loss,

$$\mathscr{E}_{\text{forw}}(\{\mathbf{f},\mathbf{I}\};\xi) = \mathbb{E}_{\mathbf{I}\sim\rho(\mathbf{I})}\left[\mathbb{E}_{\mathbf{f}\sim\rho(\cdot|\mathbf{I})}||U_\xi(\hat{\mathbf{f}})-\mathbf{I}||_{l_2}^2\right]. \tag{4.11}$$

Once pretraining is completed, the forward network is used to evaluate the Fourier matrix candidates $\hat{\mathbf{f}}$ proposed by the generative models in terms of their likely intensity pattern $\tilde{\mathbf{I}} = U_\xi(\hat{\mathbf{f}})$ (note the difference to $\hat{\mathbf{I}}$, which is the actual setup output for $\hat{\mathbf{f}}$). During the actual training, I continue to minimise Eq. (4.11) with respect to ξ in a separate training step as shown in Fig. A.5**c**. Importantly, the gradient descent steps of the VAE are taken only with respect to the encoder and decoder parameters ϕ, θ, not ξ.

In principle, one can train any conditional generative model in this way. I here choose a cVAE as an example. This forward-cVAE is updated using the following loss function

$$\mathscr{E}_{\text{cVAE + forw.loss}}(\{\mathbf{f},\mathbf{I}\};\xi,\phi,\theta) = \alpha\mathscr{E}_{\text{forw}}(\{\mathbf{f},\mathbf{I}\};\xi) + \mathscr{E}_{\text{cVAE}}(\{\mathbf{f},\mathbf{I}\};\phi,\theta). \tag{4.12}$$

where the new hyperparameter α controls the relative strength of the forward loss. After a minimal comparison of results of different combinations of α and β, I set $\alpha = 1$ and $\beta = 0.01$. The forward error is thus significantly more important than the reconstruction error. Again, once training is completed, the encoder and the forward networks are obsolete; only the decoder is retained.

Example: Inverting a square function

In order to illustrate the difference of behaviour between cVAE and cGAN in inverse problems, I here study a simple model of a non-invertible function. Furthermore, I demonstrate the benefits of a forward-loss introduced in the previous section. The task is to invert the noisy-square function $s(x) = x^2 + \xi$ on $x \in [-1,1]$ where ξ is zero-mean Gaussian noise with a standard deviation of $\sigma_\xi = 0.05$.

In Fig. 4.5**a**, I sketch the variable-flows through a standard cVAE (top) and cVAE with forward loss (bottom). In Fig. 4.5**b**, I show the results for four simple conditional generative models, a cVAE and a cGAN without (top row) and with forward loss (bottom row). All four panels in Fig. 4.5**b** show the predicted $\hat{x}$ values over the respective original forward value y. The colour code differentiate the latent space values that feed into the decoder (left two panels in **b**) and generator (right to panels in **b**): red colours correspond to $z \in [-3,0]$ while blue-coloured points are the result of $z \in [0,3]$. As the figure shows, the cVAE approximately separates the latent space into disjoint sets where the corresponding $\hat{x}$-sets appear to touch only at the $x = 0$-axis. The cGAN on the other hand has an unstructured

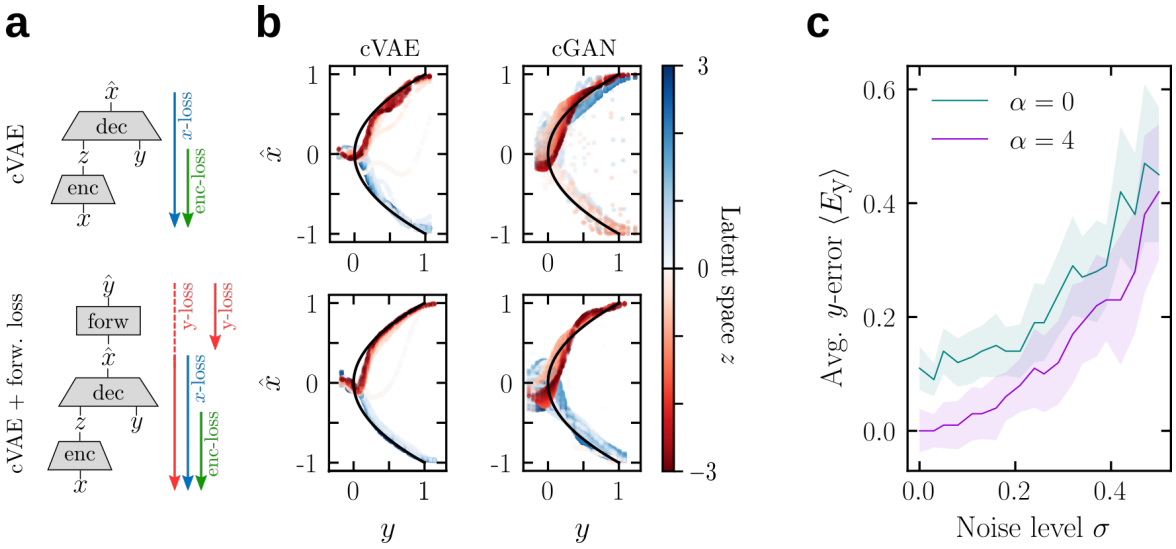

Fig. 4.5 Inverting a noisy square function $s(x) = x^2 + \xi$ using conditional generative models. **a** Sketch of the gradient flow during training for standard cVAE (top) and cVAE with forward loss (bottom). **b** Examples of cVAE (left column) and cGAN (right column) trained without (top row) and with (bottom row) forward loss for a noise level of $\sigma = 0.05$. **c** Average y-reconstruction error of a cVAE trained with (violet, $\alpha = 4$) and without (teal, $\alpha = 0$) forward loss as a function of the level of noise. The shaded areas indicate the respective standard error of the mean.

latent space where $\hat{x}$-values appear to be grouped, but the two sets are overlapping (the blue branch interfuses with the red branch). In an ideal cVAE one has thus some control over which branch is realised. In reality, the sign of the branch would, of course, be hidden to the observer.

All models in this example have been trained over $n = 10^5$ steps using the ADAM optimiser [213] at a learning rate of $\eta = 3 \cdot 10^{-4}$. The generator in the cGAN and the decoder in the cVAE model each consist of three layers with 100, 100, and 50 nodes respectively. In cGAN and cVAE models, the input layer has two input nodes to accept the latent space variable z and the conditioned-on variable y. The output layer returns a single value, the prediction, $\hat{x}$. The output of each but the final layer is subject to a rectified-linear (ReLu) function. The discriminator and encoder share the same simple architecture with the difference that they accept the true function input x instead of the latent space variable z.

The latent space variables in the cVAE model are drawn from a standard normal distribution $z_{\text{VAE}} \sim \mathcal{N}(0,1)$ after training. In the case of the cGAN, latent space variables are drawn from a standard normal distribution $z_{\text{GAN}} \sim \mathcal{N}(0,1)$ both during training and afterwards.

In Fig. 4.5**c**, I plot the y-error that a forward operation using the actual function s on a proposed input $\hat{x}$ achieves. This y-error constitutes the analogue of the relevant metric in my intensity-reconstruction experiments. The cVAE trained with a forward loss achieves a lower forward error as its standard counterpart, especially for low noise variances. For noise standard deviations approaching $\sigma = 0.5$, the two square-branches begin to overlap significantly which results in a deteriorating performance across all models.

Direct parametrisations of the conditional density

A complementary neural-network-based approach to inverse problems is to directly parametrise a model of the probability distribution $\rho(x|y)$ and explicitly allow for multimodality, such as a sum of Gaussians with learnable relative amplitudes, means, and variances. This is the idea behind mixture-density networks (MDN) [214, 215]. In the original implementation of MDN, the network outputs $(L+2)K$ parameters $\{\mu_1^{(l)},\ldots,\mu_K^{(l)}\}$, $\{\sigma_1,\ldots,\sigma_K\}$, and $\{\pi_1,\ldots,\pi_K\}$ for a L-dimensional problem and K Gaussian modes. Since the network accepts y as input, all parameters are functions of y. The model for the density assumes output dimensions to be uncorrelated and hence reads

$$\rho(x|y) = \sum_{k=1}^{K} \pi_k(y) e^{-\frac{1}{2}\sum_{l=1}^{L} \frac{(x^{(l)}-\mu_k^{(l)}(y))^2}{\sigma_k(y)^2}}. \tag{4.13}$$

The details of MDNs are described in Christopher Bishops book [186]. Recently, MDNs have been generalised to models of continuous mixtures of Gaussians, called compound density networks [216].

A parametrisation has the advantage that the model is able to quantify the uncertainty of its prediction of a given input. However, at least in the case of discrete-mixture-density networks, this approach may fail to invert high-dimensional mappings, where a continuous range or manifold of inputs is mapped to the same output. Moreover, the parametrisation of the density often requires a larger output dimension of the neural network, as shown above, which may be prohibitive. And, finally, even if the network-parametrised distribution converges, the problem of sampling from it remains. Admittedly, in the case of Gaussian models, sampling is easily accomplished by sampling from standard normal distributions $\vec{\varepsilon} \sim \mathcal{N}(\vec{0},\mathbf{1})$. The obtained random numbers can then be shifted to the desired distribution $\mathcal{N}(\vec{\mu},\mathbf{S})$, i.e. $\vec{z} = \sqrt{\mathbf{S}}\vec{\varepsilon} + \vec{\mu}$ where the square-root denotes the matrix-square-root.

In the case of more involved distributions, generative models have the clear advantage that they directly generate samples. The probability distribution they learn is encoded implicitly in their network weights.

I did attempt to apply mixture density networks to the problem at hand, but the model failed to converge on 8×8-Fourier matrices (explained in Sec. 4.2.1). It only visibly converged for a much reduced problem of maximum dimensionality of 4×4-Fourier components. Furthermore, I noticed that in higher dimensions it becomes increasingly difficult to avoid floating point underflows in mixture-density networks.

An expert system for Fourier-matrix inference

Finally, I should add that I also attempted naive approaches to finding the Fourier matrix. The first idea was to find a convolution matrix $\mathbf{g}$ that transports the 100x100 intensity image into the $8x8$-Fourier plane. To this end, I minimised the $||\mathbf{f}_{i,j} - (\mathbf{g} \circledast \mathbf{I})_{i,j}||_{l_2}$-norm of the convolved intensities (second term) and the Fourier matrix (first term) to estimate the elements of the convolution matrix. This approach, however, leads to non-sparse Fourier matrices, which, when transferred into holograms, fail to result in any measurable output intensity, due to overloaded holograms.

I conclude that the problem requires a thresholding non-linearity: Using a peak-finding technique, I inferred affine transformations, $i = a_y I_y^{peak} + b_y$ and $j = a_x I_x^{peak} + b_x$, which transform peak-positions in the intensity plane into Fourier-matrix indices i, j. The elements at these indices can now be set to some value, e.g. $f_{i,j} = c \cdot I_{y,x}^{peak}$, where c is a linear scaling factor (I used $c = 1$). The peak-finder requires a threshold to binarise the image, which must be seen as a model parameter. In all experiments, I used thr=0.5·max($\mathbf{I}$). The peak position is then inferred using a centroid from the non-binarised image within a region of interest extracted from the binarised image, similar to the colloid-localization routine described in Sec. 2.2.4. Of course, one cannot hope that the shape of the peak gets replicated with this technique. Since the Gerchberg-Saxton or other algorithms cannot be directly applied to the problem at hand, this naive approach serves as an "expert-system" or baseline against which I can compare the generative models.

4.3 Results

4.3.1 Validation data set results

In this study, I compare (1) spectrally normalised conditional GAN with reconstruction loss ($\beta = 1$), (2) without reconstruction loss ($\beta = 0$), and (3) conditional Wasserstein GAN

with gradient penalty without reconstruction loss ($\beta = 0$) with (4) standard conditional VAE, and (5) conditional VAE with forward loss, and, finally, the (6) the *naive expert*-system described in Sec. 4.2.2.

The architecture used for the spectrally normalised cGAN models can be found in the appendix in Figs. A.2 and A.3. The architecture of the cWGAN, but also of the cVAE are quite similar. As the figures show, every convolutional and dense layer is subject to either batch normalisation [217] (cGAN, cVAE) or, in the case of WGAN, layer normalisation [218]. Furthermore, to avoid overfitting, I generally applied dropout to dense layers with rates between $r = 0.2$ and $r = 0.3$ [219].

Hyper parameters, such as the gradient-descent optimiser used, or the number of epochs trained are summarised for all models in Tbl. A.1.

I trained all models on my "040319"-data set, which comprises 99801 samples and was acquired at $P = 1$ W laser power with a camera exposure time of 0.001s, a gain of 8x, and with a Fourier-sparseness parameter of $C = 97$. The camera settings are carefully chosen to ensure that strong, single-frequency spots do not reach the maximum dynamic range of the camera. At the same time, the weaker intensity distributions created by multiple frequencies need to be visible as well.

Before recording the intensity reconstruction error in Fig. 4.6, I briefly realigned the setup to reproduce the conditions of the training data set as well as possible. However, a significant error $\mathscr{E}_{\text{forw}}^{\text{misalign}} \sim 500 - 1000$ depending on overall intensity remains, which all results are subject to. In addition, before the experiment, I wait 30 min in order for the laser to stabilise and ensure that all optical elements involved reach stable equilibria.

The error that I am interested in, is the minimal forward error over five latent space redraws $z_0, z_1, \ldots, z_4$

$$E_{\text{forw}} = \min_{z_0,\ldots,z_4} ||\mathbf{I} - \hat{\mathbf{I}}(z_i)||_{l_2}. \tag{4.14}$$

In Eq. (4.14), $\mathbf{I}$ denotes the target intensity $\mathbf{I}$, while $\hat{\mathbf{I}}(z)$ represents the intensity corresponding to the Fourier-estimate $\hat{\mathbf{f}}(z)$, i.e. $\hat{\mathbf{I}} = S \circ F(\hat{\mathbf{f}}(z))$ (see Fig. 4.3**c**). I will therefore compare the performance of all models with respect to the forward error.

In Fig. 4.6, I plot the error in Eq. (4.14). The validation data set on which the models are tested here is separate from the training data set, but is created following the same workflow as sketched in Fig. 4.3**b**. The average error for each model is shown in th figure as well: The cVAE with forward loss (panel **a**) achieves the lowest forward error with $\langle \mathscr{E}_{\text{forw}}^{\text{min}} \rangle = 798$. The expert system, for comparison, results in an average forward error of $\langle \mathscr{E}_{\text{forw}}^{\text{min}} \rangle = 1029$, worse than both versions of cVAE.

Although the forward error $\mathscr{E}_{\text{forw}}$ provides a metric by which I can compare different models, visual inspection is still useful to assess the quality of intensity reconstructions. This is especially true for this project, since the precise position of laser spots is not as important as their shape.

For the sake of avoiding any selection bias, I plot the first 20 samples of the validation data set in Figs. 4.7-4.11 (with the respective minimum forward-error). The brightness of the intensity distributions in these figures is slightly increased for better visibility.

Especially the cVAE-models manage to recreate not just position and intensity of the intensity rather well, but in addition, in some cases, reproduces the shape of the laser spots. An inspection of Figs. 4.7 and 4.8 shows that the shapes reproduced by the standard cVAE appear more similar to the input patterns than those reproduced its forward-trained counterpart (compare e.g. 6, 10, 15, or 17 in Figs. 4.7 and 4.8).

Moreover, the Fourier matrix candidates proposed by the cVAE with forward loss in Fig. 4.7 appear remarkably different from the originals. The standard cVAE, on the other hand, essentially recreates the original Fourier matrices (see Fig. 4.8). This underlines the difficulty of the problem: very different Fourier matrices may result in quite similar output patterns.

All generative models retained some, albeit in some cases low variability in their predictions as can be seen from the spread of grey points in Fig. 4.6. In the forward-direction, it is difficult to asses whether the models have actually learned the true distribution $\rho(\mathbf{f}|\mathbf{I})$, since all proposed Fourier candidates $\hat{\mathbf{f}}$ should, in theory, lead to similar intensities $\hat{\mathbf{I}}$. However, the impression of lower variance in both cVAE models and the cGAN with reconstruction loss is correct. Indeed, especially the cVAE with forward loss yielded almost always the same Fourier prediction $\hat{\mathbf{f}}$ independent of the latent space vector z drawn. The model might suffer from a conditional mode collapse, which would see it only output a single prediction $\mu_{\hat{\mathbf{f}}}(\mathbf{I})$ given an input intensity $\mathbf{I}$. There are several possible reasons for this: As mentioned before, mode collapse is a well-described problem in the machine learning literature. The fact that it occurs here should therefore not be too surprising. It is also possible that I sufficiently reduced the complexity of the problem as described in Sec. 4.2.1 and successfully eliminated most of the input-invariance in the data. This, however, is unclear since I studied the phase problem in Sec. 2.2.5 only for isolated Fourier components.

Furthermore, models without l_2-loss terms, such as the cWGAN and the cGAN, however, do exhibit significant conditional variability. The collapse could thus simply originate from the unrelenting constraints of the l_2-loss.

In any case, a conditional-mode-collapsed model is still useful for productive purposes as long as one wants to reproduce intensity patterns - the original aim of this project. All models except the cGAN manage to produce an almost black (zero) intensity when asked to in example 9. I find this noteworthy, because many of the models and architectures I tested until this point failed to do this reliably.

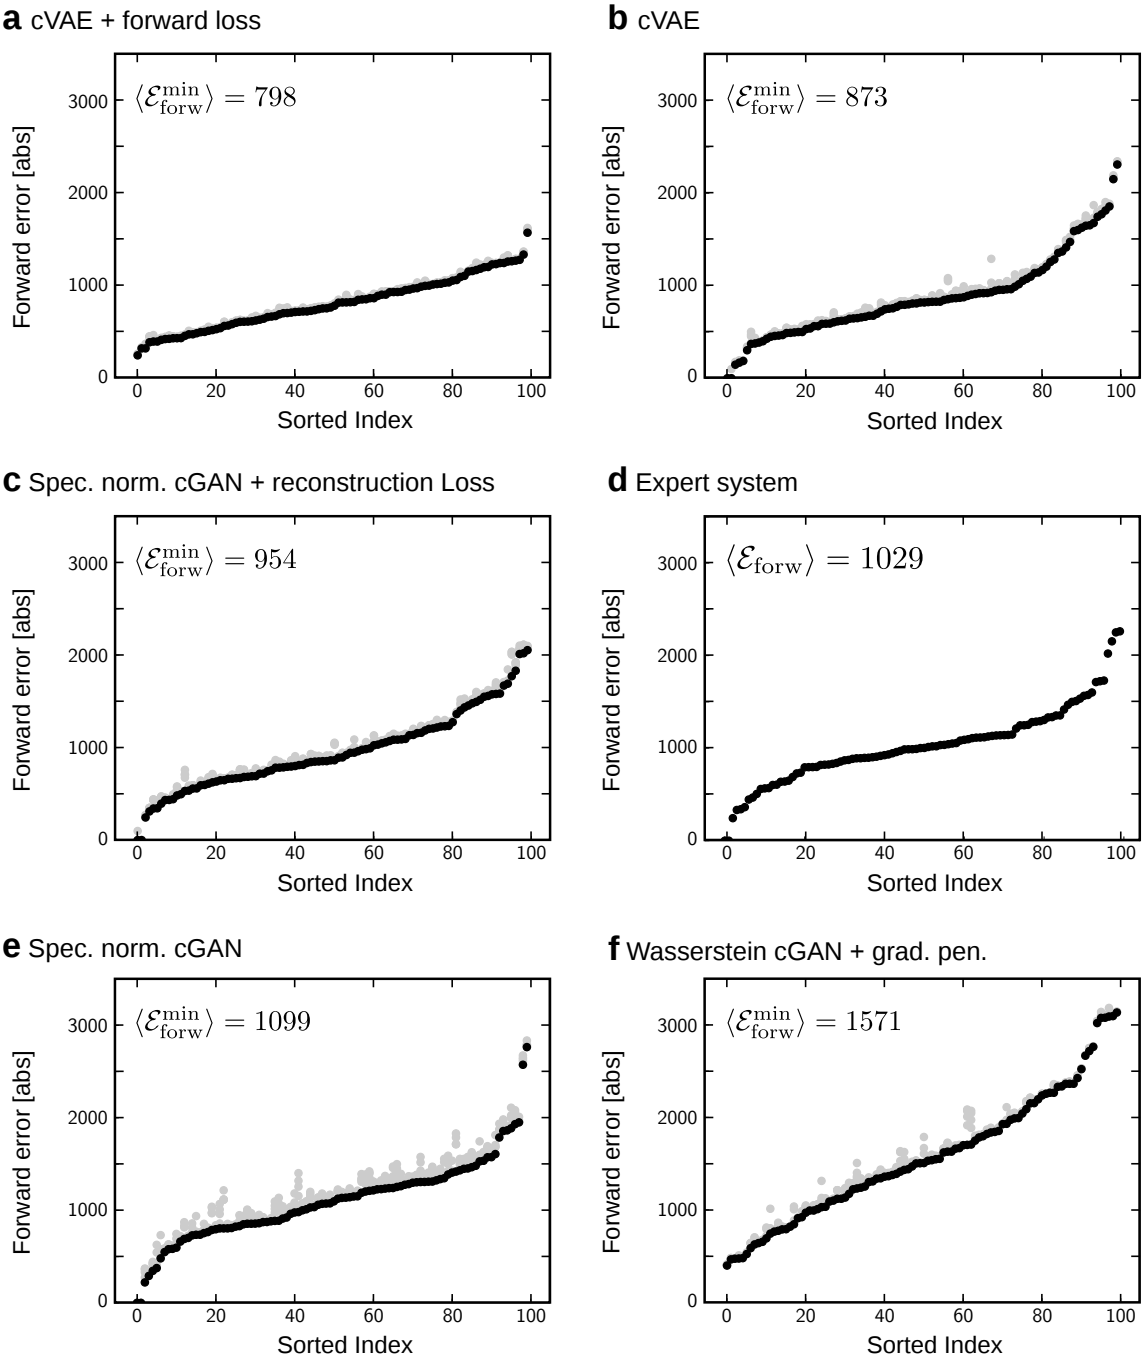

Fig. 4.6 Overview of the the minimum forward error Eq. (4.14) achieved by the different models (black) on the validation data set. The other 4 draws from the model are plotted in grey to give a sense of the spread in intensity error. The ordinate is sorted according to minimum error. The panels are sorted according to the average minimal forward error.

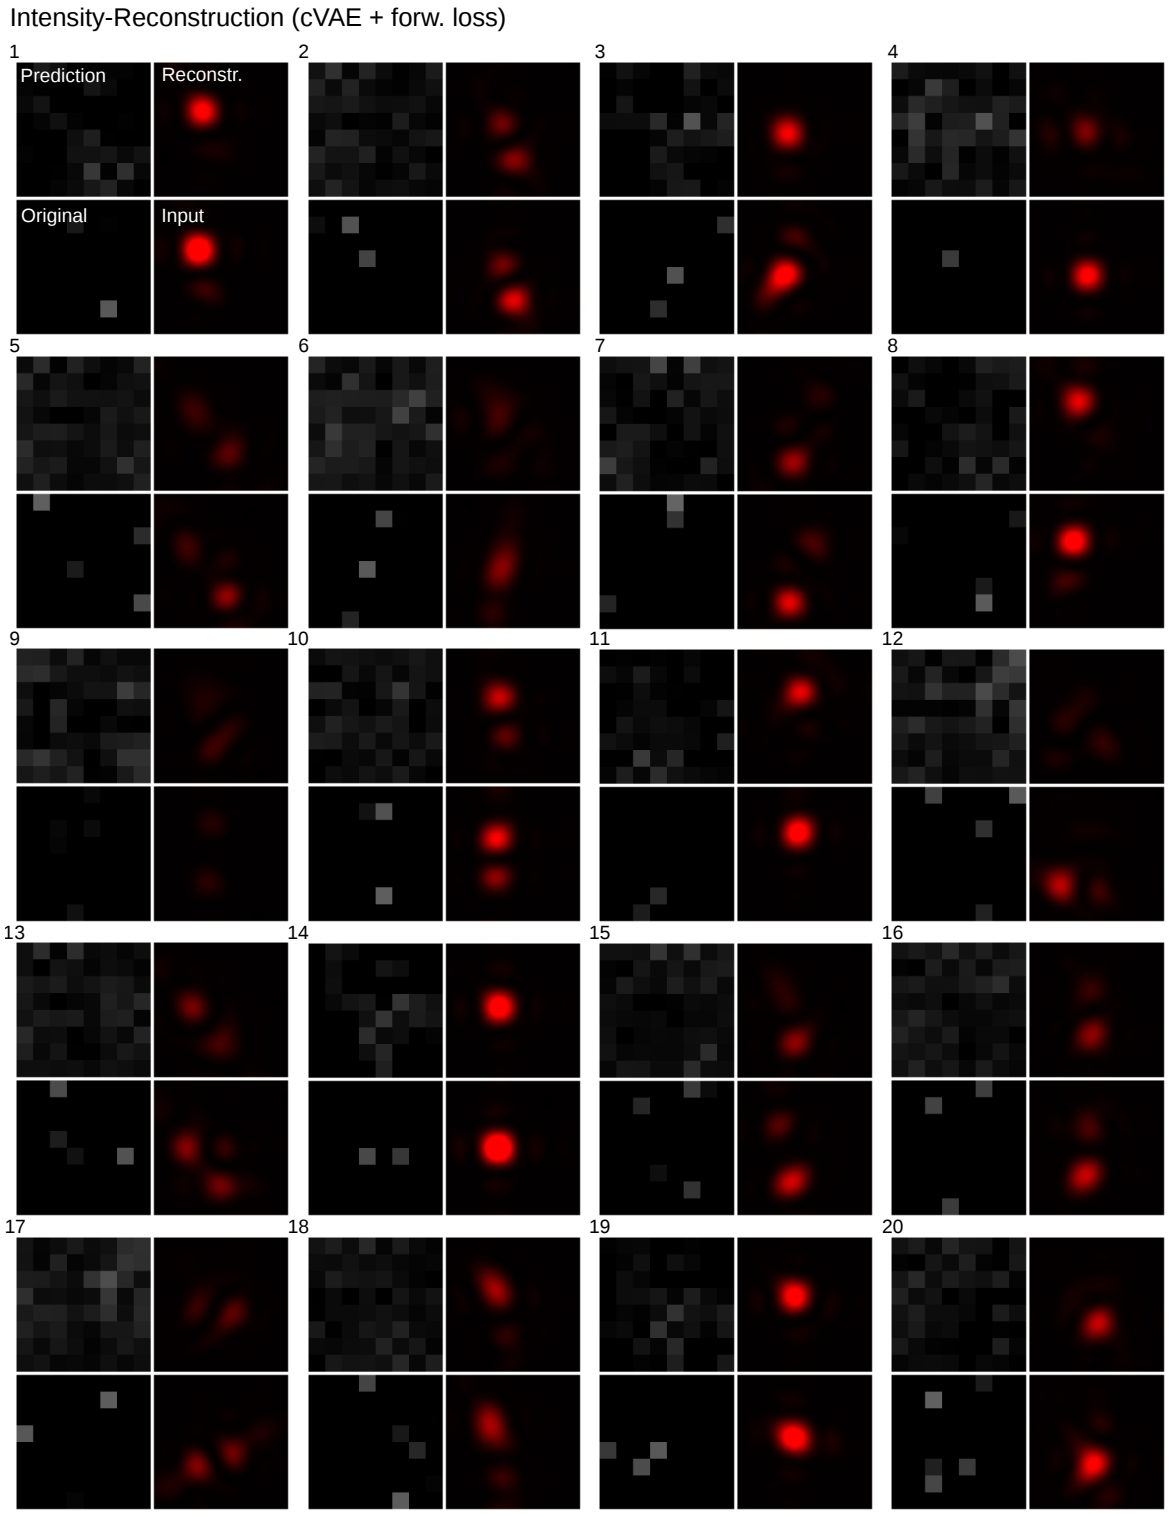

Fig. 4.7 First 20 cVAE + forw. loss minimum forward-error (see Eq. (4.14)) samples of the validation data set.

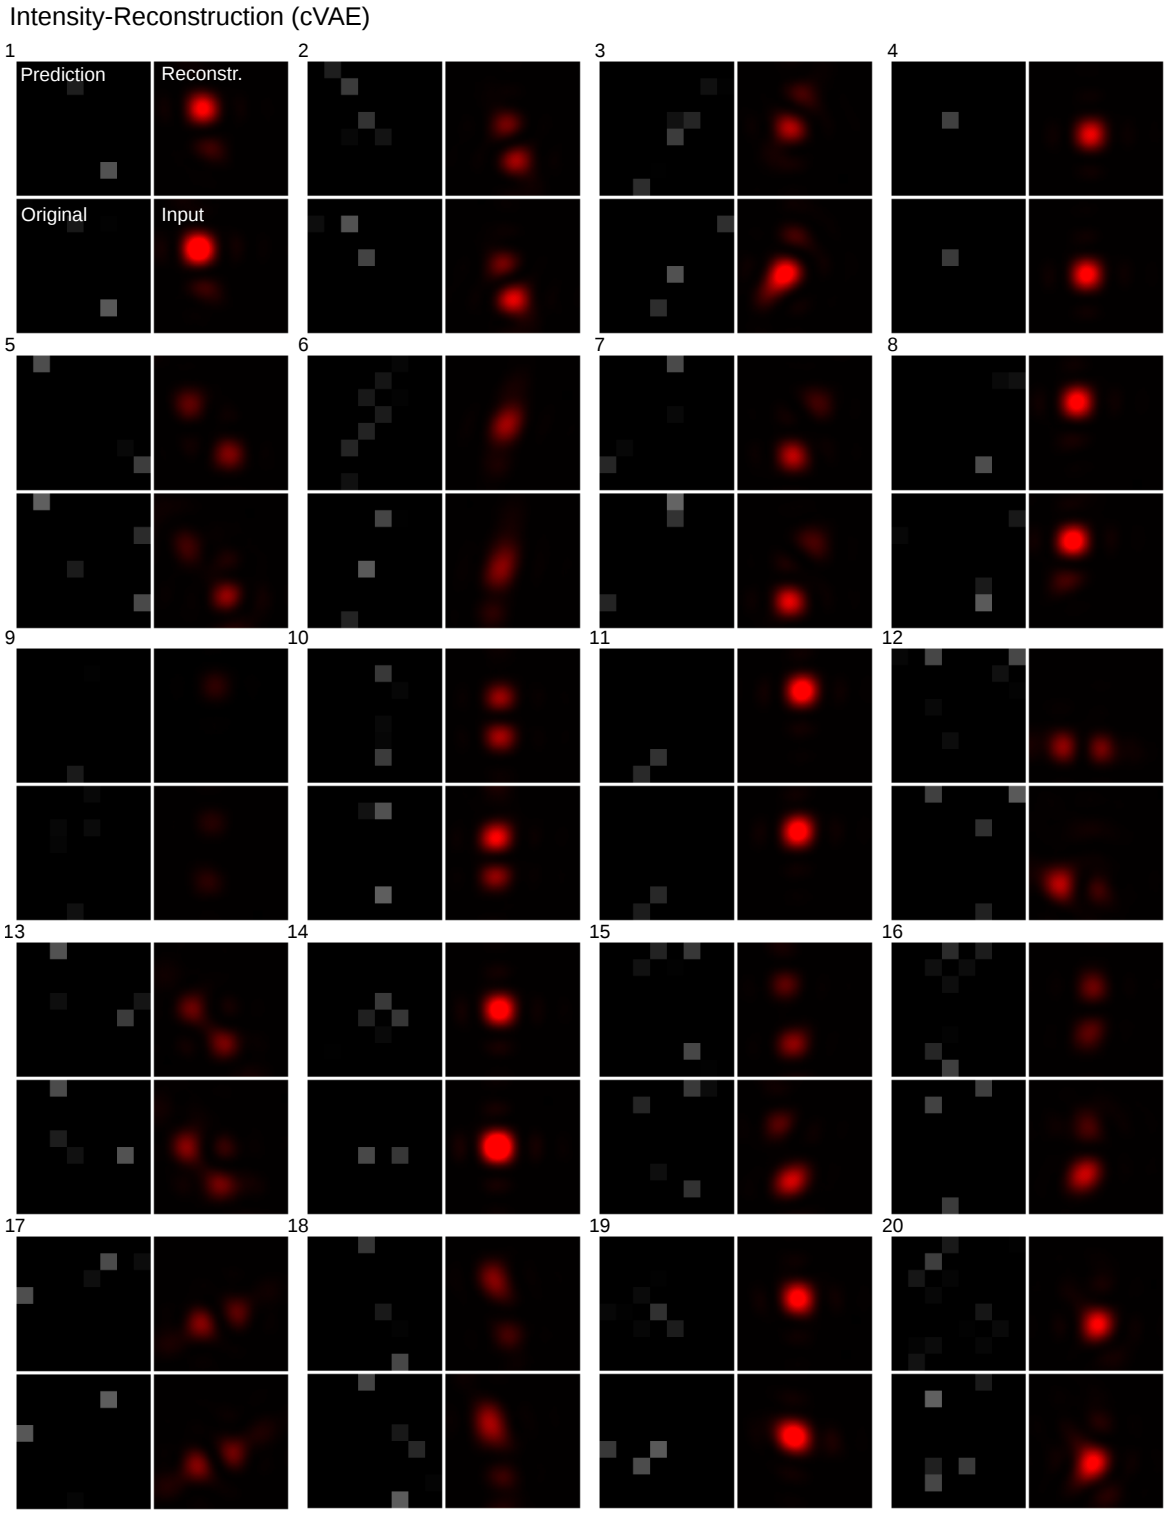

Fig. 4.8 First 20 cVAE minimum forward-error Eq. (4.14) samples of the validation data set.

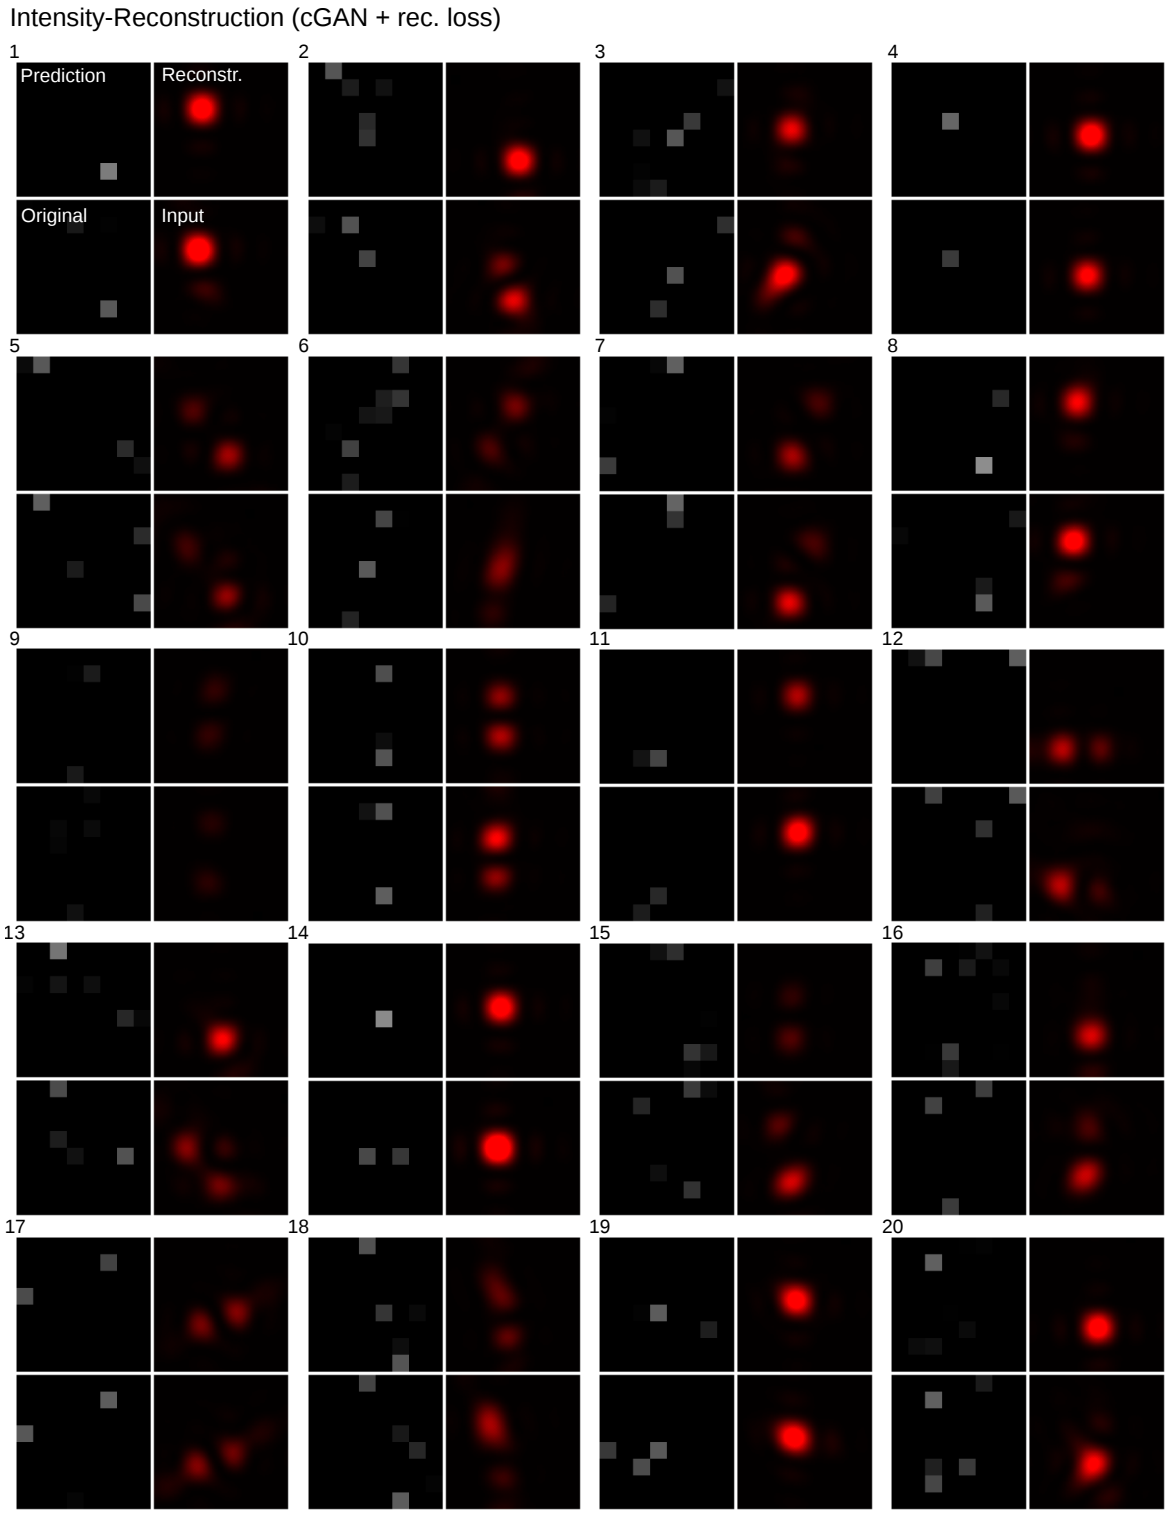

Fig. 4.9 First 20 cGAN + rec. loss minimum forward-error Eq. (4.14) samples of the validation data set.

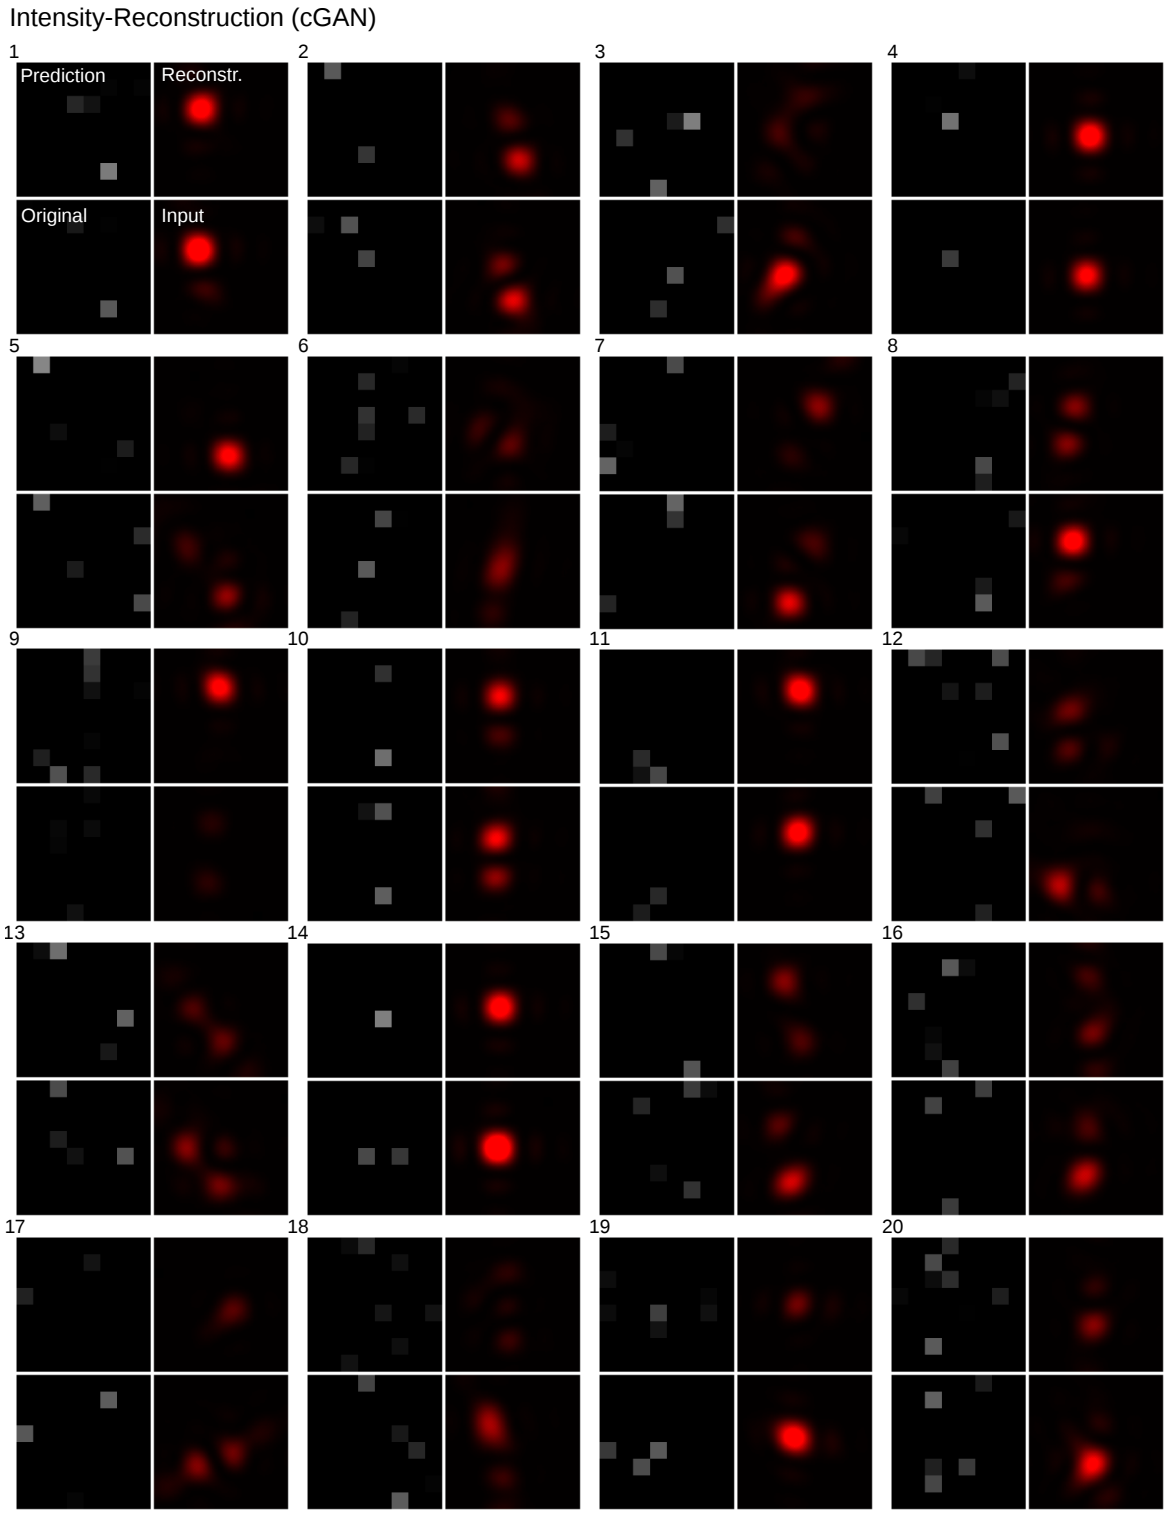

Fig. 4.10 First 20 cGAN minimum forward-error Eq. (4.14) samples of the validation data set.

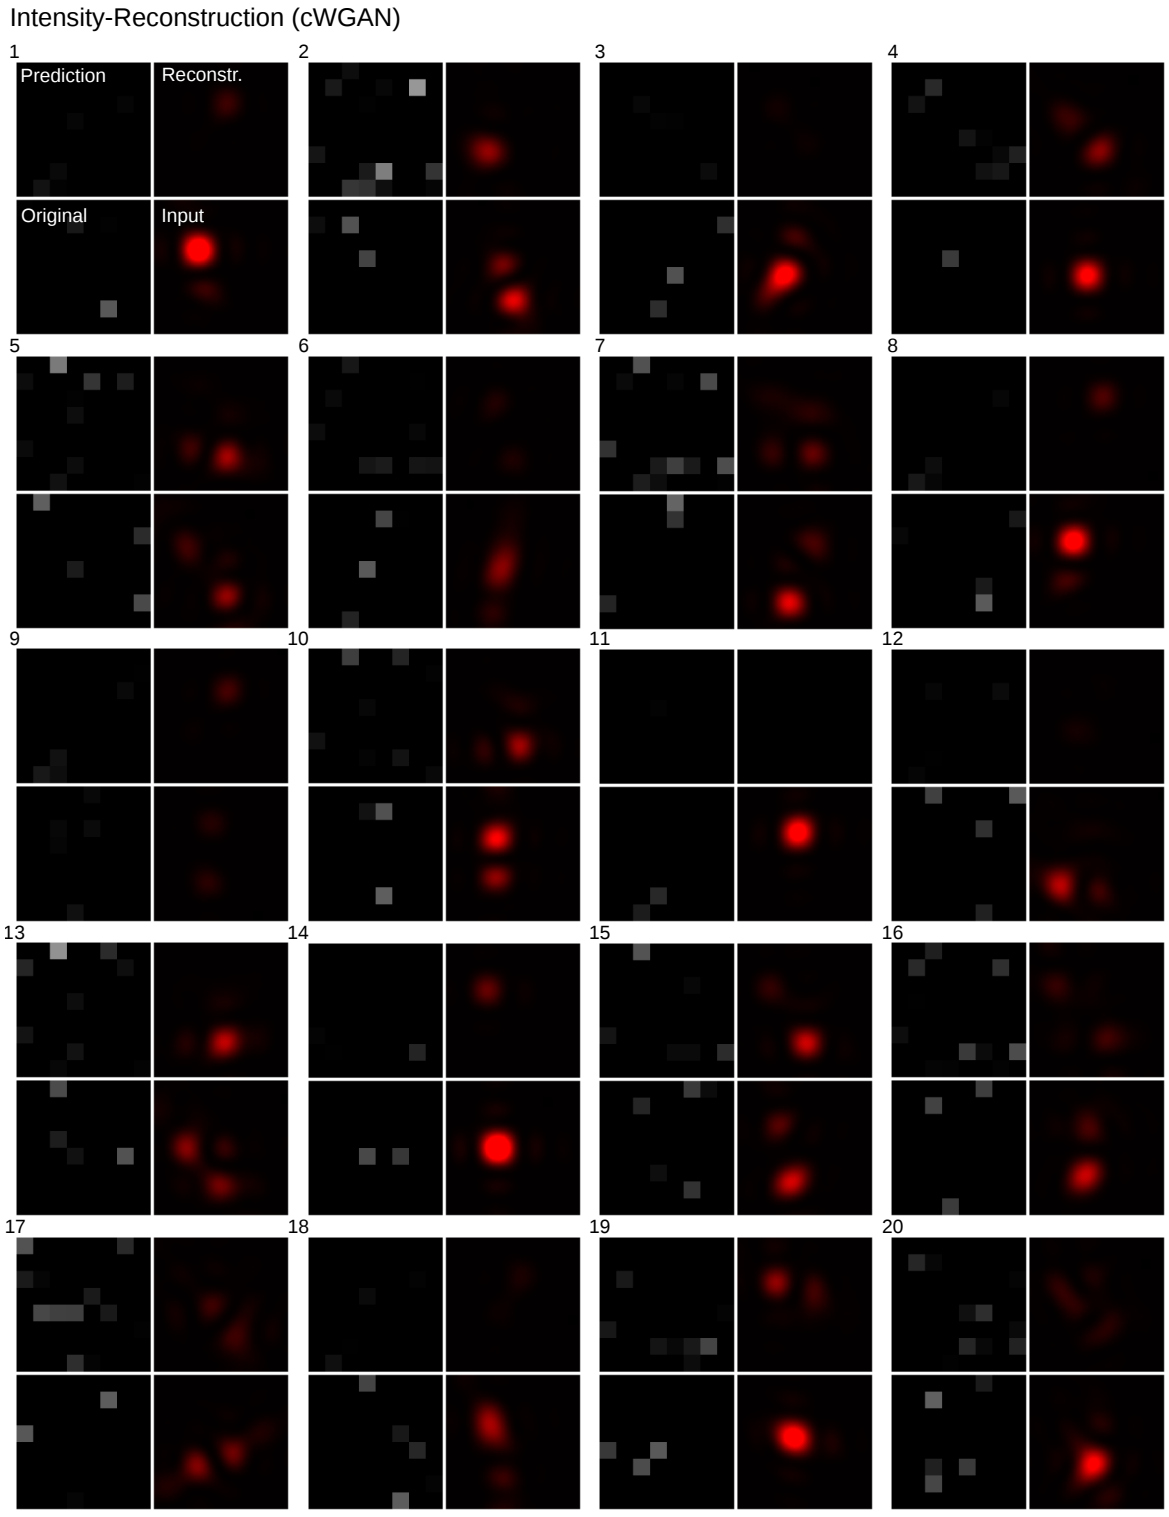

Fig. 4.11 First 20 cWGAN minimum forward-error Eq. (4.14) samples of the validation data set.

4.3.2 Test set results

In many applications of digital holography, the target intensity distribution will vary across a wide range of possible patterns. Importantly, from an applied perspective, it is often not *a priori* clear whether a hologram that results precisely in the target intensity even exists. A desirable property of a generative model for inverse digital holography is therefore to propose Fourier matrices or holograms, that result in an intensity that is as close as possible to the target. However, with the exception of the cVAE with forward loss, all models discussed here optimise for similarity in Fourier space, not intensity space. Moreover, the models are trained on holography data sets which may not include the patterns necessary to create the desired intensity.

In the following, I test the generative models on entirely synthetic intensity distributions **I**. I create these distributions using a Gaussian-mixture model with a variable number of peaks $0 < k \leqslant N_p$, peak-positions $(\mu_x^{(k)}, \mu_y^{(k)})$, amplitudes A_k, peak variances $(\mathrm{var}_x^{(k)}, \mathrm{var}_y^{(k)})$, and x-y-covariances $\mathrm{cov}_{x,y}^{(k)}$,

$$\rho_{\mathrm{synth}}(x,y) = \sum_{k=0}^{N_p} A^{(k)} e^{-\frac{1}{2}\frac{\mathrm{var}_y^{(k)}(x-\mu_x^{(k)})^2 - 2\mathrm{cov}_{x,y}^{(k)}(x-\mu_x^{(k)})(y-\mu_y^{(k)}) + \mathrm{var}_x^{(k)}(y-\mu_y^{(k)})^2}{\mathrm{var}_x^{(k)}\mathrm{var}_y^{(k)} - \mathrm{cov}_{x,y}^{(k)}\mathrm{cov}_{x,y}^{(k)}}}. \tag{4.15}$$

In a loose attempt at creating realistic intensity patterns which could have been part of the training data set, I required peak positions μ_x, μ_y in Eq. (4.15) to respect a margin to the image boundaries of 24 pixels (left, top) and 30 pixels (right, bottom). The variances are restrained by $50 \leqslant \mathrm{var} \leqslant 65$. Furthermore, the maximum achievable intensity for each test set sample is restricted to $I_{\max} = 200/\sqrt{N_p}$, where N_p is a discrete uniformly-distributed random number drawn initially $N_p \in \{1,2,3\}$. Despite these restrictions, the intensities might still be unrealistic: in reality, just to give one example, spot amplitudes and spot-shape variances are not independent parameters.

The resultant minimum forward error achieved by the generative models is shown in Fig. 4.12. I show ten examples for each model of the best three models in Fig. 4.13.

Surprisingly, the standard cVAE fares worse on synthetic data (Fig. 4.12**e**). Perhaps less surprising is the fact that the cVAE with forward loss still performs well and, again, achieves the lowest forward error (Fig. 4.12**a**). The examples shown in Fig. 4.13**b** underline its strong performance. It generally attempts to produce an approximative intensity distribution, as the double-peak example 7 in Fig. 4.13**b** shows.

The data-set-agnostic expert system, too, performs well, but slightly worse than the best trained model (panel **b**). It fails, for instance, to realise that there are two close-by spots in example 7.

The cGAN with reconstruction loss (panel **c**) exhibits the third best performance with $\langle \mathscr{E}_{\text{forw}}^{\text{min}} \rangle = 1163$. The cGAN version without reconstruction loss (not shown) performs just slightly worse with $\langle \mathscr{E}_{\text{forw}}^{\text{min}} \rangle = 1170$.

4.4 Conclusion

In this chapter, I introduce a, to my knowledge, novel approach to inverse holography using conditional generative models. Importantly, my approach takes optical aberrations or alignment artefacts into account and directly relates desired intensity distributions to Fourier tables, which can then be used to compute holograms.

The approach is geared towards applications in holographic optical tweezers where required intensity landscapes may only span a few diffraction-limited-spot sizes. However, as my experiments in chapter 3 show, it is is necessary to construct such intensity landscapes with high precision. High precision on small patterns is often not achieved by phase-retrival algorithms. The standard Gerchberg-Saxton, for instance, assumes a direct correspondence between Fourier-components and resultant intensities and disregard point-spreading effects, non-linearity of phase-SLM pixel value relations, interference, or other optical aberrations.

My results paint a clear picture, in terms of which generative model is best suited for the purpose. The cVAE model with forward loss proved excellent at all tasks. The standard cVAE, by contrast, perfomed well when asked to reconstruct intensities that were created by a given workflow and Fourier-component distribution, it failed to generalise well to synthetic data. Since it only learns the structure of the training data set, it tends to react to synthetic data with zero output. GAN, on the other hand, have a more fuzzy, semi-supervised training objective: in GAN with reconstruction loss, the generator at the same time needs to fool the discriminator and satisfy the reconstruction error. In earlier stages of training, the discriminator itself will not have a precise idea about which Fourier components should result in which spot position. Under these circumstances, approximately correct outputs by the generator will likely satisfy the discriminator. Admittedly, mathematical properties of GAN, let alone conditional GAN, are subject of ongoing research and more detailed explanations might be found in the future.

What becomes clear in this chapter is the fact that in order to deploy (non-Bayesian) deep learning models one may have to conduct a comprehensive model and hyperparameter search.

The architectures and hyperparameters identified in this study are the result of a number of educated trial-and-error over several generations of data sets with different

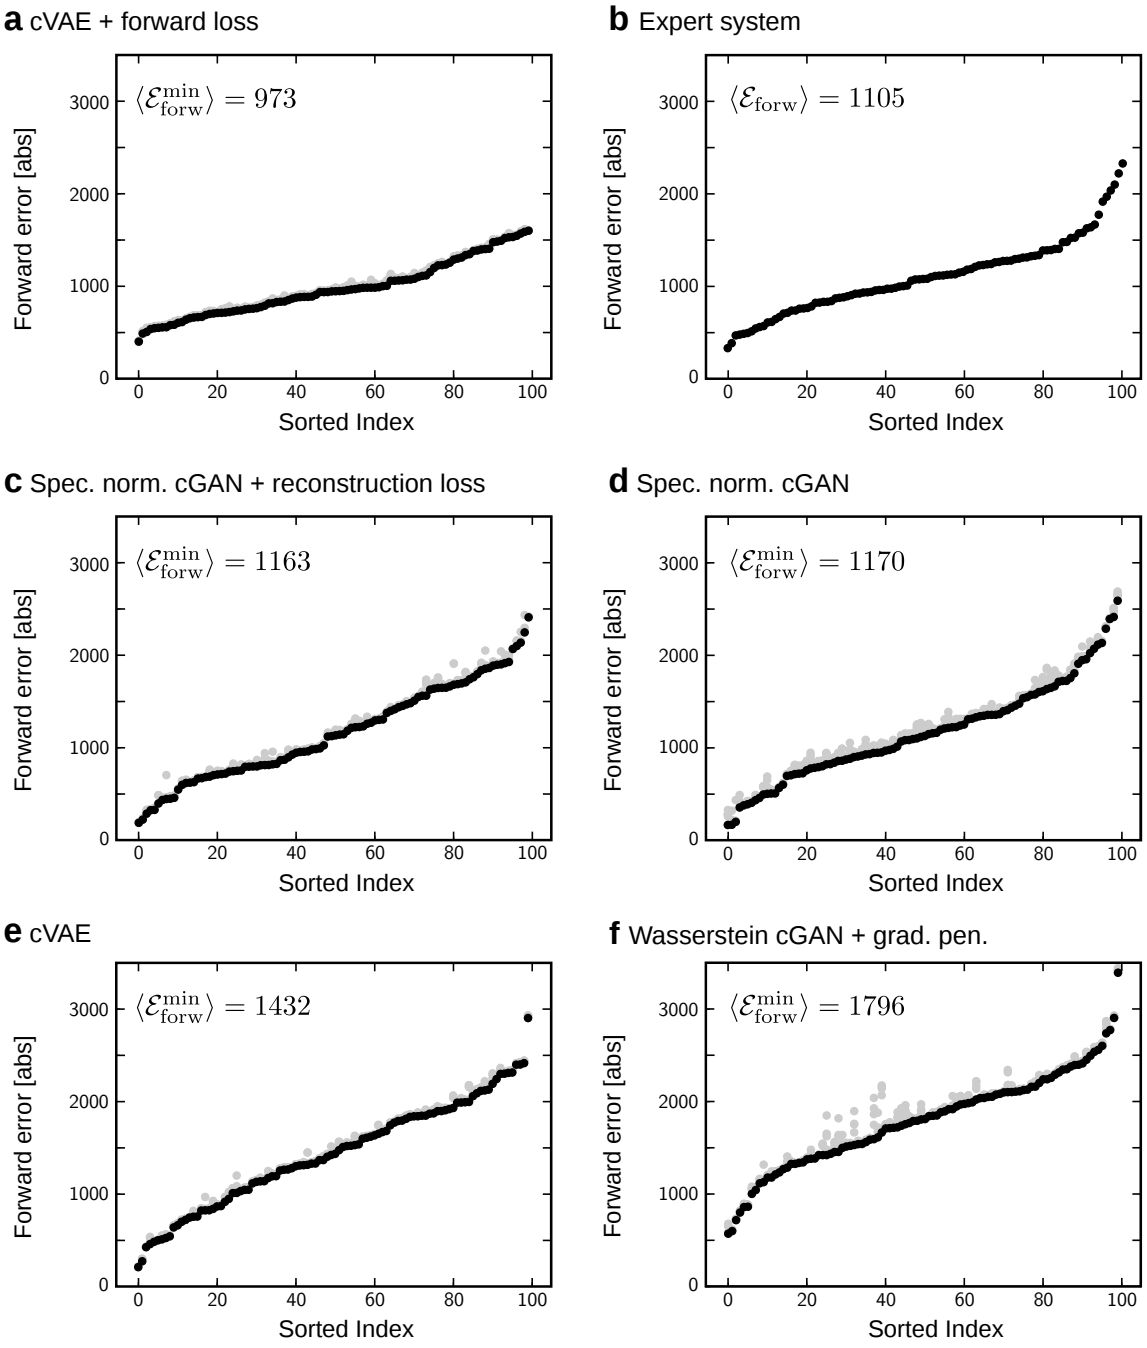

Fig. 4.12 Similar to Fig. 4.6 but for a synthetic intensity data set, created by a mixture of Gaussians, see Eq. (4.15). The forward error measured here is defined in Eq. (4.14).

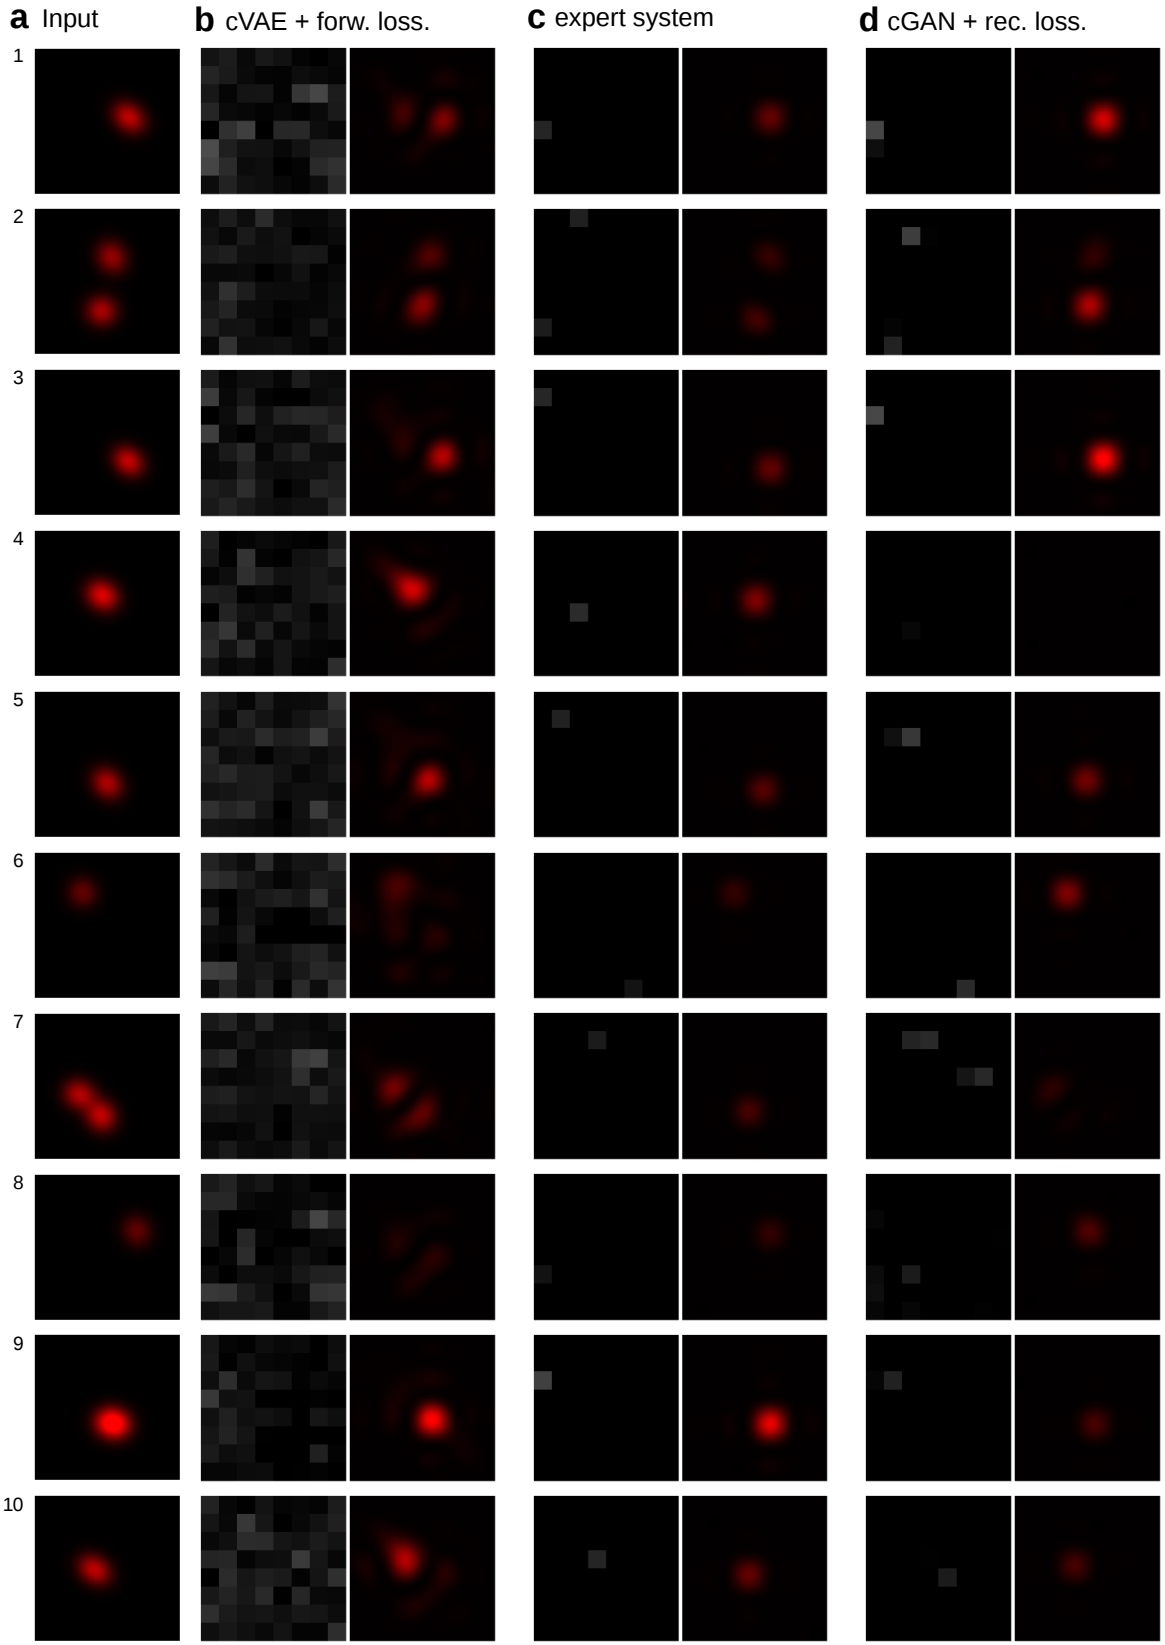

Fig. 4.13 **a** First ten examples of input synthetic intensity distributions. **b-d** Corresponding minimum-forward-error samples of intensity reconstruction.

workflows and camera settings. The setting that I converged to provides the required dynamic range of the camera for the intensity patterns while retaining a sufficient degree of permissible hologram complexity such that the generative models could be used for tweezing purposes.

Irrespective of the type of approximation method, the focus in this chapter firmly lies on (re)producing output intensity patterns. Indeed, by fixing all Fourier components to the real plane, I fix all relative phases between the Fourier components. Nevertheless, the method presented here should still be classified as phase-retrieval since the SLM is still a phase-only device.

From a thermodynamics-perspective, one wants to be able to realise a given force landscape; the nature of the force is often secondary. The idea here is to create models that are powerful enough to create force landscapes solely by shaping intensities.

While the restriction to an 8×8-Fourier matrix is a severe limitation, it also makes the model more portable: Should, for instance, the angle of the SLM change (see Fig. 4.3**a**), the model would not need to be retrained. One would rather embed the Fourier matrix at a different position in Fourier space.

The choice I made in Sec. 4.2.1 to consider Fourier gratings instead of blazed gratings (see Sec. 2.2.1) is mainly out of convenience. The technique discussed in this chapter should be applicable to blazed gratings as well. In my arXiv submission, I apply the technique discussed here to data sets created from blazed gratings [220].

Finally, I want to stress that my work in this chapter is merely a proof of concept. The models used were not too complex: For instance, the generator in Fig. A.2 "only" has on the order of 800k parameters, which pales in comparison with modern deep vision systems, such as the ResNet-class [221]. There is no fundamental reason, why this study could not be repeated on the direct and unrestricted mapping between 800×600-sized holograms to all 1280×1024-pixels on the camera. Such models, however, appear to be currently beyond the computational reach of standard deep-learning on desktop GPUs.

Chapter 5

Towards massively parallel fluorescence-based nanopore experiments

One of the aims of this thesis is to bridge the gap between the micron- and the nano-scale in experimental control and availability of data. This chapter describes an ongoing project which is thus, again, part of the outlook. In contrast to previous chapters, the challenges involved in this project are mainly of technical nature. The idea here is to recreate the setup described by our colleague F. Montel at ENS Lyon [222] and combine it with recently designed FRET-DNA origami platforms [223, 224]. Crucially, the DNA-origami study by Hemmig *et al* [224] was carried out on single glass nanopores and thus on single DNA-origami platforms. While this single-pore approach offers a high degree of control, it suffers from an inherently low signal-to-noise ratio, particularly in FRET measurements.

Even though I carried out exploratory measurements with the mentioned DNA-origami system, I here focus on simpler DNA-translocation experiments, which demonstrate the functionality of the setup and established protocols.

Polymer capture and translocation through confining pores is a problem of theoretical and practical interest [225–228], particularly in the context of nanopore sensing and sequencing of DNA, RNA, and proteins [148]. If electrical fields serve as driving force for translocations, the non-equilibrium thermodynamics of charged-pore systems in aquaeous electrolyte solutions becomes relevant. In particular the phenomenon of electroosmotic flow (EOF), which arises as an Onsager reciprocal flux in response to applied gradients in electrical potential, complicates the force balance on particles attempting to translocate through nanopores [229]. Typically, EOF becomes pertinent at low salt concentrations where the charges on the walls of the nanochannel are insufficiently screened

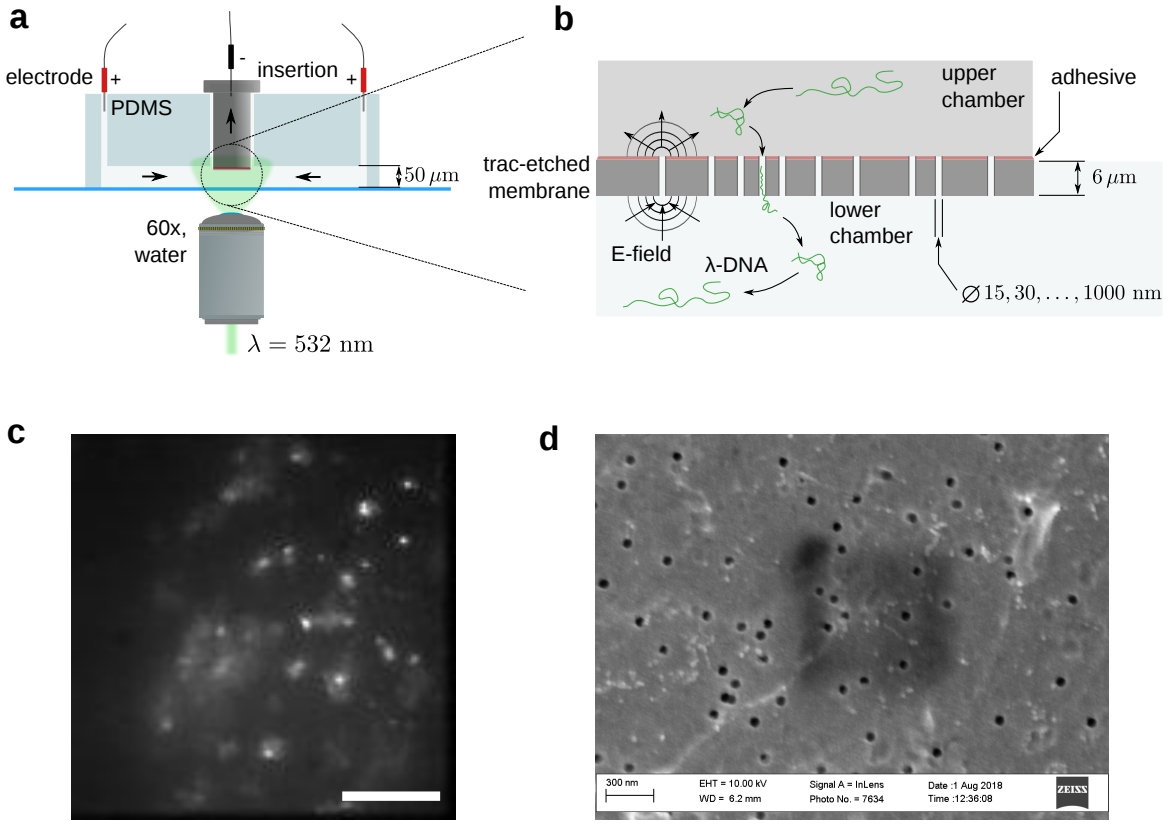

Fig. 5.1 Overview sketch of the DNA translocation experiments on track-etched membranes. **a** Sketch of the microfluidic chip. Arrows indicate direction of the electrical field applied. **b** Close-up sketch of the dynamics of λ-DNA near and within the pores. **c** Time-averaged microscope image to show translocation-event density. The scale bar represents 8 μm.**d** SEM image of gold-labelled 50 nm membranes.

by surrounding ions in the solution, a phenomenon conceptualised in the theory of Debye layer. The associated length scale of this layer is the Debye length, which for a single species of monovalent ions is given by $\lambda_D = \sqrt{\varepsilon k_B T/(2n_\infty e^2)}$ where ε denotes the electrostatic permittivity of the medium and n_∞ represents the bulk concentration of ions with elementary charge e.

The flow strength of the EOF critically depends on λ_D and the surface charge density of the pore. In order to reduce the influence of EOF, one can therefore either attempt to increase the salt concentration to improve screening of wall charges or reduce the surface charge that gives rise to the effect in the first place. This is typically achieved by reducing the pH of the solution, since surface charges in aquaeous electrolytes can often be neutralized or their origin eliminated at low pH [230].

Both strategies, however, may come at a cost in terms of their compatibility with stability or functionality of fluorescent reporters such as intercalating dyes [231].

5.1 Microfluidic design and methods

The microfluidic system that I designed is shown in Fig. 5.1**a**. It consists of two parts: a 3D-printed insertion (shown in dark grey) and a microfluidic chip which accomodates it (shown in light blue). The membrane perforated by the nanopores is glued to the bottom of the insertion, such that it can be placed within $50-100\,\mu$m of the surface of the cover slip. These polycarbonate membranes are commercially available (Whatman) in large quantities and are fabricated using the track-etching technique [232]. This technique has the advantage of producing well-defined nanopores with low dispersion in pore sizes. Available pore diameters range from 15 nm to 1 μm. The pores are irregularly scattered across the membrane with an average distance of $d = 1/\sqrt{\rho_{\text{pore}}} \approx 400$ nm (the pore density is specified by the manufacturer as $\rho_{\text{pore}} = 6 \cdot 10^8/\text{cm}^2$ [233]). The total imaged area in my setup corresponds to $31 \times 31\ \mu\text{m}^2$, which encompasses around 6000 pores (see panel **c**).

However, a number of pores might be blocked by debris. An image of the (gold-coated) surface obtained using a scanning electron microscope is shown in Fig. 5.1**d**.

The electrodes shown in Fig. 5.1**a** consist of platinum to minimise the risk of electrochemical degradation during the experiment. The voltage is provided by a source-meter (2450 Keithley), which is controlled via USB by a custom-written LabView program.

The choice of adhesive in Fig. 5.1**a** proved to be critical and colleagues and I had to go through a number of iterations before we converged to Loctite 3430 (Henkel).

Technical drawings containing all relevant dimensions of the metal mold for the microfluidic chip and the insertion can be found in the appendix in Fig. A.6.

As far as this project is concerned, the optical setup is a standard widefield fluorescence setup. It is, however, in principle capable of performing FRET measurements as described in Hemmig *et al* [234, 224].

The resistance of the track-etched membrane scales inversely with the number of active pores and thus varies considerably from sample to sample due to variations in the amount of excess adhesive. However, for 500 mM KCl and 10 mM MES at pH 5, I measure a resistance of $R \approx 13.7\,k\Omega$. The theoretical value for a single pore with diameter $\varnothing = 50$ nm in this salt solution ($\sigma = 5.4$ S/m) [235], and a membrane thickness of $d = 6\,\mu$m is given by [229]

$$R_{\text{theo, single}} = \sigma^{-1}\left(\frac{d}{\pi\varnothing^2} + \frac{1}{2\varnothing}\right) \approx 6\cdot 10^{10}\,\Omega. \tag{5.1}$$

An active area of a disk with radius $r = 0.1$ mm corresponds to around $n_{\text{pores}} \approx 5\cdot 10^4$ pores, assuming the aforementioned pore density. The corresponding total resistance amounts to $R_{\text{theo, total}} = R_{\text{theo, single}}/n_{\text{pores}} \approx 11887\,\Omega$, which is in the correct regime. However, the estimate of the radius r is probably too small. A realistic estimate would be $r = 0.5$mm, which results in $R_{\text{theo, total}} = 475\,\Omega$. This suggests that a significant number of pores is somehow blocked or other, unknown factors create such high resistance.

5.1.1 Chip assembly protocol

The protocol I used to assemble the microfluidic system in Fig. 5.1**a** is summarised below.

1. (Day before) Place insertion upright on double-sticky tape under 4x-microscope. Apply adhesive in low quantity to top mouth of insertion. Punch out a 3 mm disk from the Whatman membrane with nanopore diameter of choice using a biopsy punch and place carefully on the adhesive. The adhesive is prone to spreading on the membrane. Ensure that the centre of the membrane is not wetted by the adhesive. Store in sealed container overnight.

2. (Before experiment) Wash membrane on insertion using dd$H_2$0, but avoid rinsing the membrane directly. Then, insert the insertion into a PCR vial filled with measurement buffer (filtered with syringe filter, $\varnothing = 220$ nm). Centrifuge for 1 min in table top centrifuge to wet nanopores and remove air bubbles. The insertion is designed so that the membrane will not touch the bottom of the vial.

3. (Before experiment) Punch out central inlet in PDMS chip using a 4 mm biopsy punch (only when PDMS has cooled down from baking). Punch out side-inlets in PDMS chip using 1mm or 1.5 mm biopsy punch.

4. Plasma-bond PDMS chip onto cover slip, then fill chip with measurement buffer.

5. Insert insertion into punched inlet of PDMS chip, avoid formation of air bubbles. The water column in the insertion should not rise, otherwise a leak is likely to have formed.

6. Place PDMS chip on optical setup, fill top or bottom chamber with target solution and commence measurement.

5.2 Experimental results

5.2.1 DNA pore insertion and escape

In this section, I show the results of a proof-of-principle experiment. In this experiment, I insert sytox-orange (Invitrogen, Thermo-fisher) labelled λ-phage DNA into the lower chamber of the setup. The density of sytox orange to DNA-basepairs was set to 1:10, the measurement buffer used contained 3 mM KCl, 0.5x TBE at pH 8 which corresponds to a Debye length of around $\lambda_D \approx 4$ nm. The exposure time is set to 50 ms at a gain of 200x. The 532-laser intensity remained fixed at a current of 67 % of maximum. The pore size used in this experiment was 30 nm.

The idea of the experiment is to observe pore-insertion and subsequent escape of the DNA under positive and negative voltage respectively. This direction of DNA movement would be in line with the anticipated flow direction of EOF. The focus of the microscope rests just below the lower surface of the membrane, ascertained by the presence of autofluorescence coming from the membrane.

Indeed, as Fig. 5.2 shows, the DNA can be observed to disappear into the membrane upon application of $U = +3$ V, but reappears once the voltage is switched to $U = -3$ V. This cycle is repeated 2-3 times as shown.

For reasons currently unknown, a similar experiment in which the DNA is loaded into the upper chamber failed to result in any visible translocation events.

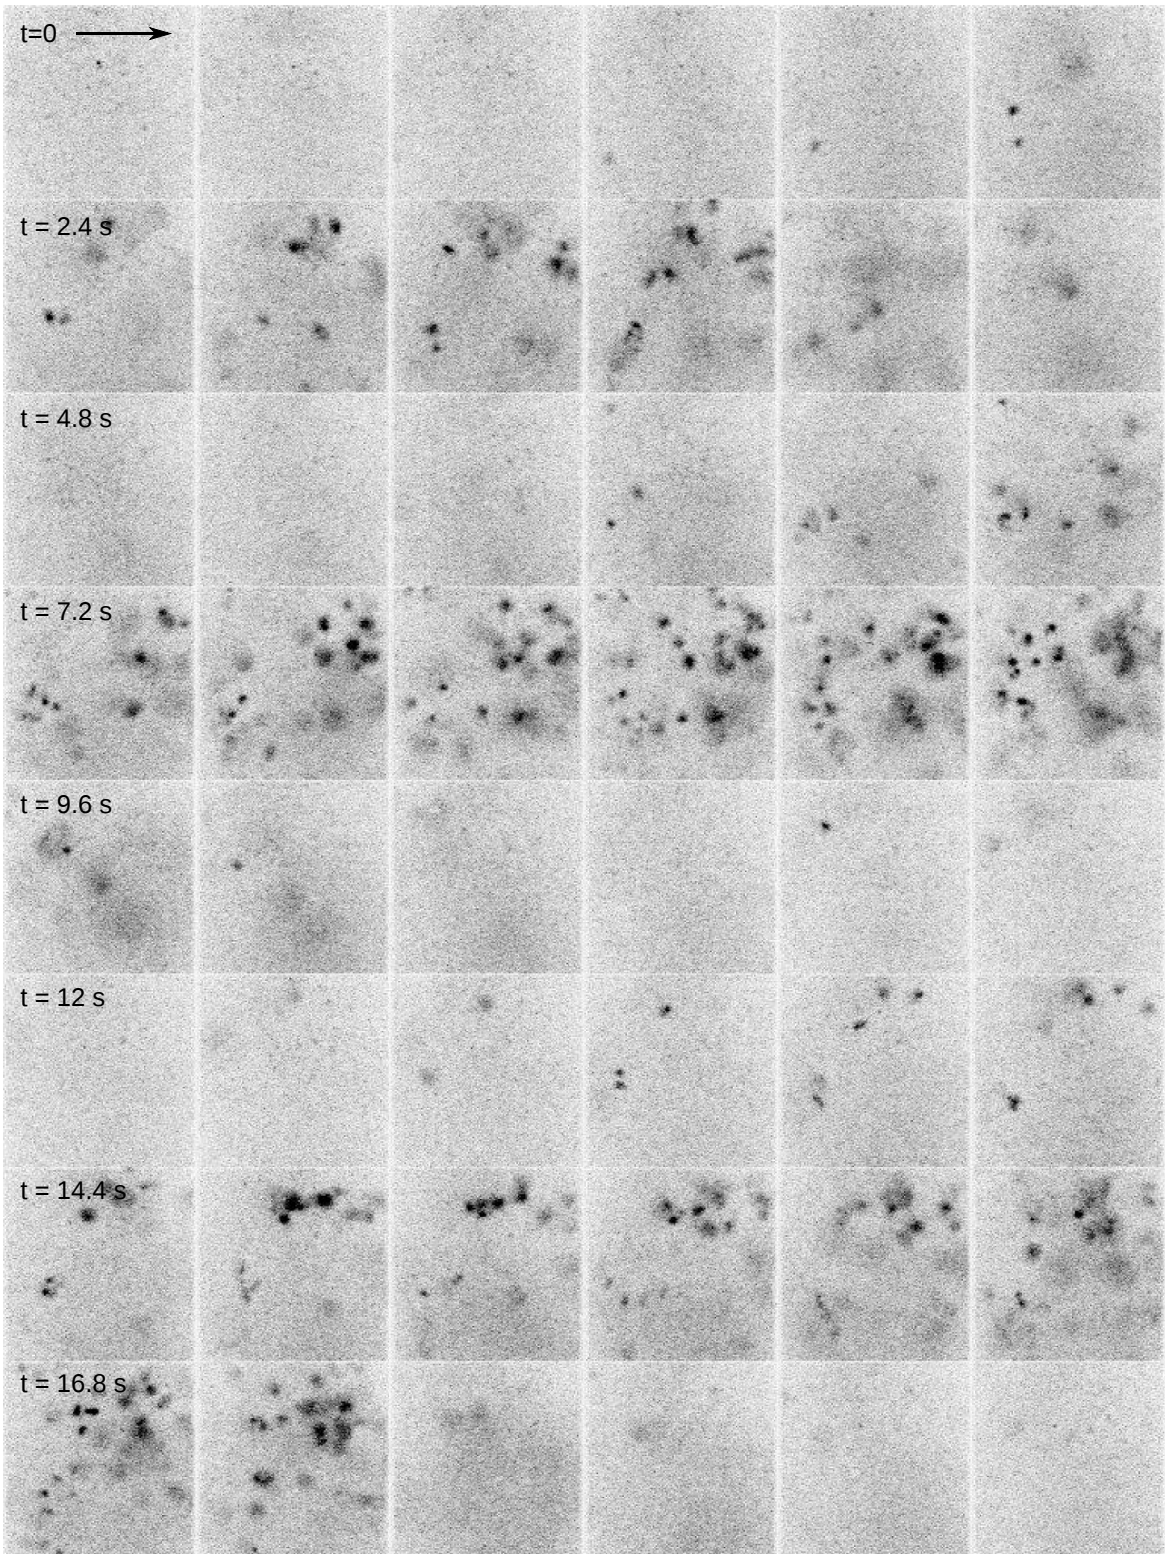

Fig. 5.2 Time series of the ping-pong DNA experiment under switching voltage $U = \pm 3$ V. The images are colour-inverted for better visibility in print. The pore size is $\varnothing = 30$ nm, the buffer contains 3 mM KCl, 0.5x TBE at pH 8.

5.2.2 Visualising DNA translocations

In a subsequent experiment, I increased the pores size to $\varnothing = 50$ nm and applied the λ-phage DNA into the top chamber of the microfluic chip. Imaging parameters are similar to the previous experiment. The voltage was set to various levels in order to scan for the right conditions under which translocations can be detected. A translocation should appear as a bright, brief flash-like event on the camera, since the focus of the microscope was set such that neither the top chamber nor much of the bottom chamber was visible. Interestingly, a few bright events could be detected for almost all voltage levels, but appeared more frequent for negative voltages.

In Fig. 5.3, I give a frame-by-frame time series of a video recorded under $U = -5$ V. I marked a particularly long-lasting event, commencing around $t_0 = 0.85$s, with a red-dashed circle.

Fig. 5.3 shows that counting λ-DNA-translocations appears now to be within reach. For instance, one could apply a Gaussian blur and subsequently subtract each frame from the previous image in the stack. The resultant image stack can be thresholded to provide an event count. Similarly, one could extract an estimate of event durations, another variable of interest in Auger *et al* [222], by measuring for how many frames a certain image region surpasses the threshold. Importantly, events in Fig. 5.3 are likely to span multiple physical processes as I discuss in the conclusion.

5.3 Conclusion

The design of the microfluidic chip has been completed and its functionality has been demonstrated in a first series of proof-of-principle experiments. However, a number of factors remain unknown in the setup. It is unclear whether the DNA interacts in a non-specific way with the polycarbonate as I speculate in Sec. 5.2.1. Once in the pore, the fluorescence of intercalating dyes seems to disappear, which might indicate a quenching interaction between dye and membrane.

For pore sizes larger than 30 nm, I do observe translocation events, albeit with sometimes suspiciously long durations. For instance, the event that I highlight in Fig. 5.3 lasts for around 600-700 ms, vastly longer than typical nanopore translocations (see e.g. [231]). The main drawback (or feature) of the method presented in this chapter consists in its inability to clearly distinguish between attempt and actual translocation. In all likelihood, the data in Fig. 5.3 captures an 'entrance waiting time' caused by entropic barrier effects [226] or EOF-caused entrance trapping [236]. Another possible explanation of

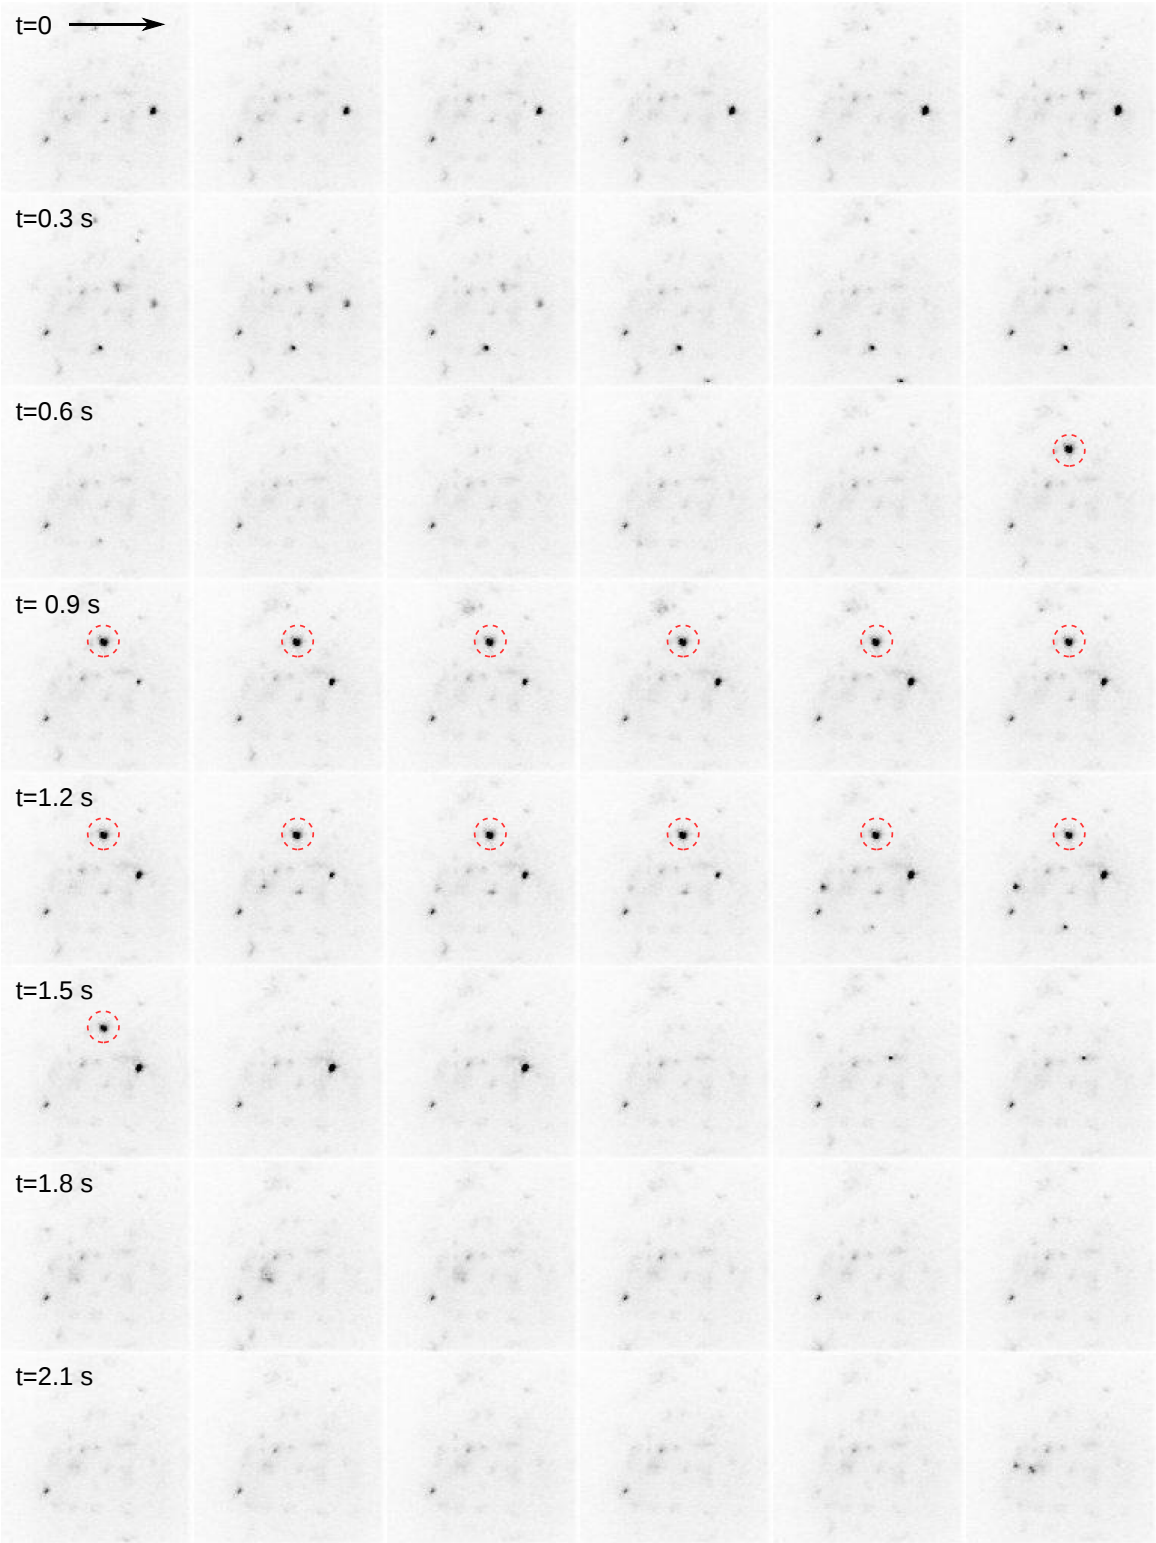

Fig. 5.3 Frame-by-frame series of images recorded under $U = -5$ V. The images are colour-inverted for better visibility in print. Bright (here dark), lasting events are believed to be translocation events from the upper chamber. A long-lasting event is marked with a red-dashed circle.

the long event duration might be the aforementioned unspecific pore-DNA interactions, which I cannot rule out.

Progress in this project was hampered considerably by unreliable adhesives, which often resulted in membranes detaching during the experiment. This, however, is not immediately noticeable, but can only be observed after the experiment, voiding all results obtained during measurements. The current adhesive (Loctite 3430) appears to be more stable under the harsh experimental conditions.

Initial FRET experiments failed to yield conclusive results, especially since the presence of autofluorescence coming from the membrane made it difficult to detect the rather weak FRET signal, despite apparent colocalisation of donor- and acceptor dyes. This problem can potentially be overcome by coating the membrane with a gold layer, which should suppress autofluorescence and increase fluorescence near channel entrances due to the zero-mode-waveguide effect [237, 222]. This, however, appears to require a homogeneous, high-quality coating with gold, which I failed to produce in the shortness of time. Attempts to use custom gold-coated membranes created by vacuum-deposition were unsuccessful due to the formation of what presumably are surface-plasmons. These plasmons resulted in intense, localised background fluorescence. Data from darkfield emission spectroscopy carried out in the group of Prof. J Baumberg corroborates this hypothesis.

Problems arising from low signal-to-noise ratios and the large depth-of-field could, however, be potentially alleviated by exchanging the optical setup for a total-internal-reflection-fluorescence (TIRF) microscope, which requires the nanopore membrane to placed within 50-200 nm of the cover slip. The membrane should, of course, not touch the cover slip but leave space for the measurement solution. I designed the microfluidic mold (see Fig. A.6**b**, **c**) such that the experimenter can opt to create a 50 μm-high PDMS-spacer between the cover slip and the membrane. In order to use a TIRF setup, this spacer cannot be used as the distance it creates would be prohibitive.

Finally, my results in Chapter 3 show that in order to accurately measure first-passage-related effects, one needs to gather large amounts of data. The multi-pore approach described in this section is designed to overcome the limitation inherent to single-pore setups, such that detailed studies of nanopore translocations now come within reach.

Chapter 6

Conclusion

In this thesis, I describe an automated holographic optical tweezers setup, which is capable of performing first-passage time experiments over several days almost without human supervision. I use holographic light shaping here not only to (re)position colloidal particles, but also to create intricate energy landscapes, which allows me to test various theoretical predictions of first-passage-time behaviour. Furthermore, I combined these light-based force landscapes with forces from pressure-driven flows, entropic forces resulting from channels of varying width, and forces resulting from externally applied electrical potentials. The versatility of the setup has grown steadily over time and is now at a point where all the aforementioned forces can be superimposed in a well controlled fashion. Due to the online localisation of colloids, almost all parameters can be driven in a feedback-mode such that the state of the system can be put under active control. For instance, one could create pre-defined force landscapes within microchannels by adjusting the external voltage depending on the position of a confined colloid. A fluctuating electrical field with a spatially-varying variance could be used to model stochastic walks under multiplicative noise. The capacitance of the microchannels appears to be such that the force-change is effectively instantaenous on colloidal timescales.

The first-passage time results desribed in Chapter 3 are encouraging, since they unequivocally support modern thermodynamic theory, often without any adjustable parameters. In this sense, this thesis fits well into the field of applied thermodynamics, which has seen a number of realisations and verifications of thermodynamic predictions, the more famous of which include verifications of fluctuation theorems [238, 239, 22] or a realisation of Landauer's principle [106]. It appears that the underlying assumptions that allow for direct comparison of confined dynamics of colloids with stochastic theory are well fulfilled: Single colloids in PDMS microchannels immersed in low-salt aquaeous solutions indeed behave as hard spheres undergoing confined Brownian walks.

Nevertheless, I regard the escape-rate optimisation that I present in Sec. 3.2 as a significant result. It casts doubt on the general applicability of Arrhenius-like scaling of rates in thermally-induced transitions. Indeed, since the barrier shape can fundamentally alter the dependence of the transition rate on the barrier height, one should consider the full integral in Eq. (1.56) in cases of doubt. Moreover, it would be interesting to see if free-energy landscapes in nature are fine-tuned to such rate-boosting effects. The theory developed by my coauthors covers a wide range of physical regimes; the effect should extend into the inertial regime and higher dimensions. The experiments show that the barrier shapes can indeed be realised, despite the limited agility of holographic tweezers. Our result is therefore quite fundamental and I suspect of wider interest to the community.

The automation routine used in Chapter 3 is supposed to turn the HOT setup into a "real-world Brownian simulator". The setup could therefore potentially benefit from strategies originally devised for molecular or Brownian dynamics simulations. A good example are "steered molecular dynamics" simulations [240], in which a weak harmonic potential is created around the solute of interest (here colloid), that is slowly dragged along the channel. Using Jarzynski's equality, one can then infer the external forces acting on the colloid. This strategy will not work if the external forces are created by the laser itself, since the hologram used to create the weak optical dragging potential will distort the intensity profile. However, the approach might be useful to efficiently map out other forces including entropic, spurious, or electrical forces.

Finally, a technical note: The biggest obstacle in microfluidic experiments turned out to be latent hydrodynamic flows, which typically do not subside even over days. In order to cope with this problem, I had to develop a workflow in which I check in the beginning of each experiment carefully the splitting probabilities in a straight channel without any externally applied forces. While the strength of such flows was markedly reduced after the introduction of macroscopic equilibration channels, which short-circuit both sides of the microchannels, the problem never completely disappeared. The most likely cause seems to be a steady absorption of water by the surrounding PDMS, which then somehow gets rectified into a directed, permanent flow.

Going forward, instead of using colloids, one could try to establish protocols for trapping of oil droplets in water or, conversely, aquaeous droplets in oil. With a sufficiently high difference in refractive indices, these systems should be amenable to optical tweezing. Crucially, the single-particle approach of optical tweezers would alleviate any problems arising from a dispersion of droplet sizes. The experimenter can choose a suitably-sized droplet and isolate it in a microchannel using the tweezers. This would open up a path to, for instance, measurements of depletion forces between oil droplets. Exploratory

experiments with sunflower oil/tween-20 droplets in water were encouraging: I managed to find a small-enough droplet, which fitted into 1 μm-wide channels. The trapping strength was high enough to reliably move the droplet around and the droplet did not stick to the PDMS or glass. In order to take this idea to the next level, one could explore tweezing Sorbitol-enriched-aquaeous droplets in flourinated oil, an established oil/water-system in microfluidic experiments which appears to have just the right ratio of refractive indices and viscosities (see Eqs. (2.1) and (2.2)).

In Chapter 4, I describe what I believe is a novel phase-retrival method based on state-of-the-art machine learning techniques. Conceptually, I treat possible invariances arising in the Fourier coefficient-to-intensity relation using a concept from machine learning known as latent space. I compare different generative models with an expert system with respect to their ability to recreate or approximate given intensity distributions by tweaking the value of 64 Fourier components. Unlike existing phase-retrival algorithms, the approach is model-free and therefore could, in theory, be used on misaligned optical setups. All one has to do is to find a set of contiguous Fourier components which direct the light towards the region of interest. Once this set is identified, one can record training data sets with relative ease. However, I find it conceivable that the approach would also work on the direct hologram-to-intensity relation. Furthermore, I propose a, to my knowledge, novel forward-looking cVAE architecture, which makes use of an additional network to directly descent the loss in the intensity plane.

The crucial question in all machine-learning approaches is how well the method generalises to unseen inputs or how well it approximates impossible intensity distributions. Using validation and test sets, I explore the generalising ability of all models. Clearly, both versions of cVAE proved to be excellent at reproducing previously unseen intensity shapes in the validation data set. Even though the standard cVAE model fared worse on synthetic test set data, I believe that the results in Figs. 4.7 and 4.8 are encouraging. In an application, one could vary the desired intensity perhaps slightly to find the best approximation within the scope of the model. A formal, algorithmic way to do this might be of great benefit.

The machine learning field is currently evolving at a breath-taking pace. Novel methods are being proposed at high frequency. Some of these methods, such as invertible neural networks [241, 242], might be even better suited than the conditional generative approach discussed here.

Perhaps, one could get even higher perfomance with adaptable algorithms, i.e. algorithms that are able to check the intensity distribution they create and propose new holograms based on their previous output. Instead of asking for the best hologram $\mathbf{h}$, one

could ask such a model for the best algorithm to arrive at that hologram. This thinking leads into the territory of reinforcement learning, an area of intense research.

The final results chapter of this thesis, Chapter 5, forms part of the outlook. Progress in this project was slowed down by various technical challenges in establishing a working, but more importantly, a reliable experimental platform for studying nanopore translocations in parallel. The microfluidic design, choice of electrodes, choice of adhesive, and handling protocols are now at a point that serious experimentation can begin: The setup is able to count translocation events in a highly parallel fashion, which should greatly improve the accuracy of experimentally-inferred variables in polymer translocation experiments, an area of intense study. As I discuss at the end of Sec. 5.3, I advocate switching to a TIRF microscope in order to overcome some of the remaining issues identified in experiments.

Even though this thesis is not directly contributing to efforts of finding a possible 'thermodynamic niche' mentioned in Sec. 1, it nevertheless establishes a setup capable of scanning for thermodynamic effects associated with a large variety of parameters and forces. An entirely different blue-sky avenue of future research would therefore be to harness the demonstrated degree of control over laser light and voltage to create controlled temperature- and voltage-gradients in order to activate thermo- and electrophoretically-driven processes within droplets or directly in microchannels.

Appendix A

Appendix

A.1 The multichannel automation routine

The following description of the Labview automation routine can be used as commentary of the source code in the newest version of the modified RedTweezers programme. Fig. A.1 provides a rough idea of the experimental situation for which the routine is designed.

The mulitchannel automation routine is executed in the interface thread. This thread has a refresh rate of 50 Hz and asynchronously requests data from other threads, notably the image processing thread, which manages and refills the list of particle localisations `r_particles`. The routine also interacts with the recording routine, which takes care of storing channel localisations, ROI corner positions, and other meta data into tdms files. Moreover, after each completed automation cycle, the automation routine notifies the parameter-list routine, which then updates the parameter under study, e.g. the phase-gradient parameter p, to the next value in a user-defined list.

The routine is implemented as a state machine, which recomputes its current state s each time the code is invoked. There are five states, each with an integer $s \in \{0$ (`default`), 1 (`moving_to_target`), 2 (`wait_for_other_ROI`), 3 (`record`), 4 (`error_clean_up`)$\}$.

The routine has a number of input parameters, the most relevant of which is a list of user-defined regions-of-interest (ROI), specified by their four corner positions. Each ROI is treated separately, that is, each ROI has its own state variable s. The array of state variables is stored locally in the loop together with a reference to the array of ROI-corner positions, and an array of the current position of particle-track ROIs `particle_follow_roi_list`. The latter are used to check for additional particles in the vicinity of currently actuated particles and there are as many of them as there are channel ROIs. Furthermore, the routine has access to the array of traps `trap_list`, which is an array of containers containing trap positions, trap IDs, and further trap parameters (e.g. trap type (line, point, etc)). The

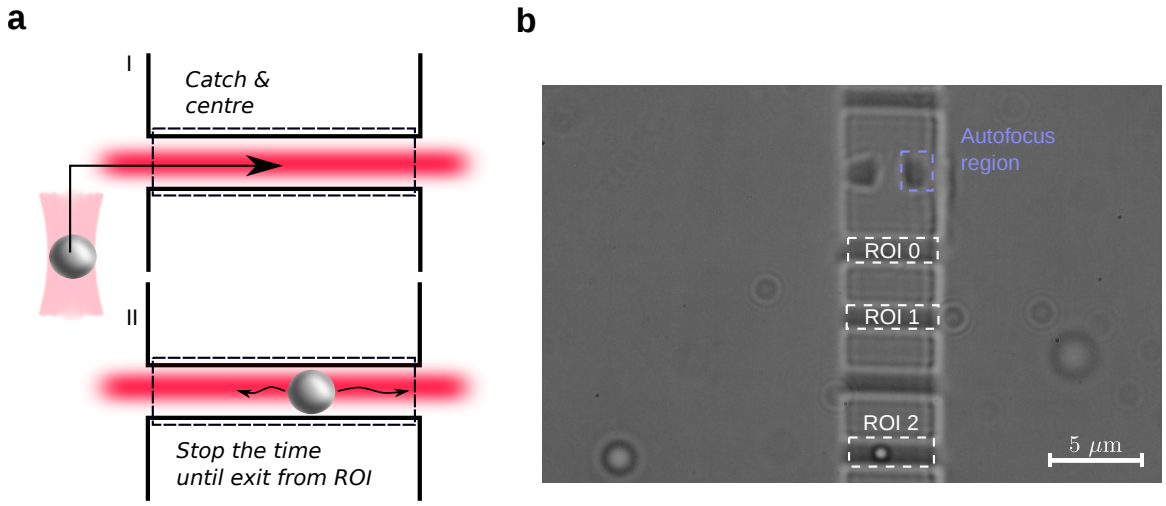

Fig. A.1 Simplified sketch of the automation routine. **a** Two example stages of the automation routine. **b** Exemplary geometry of multi-channel microfluidic mask. Here, the 5 μm mask is shown. The autofocus region (light blue) is used to measure the extend of a dark-appearing cavity, which provides an approximately linear scale of the current axial position of the focal plane relative to the chip.

range of trap IDs used by the automation begins at 666, where 666 itself is reserved for traps engaged in clean-up operations. Each ROI has its own drag-and-drop-trap ID $667+r$ where $r \in \{0,\ldots,R-1\}$.

The user needs to specify a number of parameters, such as (float) `trap_increment`, (bool) `randomise_initial_x_position`, (bool) `randomise_initial_y_position`, (float) `rand_init_x_range`, (float)`rand_init_y_range`, (float) `deterministic_init_x`, (float) `deterministic_init_y`, (bool) `kickout`, (unsigned int) `permissible_out_takes`.

The following list of operations is executed for each ROI $r \in \{0,\ldots,R-1\}$:

(prep) Check the state variable $s[r]$ and switch to that state.
Initialise `out_takes`=`permissible_out_takes` (left assign).

s[r] == 0 Check the current number of particles in the r^{th} ROI using `r_particles`. If this number is smaller or equal to one, find the localisation in `r_particles` that is (i) closest to the centre point of the r^{th} ROI, and (ii) not currently actuated by other ROIs $r' \neq r$. In order to comply with (ii), check if the closest particle is contained in any of the ROIs in `particle_follow_roi_list` if so, find the next best particle.

Get the position of the so-obtained particle (x_p, y_p) and create a particle-follow ROI around that position with a margin of 15 pixels, overwriting the r^{th} ROI in `particle_follow_roi_list`. Furthermore, purge `trap_list` of any traps with ID=$667 + r$, then request a new point trap to be dispatched to (x_p, y_p) by appending a new point trap object to `trap_list` with ID=$667 + r$. The trap-handling routine will then take care of creating the trap.
If, however, there are more than one particle in the ROI, and if the user has set `kickout`== True, then request a new point trap with ID=666 to be dispatched to the left edge on the y-centreline of the ROI. Set $s[r] = 4$. If the user has set `kickout`== False, set $s[r] = 0$.

s[r] == 1 If the number of localisations in `r_particles` that falls within the r^{th}-particle track ROI in `particle_follow_roi_list` is equal to one, find the corresponding trap which is the one with ID==$667 + r$. There should be only one at a time. Then, compute the trap-increment vector $(\Delta t_x, \Delta t_y)$ to the target position in the r^{th} ROI. The step size along x and y is determined by `trap_increment`. The trap completes the y-part of the drag-and-drop operation first to avoid the barrier in Fig. A.1. Update the position of the r^{th} particle-follow ROI as well. If the trap position is closer to the target than a predefined ε, set $\mathbf{s}[r] = 2$, otherwise $s[r] == 1$. If however, the particle number in the r^{th}-particle track ROI in `particle_follow_roi_list` is not equal to one, then delete the trap with ID=$667 + r$ from `trap_list` and set $s[r] = 0$. The moving trap has probably hoovered up several particles, so start anew.

s[r] == 2 Check the number of localisations $n_{\text{ROI}}(r)$ in the r^{th} ROI. If $n_{\text{ROI}}(r) == 1$, set `ready_for_recording`(r) = True. Also, request the `ready_for_recording` state of the other ROIs. If at least one of the other ROIs has `ready_for_recording`== False and $n_{\text{ROI}}(r) \neq 1$, then set $s[r] = 0$, if, however, $n_{\text{ROI}}(r) == 1$, then wait by setting $s[r] = 2$.

If, on the other hand, all `ready_for_recording`(r) == True, set $s[r] = 3$. Furthermore, if $n_{\text{ROI}}(r) = 1$, and $r == R$ with R being the total number of ROIs, then purge all traps with IDs between 666 and $666 + R$ from `trap_list`, wait 1 s, and start the recording routine.

s[r] == 3 Check the current number n_{ROI} of particles in the r^{th} ROI using `r_particles`. If it is greater than one, stop the recording and set $s[r] = 0$.

If is not greater than one and if data is still being recorded, check if all ROIs are empty. If so, check the current value of `out_takes`, which tracks the number of

reported frames in which no particle was localised in any of the ROIs. Decrement this variable. Before that, check if `out_takes`$== 0$, and if so, set $s[r] = 0$ and stop the recording. If `out_takes`> 0, just set $s[r] = 3$ and move on.

If, however, data is not being recorded, set $s[r] = 3$ if $n_{\mathrm{ROI}} == 1$, and $s[r] = 0$ if $n_{\mathrm{ROI}} \neq 1$. This happens, if one of the other ROIs has terminated data acquisition.

s[r] == 4 Get the position of the trap where ID==666 and advance it to the right along the channel with a stepsize determined by `trap_increment`. If the trap has reached its target (ROI centre+500 pixel in x-direction), delete it and set $s[r] = 0$. Otherwise keep this ROI in the error state by setting $s[r] = 4$.

The algorithm outlined above ensures a smooth operation of the localisation routine and, importantly, makes sure that the recording is only terminated once all channels are empty. Dragging new particles into ROIs, while other ROIs record data is a situation that has to be avoided, since the trap-shapes used to create energy-lanscapes in the recording ROIs would be distorted by the dispatched drag-and-drop traps and weakened because of the laser power diverted to them.

A.2 Details of cGAN and cVAE training and architecture

A.2.1 What kind of functions are the generator and discriminator?

The publications of Goodfellow *et al* [198] and Kingma [212] define GAN and VAE as general principles, which do not necessarily rely on neural-networks as underlying regression

Property	S.n. cGAN (rec. loss)	cWGAN	cVAE (forw. loss)
Optimizer:	ADAM [213]	ADAM	RMSprop
Learning rate (η):	10^{-4}	10^{-4}	10^{-4}
Disc. : Gen. updates:	5	5	N.A.
Init. weight std:	0.02	0.02	0.02
Latent space dim.	64	64	64
Random seed:	42	42	42
Batch size:	60	60	60
Epochs:	40 (20)	40	20

Table A.1 Training parameters for all generative models tested in this study. The forward network in the cVAE with forward loss was trained using ADAM, other parameters are similar.

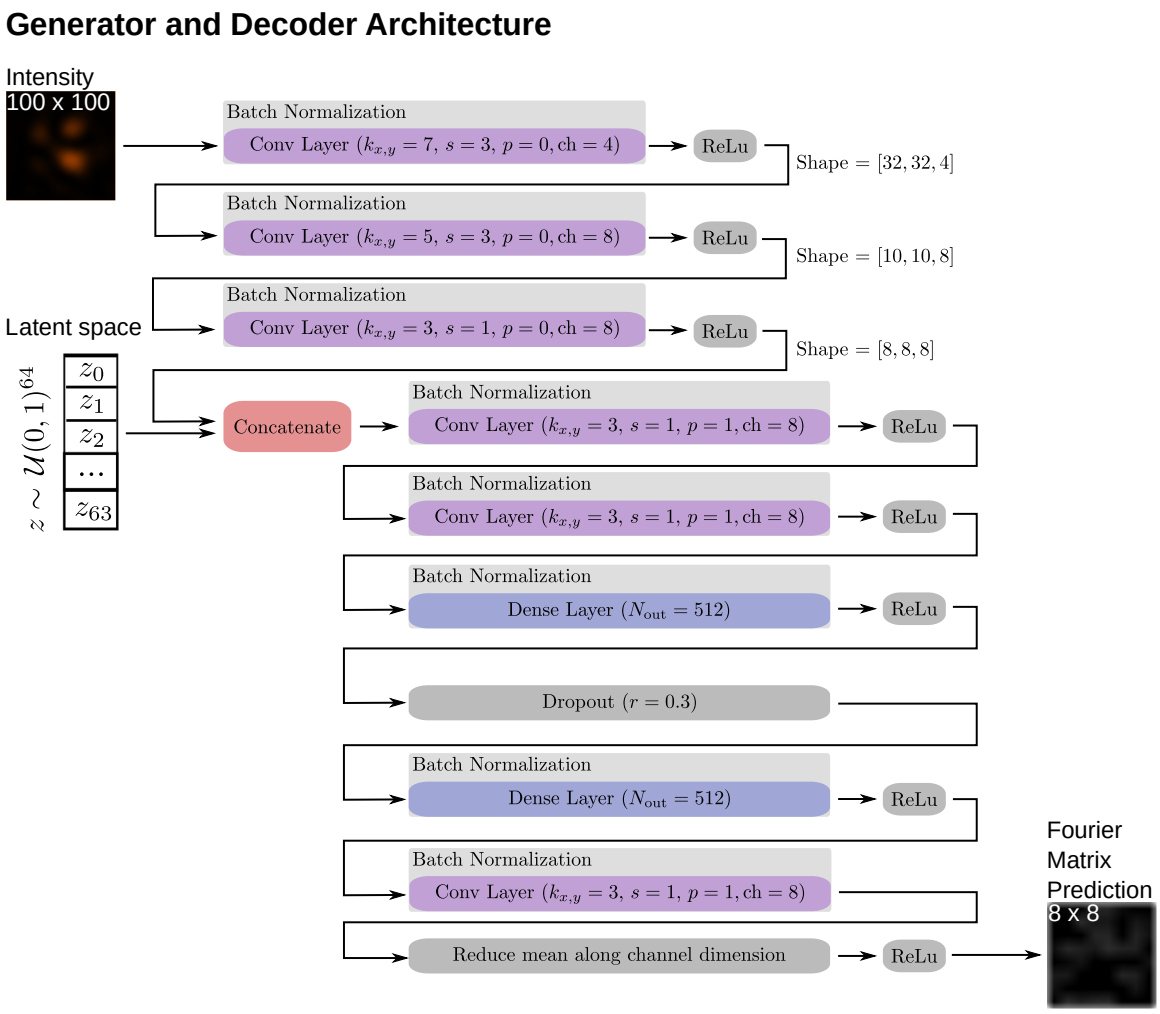

Fig. A.2 Generator and decoder architecture of the spectrally-normalised cGAN model and cVAE model (with and without forward or reconstruction loss).

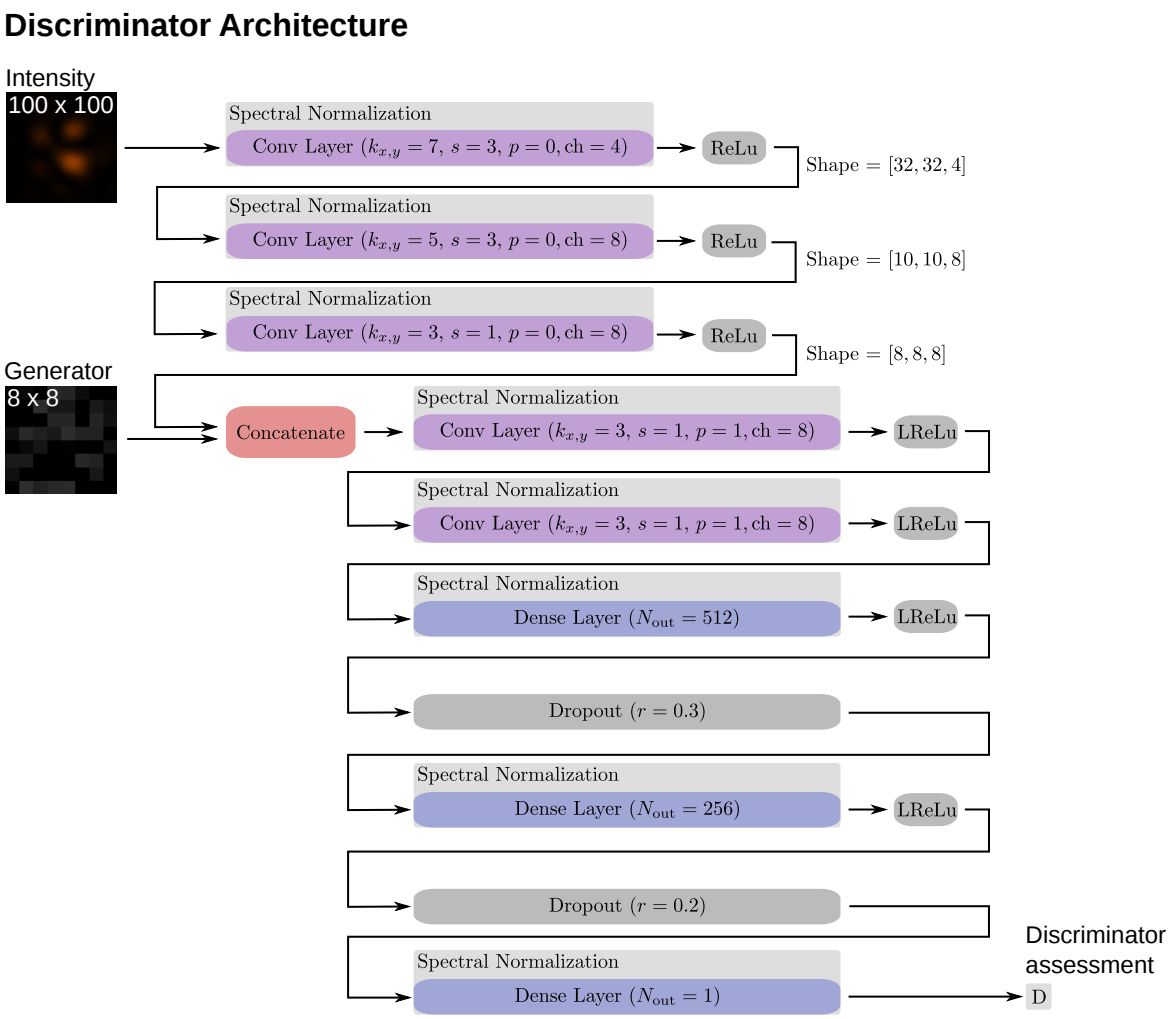

Fig. A.3 Discriminator archicture of the spectrally-normalised cGAN model (with and without reconstruction loss).

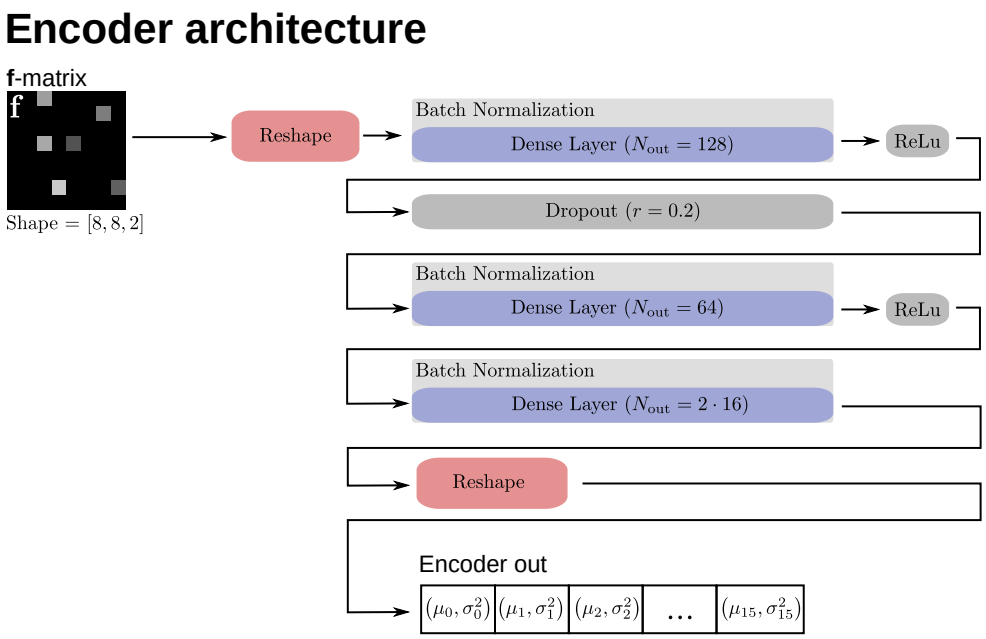

Fig. A.4 Encoder archicture of the cVAE model (with and without forward loss).

engines. In practice, however, GAN and VAE are usually implemented using neural networks. This abstraction highlights an important point: irrespective of their precise form, neural networks behave as mathematical functions, that is, they map inputs to outputs like any other function.

In this case, the generator in a conditional GAN is a function which maps a pair $(\mathbf{I}, z)$ to $\hat{\mathbf{f}}$. The cardinality of its input is thus the sum of cardinalities of $\mathbf{I}$, which is problem-dependent (here 10^4) and z, which is down to choice. The cardinality of the latent variable z is the cardinality of the latent space, which is here fixed to 64. Of course, the precise form of this high-dimensional function called generator depends on the number of layers and the number of nodes chosen for each layer in between input and output (hidden layers). The architecture of the generator used here is shown in Fig. A.2 and is the result of a design decision. I ensured correct in- and output dimensions, which requires a certain minimal number of layers. For instance, as can be seen in Fig. A.2, I need at least three convolutional layers to reduce the 100×100 intensity input to an 8×8-sized tensor using a stride of $s = 3$. Of course, input cardinality could be reduced with fewer layers at the cost of having to use a higher kernel stride, which reduces resolution. The two convolutional layes which follow the concatenation in Fig. A.2, are then used to offer the network the ability to establish a sense of locality over the latent space.

Importantly, purely convolutional networks lack the ability to model differences in behaviour between different locations of the image. While the mapping I am interested in

should, in principle, be translationally invariant from **f**-matrix space to intensity space, optical aberrations and misalignment may result in small differences in laser-spot shapes between, say, $S \circ F(f\mathbf{e}_{1,1})$ and $S \circ F(f\mathbf{e}_{8,8})$ with $\mathbf{e}_{ij}$ denoting a Kronecker-matrix. In order to give the network the ability to accomodate such differences, I chose to endow the network with fully-connected layers.

The conditional discriminator is a function which accepts pairs of the form (x, y) and returns a single scalar, D. The architecture chosen here for the discriminator largely followed similar a chain of reasoning as for the generator (e.g. the convolutional side-chain for the intensity input is the same). The discriminator architecture is shown in Fig. A.3.

However, I admit that a more principled hyperparameter and network architecture search as it is done in some studies (e.g. [243, 244]) is beyond the scope of this study and would also require vastly more computational resources.

A.2.2 How to train a conditional GAN

Training of cGAN proceeds in several distinct gradient-descent steps, which I outline below. A graphical version of these steps showing the flows of gradients can be found in Fig. A.5**a**. Every forward pass through any of the networks will cause the network to store the activations required for calculating the gradient (see Sec. 4.1.1). The following is a mixture of pseudocode and text:

For $i = 0$, $i < n_{sample}/n_{batch}$, $i{+}{=}1$
$\mathscr{D}_i \leftarrow$ *Batch sample*$(\{\mathbf{f},\mathbf{I}\}, i, n_{batch})$
For each $(\mathbf{f}_j, \mathbf{I}_j)$ *in* $\mathscr{D}_i$

1. **Discriminator step**:
   (a) (Real sample) Calculate $D_{\text{real}} \leftarrow \text{Disc}(\mathbf{f}_j, \mathbf{I}_j)$.
   (b) Backpropagate result and calculate partial derivative of first term in Eq. (4.9) (without the expectations) with respect to discriminator parameters ϕ and store in variable "Grad_disc".
   (c) (Fake sample) Draw random $z_j \sim \mathscr{U}^{64}$, then generate fake sample $\hat{\mathbf{f}}_j \leftarrow \text{Gen}(\mathbf{I}_j, z_j)$ and calculate $D_{\text{fake}} \leftarrow \text{Disc}(\hat{\mathbf{f}}_j, \mathbf{I}_j)$.
   (d) Backpropagate result and calculate partial derivative of second term in Eq. (4.9) (without the expectations) with respect to discriminator parameters ϕ and add it to Grad_disc. Store the result in batch-gradient store.

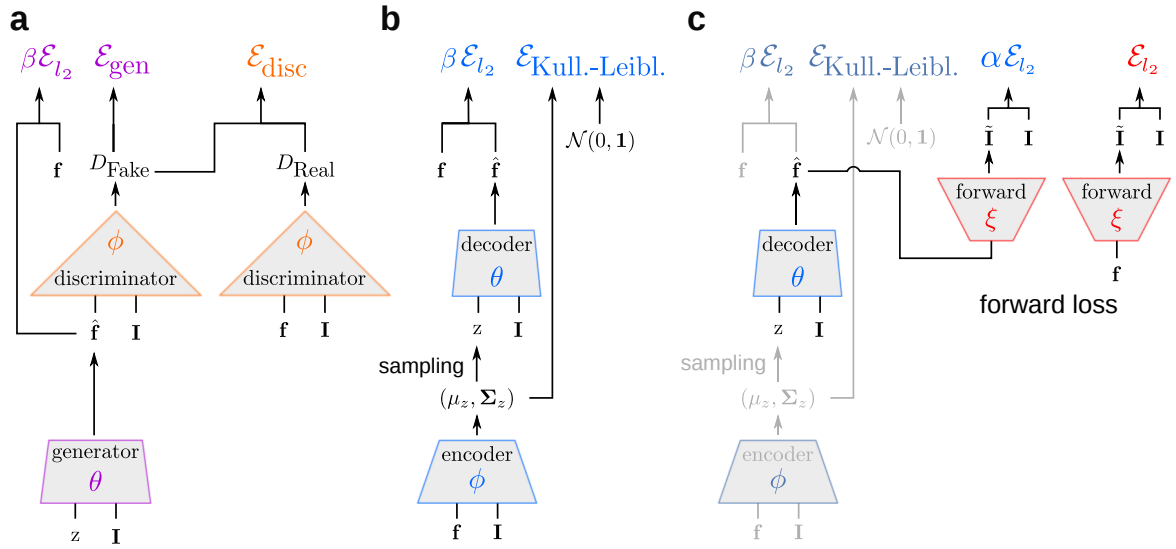

Fig. A.5 Overview of training procedures and gradient flows. The widths of the network-elements (grey shapes) indicate the dimensionality of input and output matrices. The colour code relates the network variable sets $\{\theta, \phi, \xi\}$ to the respective loss terms. For a variable set, such as θ, to be subject to the gradient of a loss, it must have the same colour and a directed connection to it. For instance, **a** Conditional GAN with ($\beta > 0$) or without reconstruction loss ($\beta = 0$). **b** Conditional VAE. **c** Conditional VAE with forward loss. The greyed-out part is the underlying original cVAE architecture shown in panel **b**.

2. **Generator step**:

   (a) If a custom generator loss is used, insert D_{fake} into the generator loss, and backpropagate through discriminator.

   (b) Pass backpropagation δ's from input end of discriminator (without applying a gradient there) to generator and calculate gradient with respect to generator parameters θ. Store this gradient in batch-gradient store.

   (c) Calculate reconstruction loss $E_{\text{rec}} = ||\hat{\mathbf{f}}_j - \mathbf{f}_j||^2_{l_2}$ and backpropagate through generator. Calculate gradient with respect to generator parameters θ and add this gradient to stored gradient.

Average all stored gradients over batch and apply them to their respective parameters.

In modern deep learning frameworks, many of these steps happen without explicit commands. Passing on backpropagation results, for instance, is implicit in concatenated functions, such as $D_\phi(G_\theta(\mathbf{I}, z), \mathbf{I})$ in tensorflow or pytorch.

A.2.3 What kind of functions are the encoder and decoder?

In the original implementation of conditional VAE, the encoder would accept a pair $(\mathbf{f}, \mathbf{I})$ and return a latent space variable z. However, here I use a restricted version which only accepts $\mathbf{f}$ (see Fig. A.4). This works, because we are approximating an actual function, namely $S \circ F$. The input dimension of the encoder is thus 64, not 10,064. The decoder on the other hand is conceptually similar to the generator in cGAN, it accepts pairs of the form $(\mathbf{I}, z)$ and returns Fourier-matrix candidates $\hat{\mathbf{f}}$. Its input dimensions are thus 10,064, since the latent space cardinality is set to 64. The output cardinality is, too, 64 since Fourier matrices are of size 8×8. The architecture chosen here for the decoder is precisely similar to the generator.

A.2.4 How to train a conditional VAE

The training steps involved in training an encoder-decoder pair are similar in some ways to the steps outlined for conditional GAN in Sec. 4.2.2. The flow of gradients is visualised in Fig. A.5**b**.
For $i = 0$, $i < n_{sample}/n_{batch}$, $i+=1$
$\mathscr{D}_i \leftarrow$ *Batch sample*$(\{\mathbf{f}, \mathbf{I}\}, i, n_{batch})$
For each $(\mathbf{f}_j, \mathbf{I}_j)$ *in* $\mathscr{D}_i$

1. **Encoder step**:
   (a) Calculate $(\mu_j, \sigma_j^2) \leftarrow \text{Enc}(\mathbf{f}_j)$.
   (b) Calculate second term in Eq. (4.10) (without expectations) and backpropagate through encoder, calculate gradient with respect to encoder parameters. Store in batch-gradient store.
   (c) Sample $z_j \sim N(\mu_j, \sigma_j^2)$ and pass to decoder.
2. **Decoder step**:
   (a) Calculate $\hat{\mathbf{f}}_j \leftarrow \text{Dec}(\mathbf{I}_j, z_j)$.
   (b) Calculate reconstruction loss (first term in Eq. (4.10)), $||\mathbf{f}_j - \hat{\mathbf{f}}_j||_{l_2}^2$, and backpropagate through decoder, calculate gradient with respect to decoder parameters θ and store in batch-gradient store.
   (c) Pass δ's from decoder input to encoder and calculate gradient there with respect to encoder parameters ϕ. Add to stored encoder gradient.

Average all stored gradients over batch and apply them to their respective parameters.

A.3 Technical sketches of the microfluidic chip in Chapter 5

Fig. A.6 shows the dimensions of the insertion (panel **a**) and the metal mold used to create the PDMS chip (panel **b**). In detail A, I highlight the 50 μm-high spacer defining the height of the micro-channel. The mold is constructed in such a way, that one can opt to use a preperforated chip using the lid shown in panel **c**. The stencil on the lid is designed to give a 3 mm-wide round perforation to accomodate a suitably-sized insertion. Importantly, this leaves a thin 50 μm PDMS layer on the edges of the microchannel on which the insertion and the membrane can rest. This way, one can create a 50μm-minimal distance between the membrane and the cover glass.

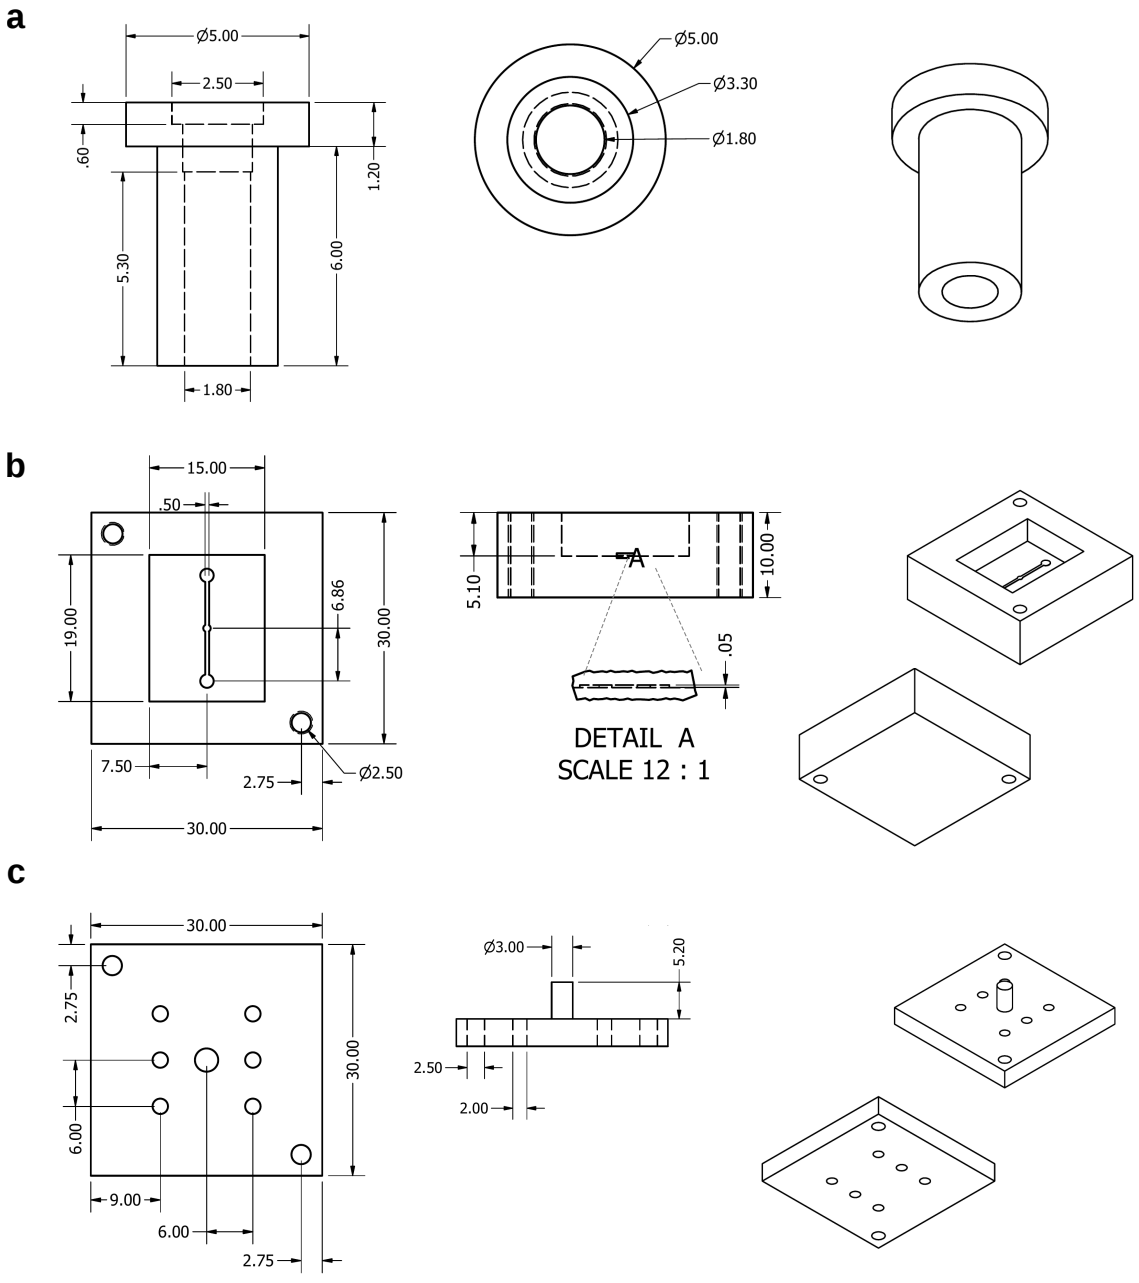

Fig. A.6 All distances are given in mm. **a** Technical sketch of the insertion (3D printable). **b** Technical sketch of the mold used to cast the microfluidic chip in PDMS. **c** Sketch of the lip with a 3 mm-wide stencil used to preperforate the PDMS chip, such that a 50 μm-thin PDMS spacer is created on the rim of the central inlet.

References

[1] S Lloyd. Going into reverse. *Nature*, 430(7003):971, 2004.

[2] CB Mast, N Osterman, and D Braun. Disequilibrium First: The Origin of Life. *J Cosmol*, 10:3305–3314, 2010.

[3] VN Kompanichenko. Inversion Concept of the Origin of Life. *Orig. Life Evol. Biosph.*, 42(2-3):153–178, 2012.

[4] JL England. Statistical physics of self-replication. *J. Chem. Phys.*, 139(12):121923, 2013.

[5] VN Kompanichenko. *Thermodynamic Inversion*. Springer International Publishing, Cham, 2017.

[6] CB Mast, S Schink, U Gerland, and D Braun. Escalation of polymerization in a thermal gradient. *Proc. Natl. Acad. Sci. U.S.A.*, 110(20):8030–8035, 2013.

[7] M Morasch, D Braun, and CB Mast. Heat-Flow-Driven Oligonucleotide Gelation Separates Single-Base Differences. *Angew. Chemie Int. Ed.*, 55(23):6676–6679, 2016.

[8] LMR Keil, FM Möller, M Kieß, PW Kudella, and CB Mast. Proton gradients and pH oscillations emerge from heat flow at the microscale. *Nat. Comm.*, 8(1):1897, 2017.

[9] N Lane. Proton gradients at the origin of life. *BioEssays*, 39(6):1600217, 2017.

[10] I Prigogine. The end of Certainty, time, chaos and new laws of nature. *Cybernetics*, 1(1), 1997.

[11] I Prigogine. Time, Structure, and Fluctuations. *Science*, 201(4358):777–785, 1978.

[12] L Boltzmann. Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. *Sitzungsberichte der Akad. der Wissenschaften*, 66:275–370, 1872.

[13] E Schrödinger. *What is life? The physical aspect of the living cell and mind*. Cambridge University Press Cambridge, 1944.

[14] FS Gnesotto, F Mura, J Gladrow, and CP Broedersz. Broken detailed balance and non-equilibrium dynamics in living systems: a review. *Reports Prog. Phys.*, 81(6):066601, 2018.

[15] R Klages, W Just, and C Jarzynski. *Nonequilibrium Statistical Physics of Small Systems*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2013.

[16] C Jarzynski. Stochastic and Macroscopic Thermodynamics of Strongly Coupled Systems. *Phys. Rev. X*, 7(1):011008, 2017.

[17] DJ Evans, EGD Cohen, and GP Morriss. Probability of second law violations in shearing steady states. *Phys. Rev. Lett.*, 71(15):2401–2404, 1993.

[18] GE Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E*, 60(3):2721–2726, 1999.

[19] C Jarzynski. Nonequilibrium Equality for Free Energy Differences. *Phys. Rev. Lett.*, 78(14):2690–2693, 1997.

[20] U Seifert. Entropy Production along a Stochastic Trajectory and an Integral Fluctuation Theorem. *Phys. Rev. Lett.*, 95(4):40602, 2005.

[21] U Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports Prog. Phys.*, 75(12):126001, 2012.

[22] TM Hoang, R Pan, J Ahn, J Bang, HT Quan, and T Li. Experimental Test of the Differential Fluctuation Theorem and a Generalized Jarzynski Equality for Arbitrary Initial States. *Phys. Rev. Lett.*, 120(8):080602, 2018.

[23] RD Astumian and M Bier. Fluctuation driven ratchets: Molecular motors. *Phys. Rev. Lett.*, 72(11):1766–1769, 1994.

[24] H Wang and G Oster. Ratchets, power strokes, and molecular motors. *Appl. Phys. A Mater. Sci. Process.*, 75(2):315–323, 2002.

[25] P Hänggi and F Marchesoni. Artificial Brownian motors: Controlling transport on the nanoscale. *Rev. Mod. Phys.*, 81(1):387–442, 2009.

[26] AC Barato and U Seifert. Thermodynamic Uncertainty Relation for Biomolecular Processes. *Phys. Rev. Lett.*, 114(15):158101, 2015.

[27] P Pietzonka, AC Barato, and U Seifert. Universal bounds on current fluctuations. *Phys. Rev. E*, 93(5):1–16, 2016.

[28] TR Gingrich, JM Horowitz, N Perunov, and JL England. Dissipation Bounds All Steady-State Current Fluctuations. *Phys. Rev. Lett.*, 116(12):120601, 2016.

[29] P Mehta and DJ Schwab. Energetic costs of cellular computation. *Proc. Natl. Acad. Sci. U.S.A.*, 109(44):17978–17982, 2012.

[30] G Lan, P Sartori, S Neumann, V Sourjik, and Y Tu. The energy–speed–accuracy trade-off in sensory adaptation. *Nat. Phys.*, 8(5):422–428, 2012.

[31] P Sartori, L Granger, CF Lee, and JM Horowitz. Thermodynamic Costs of Information Processing in Sensory Adaptation. *PLoS Comput. Biol.*, 10(12):e1003974, 2014.

[32] Y Cao, H Wang, Q Ouyang, and Y Tu. The free-energy cost of accurate biochemical oscillations. *Nat. Phys.*, 11(9):772–778, 2015.

[33] JMR Parrondo, JM Horowitz, and T Sagawa. Thermodynamics of information. *Nat. Phys.*, 11(2):131–139, 2015.

[34] JJ Hopfield. Kinetic Proofreading: A New Mechanism for Reducing Errors in Biosynthetic Processes Requiring High Specificity. *Proc. Natl. Acad. Sci. U.S.A.*, 71(10):4135–4139, 1974.

[35] L Bachelier. Theorie de La Speculation. In *Louis Bachelier's Theory Specul. Orig. Mod. Financ.* 1900.

[36] P Wilmott. *Paul Wilmott introduces quantitative finance*. John Wiley & Sons, Ltd, London, 2nd edition, 2007.

[37] C Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer-Verlag, 2009.

[38] A Einstein. Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Ann. Phys.*, 322(8):549–560, 1905.

[39] W Sutherland. LXXV. A dynamical theory of diffusion for non-electrolytes and the molecular mass of albumin. *London, Edinburgh, Dublin Philos. Mag. J. Sci.*, 9(54):781–785, 1905.

[40] JL Lebowitz and E Rubin. Dynamical study of brownian motion. *Phys. Rev.*, 131(6):2381–2396, 1963.

[41] L Bocquet and JL Barrat. Hydrodynamic boundary conditions, correlation functions, and Kubo relations for confined fluids. *Phys. Rev. E*, 49(4):3079–3092, 1994.

[42] I Pagonabarraga, MHJ Hagen, CP Lowe, and D Frenkel. Short-time dynamics of colloidal suspensions in confined geometries. *Phys. Rev. E*, 59(4):4458–4469, 1999.

[43] G Hummer and IC Yeh. System-size dependence of diffusion coefficients and viscosities from molecular dynamics simulations with periodic boundary conditions. *J. Phys. Chem. B*, 108(40):15873–15879, 2004.

[44] SL Dettmer, S Pagliara, K Misiunas, and UF Keyser. Anisotropic diffusion of spherical particles in closely confining microchannels. *Phys. Rev. E*, 89(6):062305, 2014.

[45] K Misiunas, S Pagliara, E Lauga, JR Lister, and UF Keyser. Nondecaying Hydrodynamic Interactions along Narrow Channels. *Phys. Rev. Lett.*, 115(3):1–6, 2015.

[46] B Fabry, GN Maksym, JP Butler, M Glogauer, D Navajas, and JJ Fredberg. Scaling the microrheology of living cells. *Phys. Rev. Lett.*, 87(14):148102, 2001.

[47] B Schnurr, F Gittes, FC MacKintosh, and CF Schmidt. Determining Microscopic Viscoelasticity in Flexible and Semiflexible Polymer Networks from Thermal Fluctuations. *Macromolecules*, 30(25):7781–7792, 1997.

[48] F Gittes, B Schnurr, PD Olmsted, FC MacKintosh, and CF Schmidt. Microscopic Viscoelasticity: Shear Moduli of Soft Materials Determined from Thermal Fluctuations. *Phys. Rev. Lett.*, 79(17):3286–3289, 1997.

[49] AWC Lau, BD Hoffman, A Davies, JC Crocker, and TC Lubensky. Microrheology, Stress Fluctuations, and Active Behavior of Living Cells. *Phys. Rev. Lett.*, 91(19):198101, 2003.

[50] M Doi and SF Edwards. *The Theory of Polymer Dynamics*. Clarendon Press, Oxford, 1. edition, 1986.

[51] R Kubo. The fluctuation-dissipation theorem. *Reports Prog. Phys.*, 29(1):306, 1966.

[52] J Prost, JF Joanny, and J Parrondo. Generalized Fluctuation-Dissipation Theorem for Steady-State Systems. *Phys. Rev. Lett.*, 103(9):90601, 2009.

[53] BJ Alder and TE Wainwright. Decay of the Velocity Autocorrelation Function. *Phys. Rev. A*, 1(1):18–21, 1970.

[54] R Zwanzig and M Bixon. Hydrodynamic theory of the velocity correlation function. *Phys. Rev. A*, 2(5):2005, 1970.

[55] L Onsager. Reciprocal Relations in Irreversible Processes. II. *Phys. Rev.*, 38(12):2265–2279, 1931.

[56] J Gladrow, CP Broedersz, and CF Schmidt. Nonequilibrium dynamics of probe filaments in actin-myosin networks. *Phys. Rev. E*, 96(2):022408, 2017.

[57] FC MacKintosh. Notes on Linear Response Theory. Technical report, Vrije Universiteit, 2006.

[58] T Franosch and S Jeney. Persistent correlation of constrained colloidal motion. *Phys. Rev. E*, 79(3):031402, 2009.

[59] T Franosch, M Grimm, M Belushkin, FM Mor, G Foffi, L Forró, and S Jeney. Resonances arising from hydrodynamic memory in Brownian motion. *Nature*, 478(7367):85–88, 2011.

[60] R Huang, I Chavez, KM Taute, B Lukić, S Jeney, MG Raizen, and EL Florin. Direct observation of the full transition from ballistic to diffusive Brownian motion in a liquid. *Nat. Phys.*, 7(7):576–580, 2011.

[61] J Gladrow, N Fakhri, FC MacKintosh, CF Schmidt, and CP Broedersz. Broken Detailed Balance of Filament Dynamics in Active Networks. *Phys. Rev. Lett.*, 116(24):248301, 2016.

[62] HS Chung, JM Louis, and WA Eaton. Experimental determination of upper bound for transition path times in protein folding from single-molecule photon-by-photon trajectories. *Proc. Natl. Acad. Sci. U.S.A.*, 106(29):11837–11844, 2009.

[63] H Yu, AN Gupta, X Liu, K Neupane, AM Brigley, I Sosova, and MT Woodside. Energy landscape analysis of native folding of the prion protein yields the diffusion constant, transition path time, and rates. *Proc. Natl. Acad. Sci. U.S.A.*, 109(36):14452–14457, 2012.

[64] HS Chung and WA Eaton. Single-molecule fluorescence probes dynamics of barrier crossing. *Nature*, 502(7473):685–688, 2013.

[65] HS Chung, S Piana-Agostinetti, DE Shaw, and WA Eaton. Structural origin of slow diffusion in protein folding. *Science*, 349(6255):1504–1510, 2015.

[66] K Truex, H S Chung, J M Louis, and WA Eaton. Testing Landscape Theory for Biomolecular Processes with Single Molecule Fluorescence Spectroscopy. *Phys. Rev. Lett.*, 115(1):1–5, 2015.

[67] K Neupane, AP Manuel, and MT Woodside. Protein folding trajectories can be described quantitatively by one-dimensional diffusion over measured energy landscapes. *Nat. Phys.*, 12(7):700–703, 2016.

[68] GR Marshall, JA Feng, and DJ Kuster. Back to the future: Ribonuclease A. *Biopolymers*, 90(3):259–277, 2008.

[69] W Nadler, AT Brunger, K Schulten, and M Karplus. Molecular and stochastic dynamics of proteins. *Proc. Natl. Acad. Sci. U.S.A.*, 84(22):7933–7937, 1987.

[70] A Borgia, PM Williams, and J Clarke. Single-Molecule Studies of Protein Folding. *Annu. Rev. Biochem.*, 77(1):101–125, 2008.

[71] AN Gupta, A Vincent, K Neupane, H Yu, F Wang, and MT Woodside. Experimental validation of free-energy-landscape reconstruction from non-equilibrium single-molecule force spectroscopy measurements. *Nat. Phys.*, 7(8):631–634, 2011.

[72] K Neupane, DAN Foster, DR Dee, H Yu, F Wang, and MT Woodside. Direct observation of transition paths during the folding of proteins and nucleic acids. *Science*, 352(6282):239–242, 2016.

[73] K Neupane, F Wang, and MT Woodside. Direct measurement of sequence-dependent transition path times and conformational diffusion in DNA duplex formation. *Proc. Natl. Acad. Sci. U.S.A.*, 114(6):1329–1334, 2017.

[74] P Cossio, G Hummer, and A Szabo. Transition paths in single-molecule force spectroscopy. *J. Chem. Phys.*, 148(12):123309, 2018.

[75] JD Wen, L Lancaster, C Hodges, AC Zeri, SH Yoshimura, HF Noller, C Bustamante, and I Tinoco. Following translation by single ribosomes one codon at a time. *Nature*, 452(7187):598–603, 2008.

[76] K Neupane, NQ Hoffer, and MT Woodside. Measuring the Local Velocity along Transition Paths during the Folding of Single Biological Molecules. *Phys. Rev. Lett.*, 121(1):018102, 2018.

[77] SW Englander and L Mayne. The nature of protein folding pathways. *Proc. Natl. Acad. Sci. U.S.A.*, 111(45):15873–15880, 2014.

[78] J Neupane, AP Manuel, J Lambert, and MT Woodside. Transition-path probability as a test of reaction-coordinate quality reveals DNA hairpin folding is a one-dimensional diffusive process. *J. Phys. Chem. Lett.*, 6(6):1005–1010, 2015.

[79] SW Englander and L Mayne. The case for defined protein folding pathways. *Proc. Natl. Acad. Sci. U.S.A.*, 114(31):8253–8258, 2017.

[80] SW Englander and L Mayne. Reply to Eaton and Wolynes: How do proteins fold? *Proc. Natl. Acad. Sci. U.S.A.*, 114(46):E9761–E9762, 2017.

[81] WA Eaton and PG Wolynes. Theory, simulations, and experiments show that proteins fold by multiple pathways. *Proc. Natl. Acad. Sci. U.S.A.*, 114(46):E9759–E9760, 2017.

[82] R Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford University Press, New York, 1 edition, 2001.

[83] I Neri, É Roldán, and F Jülicher. Statistics of Infima and Stopping Times of Entropy Production and Applications to Active Molecular Processes. *Phys. Rev. X*, 7(1):011019, 2017.

[84] S Pigolotti, I Neri, É Roldán, and F Jülicher. Generic Properties of Stochastic Entropy Production. *Phys. Rev. Lett.*, 119(14):140604, 2017.

[85] É Roldán, I Neri, M Dörpinghaus, H Meyr, and F Jülicher. Decision Making in the Arrow of Time. *Phys. Rev. Lett.*, 115(25):250602, 2015.

[86] JB Weiss. Coordinate invariance in stochastic dynamical systems. *Tellus, Ser. A Dyn. Meteorol. Oceanogr.*, 55(3):208–218, 2003.

[87] BW Zhang, D Jasnow, and DM Zuckerman. Transition-event durations in one-dimensional activated processes. *J. Chem. Phys.*, 126(7), 2007.

[88] WK Kim and RR Netz. The mean shape of transition and first-passage paths. *J. Chem. Phys.*, 143(22), 2015.

[89] JO Daldrop, WK Kim, and RR Netz. Transition paths are hot. *EPL*, 113(1):18004, 2016.

[90] AM Berezhkovskii, MA Pustovoit, and SM Bezrukov. Channel-facilitated membrane transport: Average lifetimes in the channel. *J. Chem. Phys.*, 119(7):3943–3951, 2003.

[91] AM Berezhkovskii, G Hummer, and SM Bezrukov. Identity of Distributions of Direct Uphill and Downhill Translocation Times for Particles Traversing Membrane Channels. *Phys. Rev. Lett.*, 97(2):020601, 2006.

[92] AM Berezhkovskii, L Dagdug, and SM Bezrukov. Mean Direct-Transit and Looping Times as Functions of the Potential Shape. *J. Phys. Chem. B*, 121(21):5455–5460, 2017.

[93] AB Kolomeisky, EB Stukalin, and AA Popov. Understanding mechanochemical coupling in kinesins using first-passage-time processes. *Phys. Rev. E*, 71(3):1–8, 2005.

[94] J Gladrow, M Ribezzi-Crivellari, F Ritort, and UF Keyser. Experimental evidence of symmetry breaking of transition-path times. *Nat. Comm.*, 10(1):55, 2019.

[95] F Stern. An Independence in Brownian Motion with constant drift. *Ann. Probab.*, 5(4):571–572, 1977.

[96] JG Wendel. Hitting Spheres with Brownian. *Ann. Probab.*, 8(1):164–169, 2019.

[97] C Yin and C Wang. Hitting Time and Place of Brownian Motion with Drift. *Open Stat. Probab. J.*, 1(1):38–42, 2009.

[98] S Redner. *A Guide To First-Passage Processes*. Cambridge University Press, Cambridge, 1st edition, 2001.

[99] HA Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, 1940.

[100] A Ashkin, JM Dziedzic, JE Bjorkholm, and S Chu. Observation of a single-beam gradient force optical trap for dielectric particles. *Opt. Lett.*, 11(5):288–290, 1986.

[101] A Ashkin. Forces of a single-beam gradient laser trap on a dielectric sphere in the ray optics regime. *Biophys. J.*, 61(2):569–582, 1992.

[102] A Ashkin. Applications of Laser Radiation Pressure. *Science*, 210(4474):1081–1088, 1980.

[103] K Svoboda and SM Block. Biological Applications of Optical Forces. *Annu. Rev. Biophys. Biomol. Struct.*, 23(1):247–285, 1994.

[104] GM Wang, EM Sevick, E Mittag, DJ Searles, and DJ Evans. Experimental Demonstration of Violations of the Second Law of Thermodynamics for Small Systems and Short Time Scales. *Phys. Rev. Lett.*, 89(5):50601–50605, 2002.

[105] DM Carberry, JC Reid, GM Wang, EM Sevick, DJ Searles, and DJ Evans. Fluctuations and Irreversibility: An Experimental Demonstration of a Second-Law-Like Theorem Using a Colloidal Particle Held in an Optical Trap. *Phys. Rev. Lett.*, 92(14):140601, 2004.

[106] A Bérut, A Arakelyan, A Petrosyan, S Ciliberto, R Dillenschneider, and E Lutz. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.

[107] M Ribezzi-Crivellari and F Ritort. Large work extraction and the Landauer limit in a continuous Maxwell demon. *Nat. Phys.*, (7):660–664, 2019.

[108] J Gladrow. *Broken detailed balance in active matter - theory, simulation & experiments*. PhD thesis, 2015.

[109] M Woerdemann. *Structured Light Fields Applications in Optical Trapping, Mainpulation and organisation*. PhD thesis, University of Münster, 2012.

[110] DG Grier. Optical tweezers in colloid and interface science. *Curr. Opin. Colloid Interface Sci.*, 2010.

[111] W Demtröder. *Experimentalphysik 2 - Elektrizität und Optik*. Springer Verlag, 2017.

[112] Y Roichman, B Sun, A Stolarski, and DG Grier. Influence of Nonconservative Optical Forces on the Dynamics of Optically Trapped Colloidal Spheres: The Fountain of Probability. *Phys. Rev. Lett.*, 101(12):128301, 2008.

[113] B Sun, J Lin, E Darby, AY Grosberg, and DG Grier. Brownian vortexes. *Phys. Rev. E*, 80(1):010401, 2009.

[114] G Gouesbet. Generalized Lorenz–Mie theories, the third decade: A perspective. *J. Quant. Spectrosc. Radiat. Transf.*, 110(14-16):1223–1238, 2009.

[115] D Gabor. Holography, 1948-1971. *Science*, 177(4046):299–313, 1972.

[116] DG Grier. A revolution in optical manipulation. *Nature*, 424(6950):810–816, 2003.

[117] ER Dufresne, GC. Spalding, MT Dearing, SA Sheets, and DG Grier. Computer-generated holographic optical tweezer arrays. *Rev. Sci. Instrum.*, 72(3):1810, 2001.

[118] Y Roichman and DG Grier. Projecting extended optical traps with shape-phase holography. *Opt. Lett.*, 31(11):1675, 2006.

[119] B Sun, Y Roichman, and DG Grier. Theory of holographic optical trapping. *Opt. Express*, 16(20):15765, 2008.

[120] JE Curtis, BA Koss, and DG Grier. Dynamic Holographic Optical Tweezers. *Opt. Commun.*, 217(1-6):169–175, 2012.

[121] J Mertz. *Introduction to Optical Microscopy*. Roberts and Company, 2009.

[122] A Jesacher. *Applications of spatial light modulators for optical trapping and image processing*. PhD thesis, Leopold-Franzens University, Innsbruck, 2007.

[123] RW Bowman, GM Gibson, A Linnenberger, DB Phillips, JA Grieve, DM Carberry, S Serati, MJ Miles, and MJ Padgett. Red tweezers: Fast, customisable hologram generation for optical tweezers. *Comput. Phys. Commun.*, 185(1):268–273, 2014.

[124] Y Roichman, B Sun, Y Roichman, J Amato-Grill, and DG Grier. Optical Forces Arising from Phase Gradients. *Phys. Rev. Lett.*, 100(1):013602, 2008.

[125] KC Neuman and SM Block. Optical trapping. *Anal. Chem.*, 67(17):565A–565A, 1995.

[126] IPG Laser GmbH. YLM-5-1064-LP User's Guide. Technical report, 2006.

[127] Y Tan. *Enhanced transport through confined channels by stationary and fluctuating potentials*. PhD thesis, University of Cambridge, 2018.

[128] Mikrotron GmbH. EoSens® CL High-Speed CMOS Camera. Technical report, Unterschleissheim, 2016.

[129] R Bowman, D Preece, G Gibson, and M Padgett. Stereoscopic particle tracking for 3D touch, vision and closed-loop control in optical tweezers. *J. Opt.*, 13(4):044003, 2011.

[130] RW Gerchberg. Holography without Fringes in the Electron Microscope. *Nature*, 240(5381):404–406, 1972.

[131] JE Curtis, CHJ Schmitz, and JP Spatz. Symmetry dependence of holograms for optical trapping. *Opt. Lett.*, 30(16):2086–2088, 2005.

[132] R Di Leonardo, F Ianni, and G Ruocco. Computer generation of optimal holograms for optical trap arrays. *Opt. Express*, 15:1913, 2007.

[133] K Dholakia and T Čižmár. Shaping the future of manipulation. *Nature photonics*, 5(6):335, 2011.

[134] P Pozzi, L Maddalena, N Ceffa, O Soloviev, G Vdovin, E Carroll, and M Verhaegen. Fast Calculation of Computer Generated Holograms for 3D Photostimulation through Compressive-Sensing Gerchberg–Saxton Algorithm. *Methods Protoc.*, 2(1):2, 2018.

[135] X Yang, C Liu, Y Li, F Marchesoni, P Hänggi, and HP Zhang. Hydrodynamic and entropic effects on colloidal diffusion in corrugated channels. *Proc. Natl. Acad. Sci. U.S.A.*, 114(36):9564–9569, 2017.

[136] B Schuler, EA Lipman, and WA Eaton. Probing the free-energy surface for protein folding with single-molecule fluorescence spectroscopy. *Nature*, 419(6908):743–747, 2002.

[137] HS Chung, K McHale, JM Louis, and WA Eaton. Single-Molecule Fluorescence Experiments Determine Protein Folding Transition Path Times. *Science*, 335(6071):981–984, 2012.

[138] MT Woodside, PC Anthony, WM Behnke-Parks, K Larizadeh, D Herschlag, and Steven M Block. Direct measurement of the full, sequence-dependent folding landscape of a nucleic acid. *Science*, 314(5801):1001–1004, 2006.

[139] LI McCann, M Dykman, and B Golding. Thermally activated transitions in a bistable three-dimensional optical trap. *Nature*, 402(6763):785–787, 1999.

[140] EM Nestorovich, C Danelon, M Winterhalter, and SM Bezrukov. Designed to penetrate: Time-resolved interaction of single antibiotic molecules with bacterial pores. *Proc. Natl. Acad. Sci. U.S.A.*, 99(15):9789–9794, 2002.

[141] JM Pagès, CE James, and M Winterhalter. The porin and the permeating antibiotic: a selective diffusion barrier in Gram-negative bacteria. *Nat. Rev. Microbiol.*, 6(12):893–903, 2008.

[142] WD Stein and T Litman. *Channels, Carriers, and Pumps*. Academic Press, 2 edition, 2015.

[143] C Dekker. Solid-state nanopores. *Nat. Nanotechnol.*, 2(4):209–215, 2007.

[144] NAW Bell, CR Engst, M Ablay, G Divitini, C Ducati, Y Liedl, and UF Keyser. DNA Origami Nanopores. *Nano Lett.*, 12(1):512–517, 2012.

[145] MJ Skaug, C Schwemmer, S Fringes, CD Rawlings, and AW Knoll. Nanofluidic rocking Brownian motors. *Science*, 359(6383):1505–1508, 2018.

[146] P Hänggi, P Talkner, and M Borkovec. Reaction-rate theory: Fifty years after Kramers. *Rev. Mod. Phys.*, 62(2):251–341, 1990.

[147] L Onsager and S Machlup. Fluctuations and Irreversible Processes. *Phys. Rev.*, 91(6):1505–1512, 1953.

[148] JJ Kasianowicz, E Brandin, D Branton, and DW Deamer. Characterization of individual polynucleotide molecules using a membrane channel. *Proc. Natl. Acad. Sci. U.S.A.*, 93(24):13770–13773, 1996.

[149] KR Mahendran, C Chimerel, T MacH, and M Winterhalter. Antibiotic translocation through membrane channels: Temperature-dependent ion current fluctuation for catching the fast events. *Eur. Biophys. J.*, 38(8):1141–1145, 2009.

[150] I Kosztin and K Schulten. Fluctuation-Driven Molecular Transport Through an Asymmetric Membrane Channel. *Phys. Rev. Lett.*, 93(23):238102, 2004.

[151] S Pagliara, C Schwall, and UF Keyser. Optimizing diffusive transport through a synthetic membrane channel. *Adv. Mater.*, 25(6):844–849, 2013.

[152] JM Huguet, CV Bizarro, N Forns, SB Smith, C Bustamante, and F Ritort. Single-molecule derivation of salt dependent base-pair free energies in DNA. *Proc. Natl. Acad. Sci. U.S.A.*, 107(35):15431–15436, 2010.

[153] H Qian and XS Xie. Generalized Haldane equation and fluctuation theorem in the steady-state cycle kinetics of single enzymes. *Phys. Rev. E*, 74(1):010902, 2006.

[154] LP Faucheux, LS Bourdieu, PD Kaplan, and AJ Libchaber. Optical Thermal Ratchet. *Phys. Rev. Lett.*, 74(9):1504–1507, 1995.

[155] C Battle, CP Broedersz, N Fakhri, VF Geyer, J Howard, CF Schmidt, and FC MacKintosh. Broken detailed balance at mesoscopic scales in active biological systems. *Science*, 352(6285):604–607, 2016.

[156] S Pagliara, SL Dettmer, and UF Keyser. Channel-facilitated diffusion boosted by particle binding at the channel entrance. *Phys. Rev. Lett.*, 113(4):048102, 2014.

[157] S Tottori, K Misiunas, UF Keyser, and DJ Bonthuis. Nonlinear electrophoresis of highly charged nonpolarizable particles. *Physical review letters*, 123(1):014502, 2019.

[158] AP Manuel, J Lambert, and MT Woodside. Reconstructing folding energy landscapes from splitting probability analysis of single-molecule trajectories. *Proc. Natl. Acad. Sci. U.S.A.*, 112(23):7183–7188, 2015.

[159] JB Masson, D Casanova, S Türkcan, G Voisinne, MR Popoff, M Vergassola, and A Alexandrou. Inferring Maps of Forces inside Cell Membrane Microdomains. *Phys. Rev. Lett.*, 102(4):048103, 2009.

[160] Jr Massey and J Frank. The kolmogorov-smirnov test for goodness of fit. *J. Am. Stat. Assoc.*, 46(253):68–78, 1951.

[161] M Chupeau, J Gladrow, A Chepelianskii, UF Keyser, and E Trizac. Optimizing brownian escape rates by potential shaping. *Proc. Natl. Acad.Sci. U.S.A.*, 2019.

[162] C Wagner and T Kiefhaber. Intermediates can accelerate protein folding. *Proc. Natl. Acad. Sci. U.S.A.*, 96(12):6716–6721, 1999.

[163] J Kubelka, J Hofrichter, and WA Eaton. The protein folding 'speed limit'. *Curr. Opin. Struct. Biol.*, 14(1):76–88, 2004.

[164] VV Palyulin and R Metzler. How a finite potential barrier decreases the mean first-passage time. *J. Stat. Mech. Theory Exp.*, 2012(03):L03001, 2012.

[165] VV Palyulin and R Metzler. Speeding up the first-passage for subdiffusion by introducing a finite potential barrier. *J. Phys. A Math. Theor.*, 47(3):032002, 2013.

[166] R Zwanzig. Diffusion past an entropy barrier. *J. Phys. Chem.*, 96(10):3926–3930, 1992.

[167] X Yang, C Liu, Y Li, F Marchesoni, P Hänggi, and HP Zhang. Hydrodynamic and entropic effects on colloidal diffusion in corrugated channels. *Proc. Natl. Acad. Sci. U.S.A.*, 114(36):9564–9569, 2017.

[168] R Zwanzig. Effective Diffusion coefficient for a Brownian particle in a two-dimensional periodic channel. *Physica*, 117A:277–280, 1983.

[169] D Reguera, G Schmid, PS Burada, JM Rubí, P Reimann, and P Hänggi. Entropic Transport: Kinetics, Scaling, and Control Mechanisms. *Phys. Rev. Lett.*, 96(13):130603, 2006.

[170] AM Berezhkovskii, MA Pustovoit, and SM Bezrukov. Diffusion in a tube of varying cross section: Numerical study of reduction to effective one-dimensional description. *J. Chem. Phys.*, 126(13):134706, 2007.

[171] G Schmid, PS Burada, P Talkner, and P Hänggi. Rectification through entropic barriers. In *Adv. Solid State Phys.*, volume 48, pages 317–328. 2009.

[172] JM Romero, O González-Gaxiola, and G Chacón-Acosta. Exact solution to Fick-Jacobs equation. *arXiv preprint, arXiv:1108.1164*, 2011.

[173] L Dagdug, AM Berezhkovskii, YA Makhnovskii, VY Zitserman, and SM Bezrukov. Force-dependent mobility and entropic rectification in tubes of periodically varying geometry. *J. Chem. Phys.*, 136(21):214110, 2012.

[174] D Reguera, A Luque, PS Burada, G Schmid, J. M. Rubí, and P. Hänggi. Entropic splitter for particle separation. *Phys. Rev. Lett.*, 108(2):1–5, 2012.

[175] T Motz, G Schmid, P Hänggi, D Reguera, and JM Rubí. Optimizing the performance of the entropic splitter for particle separation. *J. Chem. Phys.*, 141(7):074104, 2014.

[176] D Reguera and JM Rubí. Kinetic equations for diffusion in the presence of entropic barriers. *Phys. Rev. E*, 64(6):061106, 2001.

[177] AWC Lau and TC Lubensky. State-dependent diffusion: Thermodynamic consistency and its path integral formulation. *Phys. Rev. E*, 76(1), 2007.

[178] G Volpe, L Helden, T Brettschneider, J Wehr, and C Bechinger. Influence of noise on force measurements. *Phys. Rev. Lett.*, 104(17):1–4, 2010.

[179] M Yang and M Ripoll. Brownian motion in inhomogeneous suspensions. *Phys. Rev. E*, 87(6):062110, 2013.

[180] O Farago and N Grønbech-Jensen. Langevin dynamics in inhomogeneous media: Re-examining the Itô-Stratonovich dilemma. *Phys. Rev. E*, 89(1):013301, 2014.

[181] U Winter and T Geyer. Coarse grained simulations of a small peptide: Effects of finite damping and hydrodynamic interactions. *J. Chem. Phys.*, 131(10):104102, 2009.

[182] T Geyer and U Winter. An O(N2) approximation for hydrodynamic interactions in Brownian dynamics simulations. *J. Chem. Phys.*, 130(11):114905, 2009.

[183] T Hohage. *Inverse Problems: Regularization theory, Algorithms and Applications*. Georg-August University of Göttingen, Goettingen, 2014.

[184] J Gladrow. https://github.com/JamesGlare/Neural-Net-LabView-DLL. 2019.

[185] Y Zhang, AM Saxe, MS Advani, and AA Lee. Energy–entropy competition and the effectiveness of stochastic gradient descent in machine learning. *Mol. Phys.*, 116(21-22):3214–3223, 2018.

[186] CM Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[187] Y LeCun, B Boser, JS Denker, D Henderson, RE Howard, W Hubbard, and LD Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Comput*, 1(4):541–551, 1989.

[188] Y LeCun, L Bottou, Y Bengio, and P Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2323, 1998.

[189] A Krizhevsky, I Sutskever, and GE Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *ImageNet Classif. with Deep Convolutional Neural Networks*, pages 1097–1105, 2012.

[190] C Szegedy, V Vanhoucke, S Ioffe, J Shlens, and Z Wojna. Rethinking the Inception Architecture for Computer Vision. In *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pages 2818–2826. 2016.

[191] K He, X Zhang, S Ren, and J Sun. Deep Residual Learning for Image Recognition. In *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, pages 770–778. 2016.

[192] C Szegedy, S Ioffe, V Vanhoucke, and A Alemi. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In *Thirty-First AAAI Conf. Artif. Intell.* 2017.

[193] C Liu, B Zoph, M Neumann, J Shlens, W Hua, LJ Li, L Fei-Fei, A Yuille, J Huang, and K Murphy. Progressive Neural Architecture Search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 19–34, 2018.

[194] R Collobert, J Weston, L Bottou, M Karlen, K Kavukcuoglu, and P Kuksa. Natural Language Processing (almost) from Scratch. *Journal of machine learning research*, 12(Aug):2493–2537, 2011.

[195] J Devlin, MW Chang, K Lee, and K Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint, arXiv:1810.04805*, 2018.

[196] A Radford, J Wu, R Child, D Luan, D Amodei, and I Sutskever. Language Models are Unsupervised Multitask Learners. *Open AI*, 2019.

[197] Y LeCun, Y Bengio, and G Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[198] I Goodfellow, J Pouget-Abadie, M Mirza, B Xu, D Warde-Farley, S Ozair, A Courville, and Y Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.

[199] K Sohn, H Lee, and X Yan. Learning Structured Output Representation using Deep Conditional Generative Models. *Adv. Neural Inf. Process. Syst.*, pages 3483–3491, 2015.

[200] B Sanchez-Lengeling and A Aspuru-Guzik. Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365, 2018.

[201] J Adler and O Öktem. Deep Bayesian Inversion. *arXiv preprint, arXiv:1811.05910*, 2018.

[202] T Karras, S Laine, and T Aila. A Style-Based Generator Architecture for Generative Adversarial Networks. *arXiv preprint, arXiv:1812.04948*, 2018.

[203] A Brock, J Donahue, and K Simonyan. Large Scale GAN Training for High Fidelity Natural Image Synthesis. *arXiv preprint, arXiv:1809.11096*, pages 1–35, 2018.

[204] S Nowozin, B Cseke, and R Tomioka. f-GAN: Training Generative Neural Samplers using Variational Divergence Minimization. 2(1):1–9, 2016.

[205] M Mirza and S Osindero. Conditional Generative Adversarial Nets. *arXiv preprint, arXiv:1411.1784*, 2014.

[206] T Salimans, I Goodfellow, W Zaremba, V Cheung, A Radford, and X Chen. Improved techniques for training gans. In *Advances in neural information processing systems*, pages 2234–2242, 2016.

[207] A Srivastava, L Valkov, C Russell, MU Gutmann, and C Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. *Adv. Neural Inf. Process. Syst.*, pages 3308–3318, 2017.

[208] T Miyato, T Kataoka, M Koyama, and Y Yoshida. Spectral Normalization for Generative Adversarial Networks. *arXiv preprint, arXiv:1802.05957*, 2018.

[209] S Lee, J Ha, and G Kim. Harmonizing Maximum Likelihood with GANs for Multimodal Conditional Generation. *arXiv preprint, arXiv:1902.09225*, 2019.

[210] M Arjovsky, S Chintala, and L Bottou. Wasserstein GAN. *arXiv preprint, arXiv:1701.07875*, 2017.

[211] I Gulrajani, F Ahmed, M Arjovsky, V Dumoulin, and A Courville. Improved Training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, 2017.

[212] DP Kingma and M Welling. Auto-Encoding Variational Bayes. *ICLR 2014*, 2013.

[213] D P Kingma and J Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint, arXiv:1412.6980*, 2014.

[214] RA. Jacobs, MI Jordan, SJ Nowlan, and GE Hinton. Adaptive Mixtures of Local Experts. *Neural Comput.*, 3(1):79–87, 1991.

[215] CM Bishop. Mixture density networks. 1994.

[216] A Kristiadi and A Fischer. Predictive Uncertainty Quantification with Compound Density Networks. *arXiv preprint, arXiv:1902.01080*, 2019.

[217] S Ioffe and C Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint, arXiv:1502.03167*, 2015.

[218] J Lei Ba, JR Kiros, and GE Hinton. Layer normalization. *arXiv preprint, arXiv:1607.06450*, 2016.

[219] N Srivastava, G Hinton, A Krizhevsky, I Sutskever, and R Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[220] J Gladrow. Digital phase-only holography using deep conditional generative models. *arXiv preprint arXiv:1911.00904*, 2019.

[221] K He, X Zhang, S Ren, and J Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[222] T Auger, J Mathé, V Viasnoff, G Charron, JM Di Meglio, L Auvray, and F Montel. Zero-mode waveguide detection of flow-driven DNA translocation through nanopores. *Phys. Rev. Lett.*, 113(2):1–5, 2014.

[223] PWK Rothemund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.

[224] EA Hemmig, C Fitzgerald, C Maffeo, L Hecker, SE Ochmann, A Aksimentiev, P Tinnefeld, and UF Keyser. Optical Voltage Sensing Using DNA Origami. *Nano Lett.*, 18(3):1962–1971, 2018.

[225] M Muthukumar. Translocation of a confined polymer through a hole. *Phys. Rev. Lett.*, 86(14):3188–3191, 2001.

[226] M Muthukumar. Polymer escape through a nanopore. *The Journal of chemical physics*, 118(11):5174–5184, 2003.

[227] UF Keyser, BN Koeleman, S van Dorp, D Krapf, RMM Smeets, SG Lemay, NH Dekker, and C Dekker. Direct force measurements on DNA in a solid-state nanopore. *Nat. Phys.*, 2(7):473–477, 2006.

[228] M Muthukumar. Theory of capture rate in polymer translocation. *J. Chem. Phys.*, 132(19), 2010.

[229] JH Masliyah and S Bhattacharjee. *Electrokinetic and colloid transport phenomena*. John Wiley & Sons, 2006.

[230] SH Behrens and DG Grier. The charge of glass and silica surfaces. *J. Chem. Phys.*, 115(14):6716–6721, 2001.

[231] VV Thacker, S Ghosal, S Hernández-Ainsa, NAW Bell, and UF Keyser. Studying DNA translocation in nanocapillaries using single molecule fluorescence. *Appl. Phys. Lett.*, 101(22):223704, 2012.

[232] RL Fleischer, PB Price, and EM Symes. Novel filter for biological materials. *Science*, 143(3603):249–250, 1964.

[233] Whatman Nuclepore Membrane Specifications. Technical report, Ltd, Whatman International.

[234] EA Hemmig. *DNA Origami Structures for Artificial Light-Harvesting and Optical Voltage Sensing*. PhD thesis, University of Cambridge, 2017.

[235] N Ermann. *Controlling the translocation process of double-stranded DNA in carrier-based nanopore sensing* . PhD thesis, 2017.

[236] DP Hoogerheide, B Lu, and JA Golovchenko. Pressure–voltage trap for dna near a solid-state nanopore. *ACS nano*, 8(7):7384–7391, 2014.

[237] MJ Levene. Zero-Mode Waveguides for Single-Molecule Analysis at High Concentrations. *Science*, (5607):682–686.

[238] J Liphardt. Equilibrium Information from Nonequilibrium Measurements in an Experimental Test of Jarzynski's Equality. *Science*, 296(5574):1832–1835, 2002.

[239] D Collin, F Ritort, C Jarzynski, SB Smith, I Tinoco, and C Bustamante. Verification of the Crooks fluctuation theorem and recovery of RNA folding free energies. *Nature*, 437(7056):231–234, 2005.

[240] MØ Jensen, S Park, E Tajkhorshid, and K Schulten. Energetics of glycerol conduction through aquaglyceroporin glpf. *Proc. Natl. Acad. Sci. U.S.A.*, 99(10):6731–6736, 2002.

[241] L Ardizzone, J Kruse, S Wirkert, D Rahner, EW Pellegrini, RS Klessen, L Maier-Hein, C Rother, and U Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint, arXiv:1808.04730*, 2018.

[242] Z Xiao, Q Yan, and Y Amit. A method to model conditional distributions with normalizing flows. *arXiv preprint arXiv:1911.02052*, 2019.

[243] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13(Feb):281–305, 2012.

[244] C Thornton, F Hutter, HH Hoos, and K Leyton-Brown. Auto-weka: Combined selection and hyperparameter optimization of classification algorithms. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 847–855. ACM, 2013.