

Università degli Studi di Torino
Scuola di Dottorato



UNIVERSITÀ
DI TORINO

**Neural and Non-Equilibrium methods for numerical
Effective String Theory calculations**

Elia Cellini

Università degli Studi di Torino
Scuola di Dottorato

Dottorato in Fisica

**Neural and Non-Equilibrium methods for
numerical Effective String Theory
calculations**

Elia Cellini

Tutor: Michele Caselle

"I haven't explored it yet, but I think I should, because it intrigues me. What fascinates me is that AI can truly extend creativity: you have a vision, and AI expands it – and then you go, 'wow!'."

Gigi D'Agostino, (2025).

Abstract

Effective String Theory (EST) is a powerful non-perturbative approach used to describe confinement in Yang-Mills theory by modeling the confining flux tube as a thin, vibrating string. EST calculations are typically performed using zeta-function regularization; however, some observables, such as the shape of the flux tube, are too complex to be addressed with this method.

This issue can, in principle, be approached by modeling the EST action as a spin system on a discretized two-dimensional lattice, making it amenable to Monte Carlo simulations. However, the strong non-linearity of the model makes conventional algorithms highly inefficient for such computations. Instead, the problem is well-suited for two distinct numerical approaches based on deep learning methods and out-of-equilibrium statistical mechanics. In the first approach, a variational density is learned using a class of deep generative models called Normalizing Flows (NFs), allowing unbiased expectation values over the EST distribution to be computed via a reweighting procedure. In contrast, out-of-equilibrium methods, such as Non-Equilibrium Markov Chain Monte Carlo (NE-MCMC) algorithms, enable exact calculations of vacuum expectation values using Crooks' theorem and Jarzynski's equality. Notably, these two approaches can be combined into an efficient sampling method known as Stochastic Normalizing Flows (SNFs).

This thesis introduces EST and its lattice regularization, followed by a brief overview of standard MCMC methods. Next, we

introduce NFs and present a proof-of-concept using Continuous NFs to demonstrate the feasibility of our method for sampling EST on the lattice. We then introduce non-equilibrium methods and SNFs, showcasing state-of-the-art results achieved with the latter. The combination of these findings enables a quantitative description of the fine details of the confinement mechanism in different lattice gauge theories. The thesis concludes by discussing the main implications of our results for EST and gauge theory, as well as potential directions for future research.

Contents

1	Introduction	1
2	Effective String Theory on the lattice	9
2.1	Nambu-Gotō effective string	10
2.1.1	Partition function	12
2.1.2	Width	16
2.1.3	The Nambu-Gotō string as a $T\bar{T}$ perturbation	17
2.2	Lattice Nambu-Gotō theory	19
2.2.1	Finite size CFT and non-universal terms	22
2.2.2	Width and Shape	23
2.3	Higher-order terms beyond the Nambu-Gotō action	24
2.4	Lattice gauge theory in the EST picture	26
3	Markov Chain Monte Carlo	30
3.1	Hybrid Monte Carlo	33
3.1.1	Leap-frog integrator	36
4	Neural Generative Samplers	37
4.1	Normalizing Flows	38
4.1.1	Coupling Layers	39
4.1.2	Continuous Normalizing Flows	40
4.2	Training	41
4.3	Estimators	43
4.4	Normalizing Flows in lattice QFT simulations	45

5	A Proof-of-Concept: simulating EST with Continuous Normalizing flows	47
5.1	Architecture	47
5.2	Simulation setup	48
5.3	Partition Function	50
5.3.1	Low-temperature regime ($L \gg R$)	51
5.3.2	High-temperature regime ($R \gg L$)	52
5.4	Width of the flux tube	55
5.4.1	Low-temperature regime ($L \gg R$)	56
5.4.2	High-temperature regime ($R \gg L$)	57
5.5	Comparison with Hybrid Monte Carlo	60
5.6	Summary	62
6	Non-equilibrium methods and Stochastic Normalizing Flows	64
6.1	Non-Equilibrium Markov Chain Monte Carlo	66
6.2	Jarzynski's equality and Non-Equilibrium iMH	68
6.2.1	Jarzynski's equality in lattice QFT simulations	74
6.3	Stochastic Normalizing Flows	75
6.3.1	SNFs in lattice QFT	77
7	EST results with Stochastic Normalizing Flows	78
7.1	Architecture and simulation details	78
7.2	Comparison between SNFs, NFs, and NE-MCMC	80
7.3	Nambu-Gotō: partition function	81
7.4	Nambu-Gotō: string width	84
7.5	Beyond Nambu-Gotō: the string width in presence of the \mathcal{K}^2 and \mathcal{K}^4 terms	85
7.6	Binder Cumulant	90
7.6.1	The Nambu-Gotō action	90
7.6.2	The Binder cumulant in presence of the \mathcal{K}^2 or the \mathcal{K}^4 terms	92

8	Conclusions and final remarks	95
A	Appendix	120
A.1	Free boson analytical solution	120
A.2	Exact sampling of the Free Boson	124
A.3	Fits of numerical results	126

Chapter 1

Introduction

Quantum Chromodynamics (QCD) is the fundamental theory describing the strong nuclear force and plays a crucial role in the Standard Model of particle physics. As a quantum gauge theory with local SU(3) symmetry, QCD governs the interactions of quarks and gluons, which carry color charge. However, in low energy regimes, these particles are never observed in isolation; instead, they are permanently bound within color-neutral hadrons, such as protons, neutrons, and mesons. This phenomenon, known as *color confinement* [1], is one of the most intriguing aspects of modern physics.

A key feature of QCD is asymptotic freedom [2, 3], meaning that at very short distances, the strong interaction becomes weaker, allowing quarks to behave almost as free particles. However, at larger distances, the interaction strength increases, leading to complex non-perturbative effects such as confinement and chiral symmetry breaking. The force between quarks grows with the separation R , following, at the leading order, a linear potential:

$$V(R) \approx \sigma R,$$

where σ represents the string tension. Unlike the Coulomb-like force in electromagnetism, the QCD force does not diminish at long distances. Instead, the energy stored between quarks is

concentrated into color flux tubes, also known as QCD strings. These tubes are thin vibrating structures of gluonic field energy that form between quarks and act as confining bonds. As a quark-antiquark pair is pulled apart, the flux tube stretches, and beyond a certain point, the energy is sufficient to create a new quark-antiquark pair rather than allowing the original pair to separate. This mechanism explains why quarks are never observed individually but always within hadrons. Quantum Field Theory (QFT) simulations on the lattice numerically depict the formation of flux tubes and the confinement of color-charged particles. However, a complete analytical proof of confinement from first principles remains one of the biggest open challenges. Understanding confinement and the behavior of flux tubes is thus crucial for advancing non-perturbative QCD and exploring the structure of the QCD vacuum.

A powerful approach to model confinement in non-abelian gauge theory is represented by Effective String Theory (EST), a non-perturbative framework that models the confining flux tube in terms of vibrating strings. In $D \neq 26$, where D are the space-time dimensions of the target Lattice Gauge Theory (LGT), the EST (at least in its simplest formulation) is anomalous at the quantum level and thus must be considered only as an effective, large-distance description of the flux tube of Yang-Mills theories. Notwithstanding this, precise Monte Carlo simulations of several different LGTs proved that it is indeed a highly predictive effective model (for recent reviews see for instance [4–6]).

The reason of this success is in the so called "low energy universality" [4, 7–13] which states that, due to the symmetry constraints imposed by the Poincaré invariance in the target space, the first few terms of the EST large-distance expansion are universal and coincide with those of the Nambu-Gotō action [14, 15]. As a consequence, the EST turns out to be much more predictive than typical effective models and its predictions depend

only on one free parameter: the string tension σ of the Nambu-Gotō model.

It is thus clear that a central role in this game is played by the Nambu-Gotō action. Its predictions can be immediately compared with the results of Monte Carlo simulations of essentially any confining LGT (two exceptions are the three dimensional U(1) model and trace deformed LGT). Thus, a precise understanding of the Nambu-Gotō physical properties could improve our knowledge of the confining regime of Yang-Mills theories.

A major progress in this context was the recent observation that the Nambu-Gotō model is actually an exactly integrable, irrelevant, perturbation of the two-dimensional free Conformal Field Theory (CFT) [11] of $D - 2$ free bosons which represent the transverse degrees of freedom of the string. This irrelevant perturbation is driven by the $T\bar{T}$ operator of the $D - 2$ free bosons [16–18]. Over the past few years, the EST community has begun to explore terms beyond the Nambu-Gotō action, one of the most intriguing and challenging open questions in the EST modeling [19–25]. One of the main tools to address this issue is the comparison between the results of high-precision LGT simulations and effective string predictions; however, very few EST results are known exactly.

We shall see below that several results are analytically known for the partition function of the Nambu-Gotō action, which can be explicitly calculated for essentially all the geometries which are relevant for the comparison with LGT results. On the contrary, much less is known on the correlation functions, and in particular on the one that measures the density of the chromo-electric flux in presence of quark sources: this quantity has an important meaning in LGT, since it allows to study the shape of the flux tube. In principle, one could address this problem by treating the Nambu-Gotō action as a spin model and regularizing it on a two-dimensional lattice, where one could

perform Monte Carlo simulations. However, due to the strong non-linearity of the model, it turns out that performing such simulations using standard algorithms is highly inefficient [26]. The problem is instead perfectly suited for a completely different numerical approach, based on deep learning architectures and non-equilibrium statistical mechanics. In this respect it is important to stress that the spin model that we obtain discretizing the Nambu-Gotō action is critical for the whole set of values of the string tension. Its continuum limit is in fact, as mentioned above, the theory of a free bosonic massless degree of freedom perturbed by the $T\bar{T}$ operator. This makes this model a perfect laboratory to test different algorithms in a particularly challenging setting.

In recent years, the exponential growth of machine learning applications has stimulated the introduction of novel methods for numerical studies of quantum field theory [27–37]. In particular, deep learning approaches [38], a branch of machine learning originally investigated in the field of representation learning [39] that strongly rely on deep neural networks, have shifted the computational paradigms in fields like computer vision [40], game playing [41], protein folding [42], and natural language processing [43]. An intriguing factor that contributed to this revolution is the emergence of phenomena such as the double descent [44, 45] or meta-learning [46, 47], which arise when neural networks operate in the so-called over-parameterized regime.

Among the various applications of deep neural networks, a recently popular approach is represented by deep generative algorithms capable of generating new data with surprisingly high quality. A remarkable example is provided by Generative Pre-trained Transformers [43, 48], a class of large language models that is significantly impacting society at different levels. Another important example is given by diffusion models [49, 50], a popular class of deep generative models that has recently been

applied to lattice field theory as well [51–54], with roots deeply connected to non-equilibrium statistical mechanics.

In theoretical physics, two classes of algorithms that have recently emerged as promising tools to overcome the problems of standard sampling methods are Variational Autoregressive Networks (VANs) [55] and Normalizing Flows (NFs) [56–58]. Both approaches are deep generative methods that provide access to the exact density of the generated samples. This feature, in combination with variational techniques, can be leveraged to obtain unbiased estimations of observables and partition functions of a target Boltzmann distribution [59–61]. Furthermore, these Neural Generative Samplers can be used to tackle problems such as *critical slowing down* [60, 62], thus overcoming the main drawbacks of standard Markov Chain Monte Carlo approaches.

VANs are popular in the statistical mechanics community since they are practical for applications involving discrete degrees of freedom [55, 63–65]. NFs cannot be used to study discrete domains, but, they are particularly well-suited and highly promising when applied to continuous systems such as lattice scalar field theories [26, 60, 61, 66–71], LGT [72–78] and QCD [35, 79, 80]. While these variational deep learning techniques have provided great results in small systems, they struggle to scale to state-of-the-art lattice studies [65, 66, 76, 81, 82]. Notwithstanding these challenges, a promising route to scaling up these approaches is by combining them with other computational frameworks to enhance their scalability. In the case of VANs, the state-of-the-art in terms of lattice volumes is the so-called Renormalization informed Generative Critical Sampler (RiGCS) [65], a hierarchical VAN whose design is inspired by the renormalization group. NFs, on the other hand, can be combined with non-equilibrium methods based on out-of-equilibrium statistical mechanics [83–85] to obtain Stochastic Normalizing Flows [86, 87]. Non-equilibrium methods are a stochastic approach with a

framework similar to NFs, widely applied in lattice field theory to compute free energies, and have recently been used to tackle critical slowing down and topological freezing as well [85, 88–93]. The most appealing feature of this class of algorithms is their scalability, ensured by the non-equilibrium picture that characterizes these methods, and, when combined with NFs, the resulting SNFs inherit these properties, thus providing a scalable hybrid flow-based sampler [78].

As we shall show in further sections, both neural and non-equilibrium methods address the problems of standard methods by rethinking how expectation values over a target ensemble are computed, using a variational and a stochastic approach, respectively.

In this thesis, mostly based on refs. [26, 78, 87, 94], we used a NF-based sampler approach to study the partition function, the shape and the width of the flux tube of a lattice-regularized version of the EST model. Thanks to the exact knowledge of the partition function of the Nambu-Gotō theory we can benchmark the correctness of this novel numerical approach and then use the results to obtain information on the shape of the flux tube. This rather unusual approach to EST calculations is interesting for several reasons that we discuss in the following.

- It allows to study large values of the string tension σ which are not accessible in LGT simulations; they are of great interest since they correspond to the perturbative regime of the $T\bar{T}$ perturbation which is controlled by $1/\sigma$ [17, 18].
- On the LGT side, it allows for a precise study of the flux-tube shape predicted by the EST model, which can be used as a benchmark to compare results from LGT simulations. In this context, for instance, we will confirm the predictions of [95–97] and the conjecture by [98].
- It is the first step toward an analytic study of the correc-

tions beyond Nambu-Gotō, which have been recently the subject of great interest in the EST community [19–24].

- It allows for the first studies of the width and shape of the higher-order corrections to the Nambu-Gotō EST.
- It allows to study in detail the role of lattice artifacts in the Nambu-Gotō predictions for LGTs.
- Last, but not least, this is the first numerical study of the lattice realization of a $T\bar{T}$ perturbed model and could provide a tool to test in a non-trivial setting various predictions obtained in the last few years on this class of models.

Among the research activities conducted during this Ph.D., two notable works, in which the author was involved, are briefly mentioned in this thesis: the RiGCS VAN architecture [65] and a study on entanglement entropy using flow-based samplers in scalar field theory [71]. The former has been discussed briefly earlier in this chapter, while the latter will be mentioned in Section 4.4. These works are not discussed in more detail, as they do not constitute the primary focus of this thesis. Additionally, although the theoretical aspects of non-equilibrium methods presented in Chapter 6 are primarily based on refs. [78, 87], the corresponding numerical results are not included here, as they specifically concern lattice scalar ϕ^4 and lattice SU(3) gauge theory, respectively.

This manuscript is organized as follows. In Chapter 2, we briefly outline EST, its lattice regularization, and the known analytical results. In Chapter 3, we introduce standard Markov Chain Monte Carlo methods and discuss their limitations, in particular in this setting. In Chapter 4, we introduce the Neural Generative Samplers approach and Normalizing Flows, while in Chapter 5, we present numerical results for a proof-of-concept

application of flow-based samplers to the study of EST. In Chapter 6, we introduce non-equilibrium methods and Stochastic Normalizing Flows, while in Chapter 7, we present numerical results obtained with the latter architecture, with particular emphasis on the width and shape of the effective string. The last chapter is devoted to some concluding remarks.

Chapter 2

Effective String Theory on the lattice

In the last few decades Effective String Theory (EST) has established itself as a very powerful and predictive approach for the understanding and modeling of the non-perturbative behavior of confining Yang-Mills theories. In this framework, the confining flux tube that connects a quark-antiquark pair is represented as a thin vibrating string [14, 15, 99–101]. In particular, the correlator between two Polyakov loops is related to the sum over all the surfaces bordered by the two Polyakov loops weighted by the EST action:

$$\langle P(0)P^\dagger(R) \rangle \sim \int DX e^{-S_{\text{EST}}[X]} \equiv Z$$

where R denotes the distance between the two Polyakov loops.

The simplest choice for S_{EST} is the well known Nambu-Gotō action: this theory is known to be anomalous at the quantum level and in fact it should be considered only a large-distance approximation of the actual (anomaly-free) EST. Despite this, the Nambu-Gotō action, whose partition function can be solved exactly, has been shown in the last few years to describe to high-precision the interquark potential in several different LGTs [4–6]. As we commented in the introduction, the reason of this

effectiveness lies in the so-called "low energy universality": even if the Nambu-Gotō action is only a first order approximation of the actual effective string describing the confining flux tube, the first few terms of any EST must coincide with those of the NG action. Exploring the terms beyond the Nambu-Gotō action is one of the most interesting and challenging open problem in the EST approach [19–25].

Besides the exact solution of the Nambu-Gotō partition function and a few one-loop calculations of other observables, one has to resort to approximate estimates or to educated conjectures, since analytical calculations are often too challenging. A well-known example is the (quantum) width of the effective string, whose counterpart in gauge theories is a correlation function that measures the thickness of the confining flux tube: for this observable the only available analytical solution is a perturbative calculation up to the second order, see refs. [95–97]. The situation is even worse if one is interested in terms in the EST action beyond the NG term, for which even the partition function (i.e. the interquark potential measured in LGT simulations) is not known.

2.1 Nambu-Gotō effective string

The simplest choice for S_{EST} fulfilling the constraints imposed by the Lorentz invariance in the target space is the Nambu-Gotō action [14, 15] that is defined as follows:

$$S_{\text{NG}} = \sigma \int_{\Sigma} d^2\xi \sqrt{g}, \quad (2.1)$$

where $g \equiv \det g_{\alpha\beta}$ and

$$g_{\alpha\beta} = \partial_{\alpha} X_{\mu} \partial_{\beta} X^{\mu} \quad (2.2)$$

is the metric induced on the reference surface Σ by the mapping $X_{\mu}(\xi)$, called *worldsheet*, and $\xi \equiv (\xi^0, \xi^1)$ denote the worldsheet




Figure 2.1: (a) Schematic representation of a flux tube, where the red and blue points represent the quark and anti-quark, respectively. (b) Schematic representation of the worldsheet $X(\xi^0, \xi^1)$, under the physical gauge in $D = 3$. As the flux tube evolves in time, it traces out a surface, which is modeled by X .

coordinates. This term has a simple geometric interpretation: it measures the area of the surface spanned by the string in the target space and has only one free parameter: the string tension σ .

S_{NG} is manifestly reparametrization-invariant and the first step is to fix this invariance. A typical choice is the so-called “physical gauge”. In this gauge the two worldsheet coordinates are identified with the longitudinal degrees of freedom of the string: $\xi^0 = X^0$, $\xi^1 = X^1$, so that the string action can be expressed as a function only of the $(D - 2)$ degrees of free-

dom corresponding to the transverse displacements, X^i , with $i = 2, \dots, (D - 1)$ which are assumed to be single-valued functions of the worldsheet coordinates. In the following, for simplicity, we assume $D = 3$ so as to have only one transverse direction. We thus eliminate the index i and simply denote as $X(\xi)$ the remaining degrees of freedom. See fig. 2.1 for a schematic representation of the flux tube and the effective string, in $D = 3$, under the physical gauge.

It is well known that this gauge fixing is anomalous at the quantum level, as it cannot account for surfaces with handles or other complex topological structures. Consequently, the resulting EST action should be regarded only as a long-distance approximation of the true string action. This is the reason why we define this approach as an "effective" string description of confinement. With a suitable redefinition of the fields, $X = \phi/\sqrt{\sigma}$, we can write explicitly this large distance expansion as follows

$$S_{\text{NG}} = \sigma RL + \frac{1}{2} \int d^2\xi \left[\partial_\alpha \phi \cdot \partial^\alpha \phi - \frac{1}{8\sigma RL} (\partial_\alpha \phi \cdot \partial^\alpha \phi)^2 + \dots \right]. \quad (2.3)$$

where R denotes, as above, the distance between the two Polyakov loops and L their length. The first term of this expansion is exactly the gaussian action, i.e. a two dimensional Conformal Field Theory (CFT) of a free bosonic field which represents the only remaining transverse degree of freedom of the string in $D = 3$. Remarkably enough, all the remaining terms of the expansion combine themselves to give an exactly integrable, irrelevant perturbation of this CFT [11], driven by the $T\bar{T}$ operator of the $D - 2$ free bosons [16–18].

2.1.1 Partition function

Using the zeta function regularization it is possible to evaluate exactly the partition function of the Nambu-Gotō theory to all

orders; the result is [8, 102]

$$\langle P(x)P^\dagger(x+R) \rangle \sim \sqrt{\sigma}L \sum_{n=0}^{\infty} w_n K_0(E_n R) \quad (2.4)$$

where K_0 is the modified Bessel function of order 0, w_n is the multiplicity of the closed string states which propagate from one Polyakov loop to the other, and E_n their energies:

$$E_n(L) = \sigma L \sqrt{1 + \frac{8\pi}{\sigma L^2} \left[n - \frac{1}{24} \right]}. \quad (2.5)$$

From the Polyakov loop correlator it is possible to obtain the interquark potential, which is defined as follows

$$V(R, L) = -\frac{1}{L} \log \langle P(x)P^\dagger(x+R) \rangle \quad (2.6)$$

and in the large- R limit we have, as expected, a linearly rising potential whose slope is controlled by the ground state energy of the EST

$$V(R, L) = \frac{E_0(L)}{L} R = \sigma(L) R \quad (2.7)$$

where

$$\sigma(L) = \sigma \sqrt{1 - \frac{\pi}{3\sigma L^2}} \quad (2.8)$$

encodes the dependence of the string tension on the physical temperature. Thus, the logarithm of the partition function Z is found to be:

$$-\log Z = \frac{1}{2} \log \frac{R}{L} + \sigma(L) R L + C \quad (2.9)$$

where the constant C has no further dependence on R or L and the term $\frac{1}{2} \log \frac{R}{L}$ comes from the open-closed string duality that we shall discuss shortly.

In the following, for our comparisons we will perform an expansion in powers of $\frac{1}{\sigma R L}$ of the action. Much earlier than

eq. (2.4), the first few terms of this expansion were derived by Dietz and Filk [103], by treating the next-to-leading terms which appear in eq. (2.3) as small perturbations of the gaussian free action. In particular the first term is the well-known partition function of the free bosonic $c = 1$ CFT

$$\langle P(x)P^\dagger(x + R) \rangle \sim \frac{e^{-\sigma RL}}{\eta(\tau)} \quad (2.10)$$

where $\eta(\tau)$ denotes the Dedekind η function

$$\eta(\tau) = q^{\frac{1}{24}} \prod_{n=1}^{\infty} (1 - q^n) \quad ; \quad q = e^{2\pi i \tau} \quad ; \quad \tau = i \frac{L}{2R} \quad . \quad (2.11)$$

To understand the meaning of this result it is useful to expand it in the two limits $R \ll L$ and $R \gg L$, in which we can neglect the q^n terms in the infinite product of the Dedekind function. These limits correspond in the LGT language to the low- and high-temperature limits of the confining regime¹ while from a string point of view, in $D = 3$, they correspond to the open and closed string channels respectively:

$$V(R, L) = \sigma R - \frac{\pi}{24R} \quad \text{for } R \ll L \quad (2.12)$$

$$V(R, L) = \sigma R - \frac{\pi R}{6L^2} + \frac{1}{2L} \log \frac{2R}{L} \quad \text{for } R \gg L \quad (2.13)$$

where the term $\frac{1}{2L} \log \frac{2R}{L}$ comes from the modular transformation $\tau \rightarrow -\frac{1}{\tau}$ of the open-closed string duality. In particular, under this transformation, the Dedekind function changes as:

$$\eta(\tau) = \left(\frac{2R}{L} \right)^{\frac{1}{2}} \eta\left(-\frac{1}{\tau}\right) \quad (2.14)$$

¹To avoid confusion let us stress that with "high-temperature" we mean a region near the deconfinement transition, but still in the confining phase.

where the new term $\left(\frac{2R}{L}\right)^{\frac{1}{2}}$ is the one whose logarithm appears in eqs. (2.9) and (2.13).

The $-\frac{\pi}{24R}$ term in the first limit (eq. (2.12)) is the well known "Lüscher term" which has been widely studied in the LGT context. The next-to-leading term (i.e. the $1/\sigma RL$ correction) can be recast in the following combination of the Eisenstein functions E_2 and E_4 [103]:

$$\langle P(x)P^\dagger(x+R) \rangle \sim \frac{e^{-\sigma RL}}{\eta(\tau)} \left(1 + \frac{\pi^2 L}{1152\sigma R^3} [2E_4(\tau) - E_2^2(\tau)] \right) \quad (2.15)$$

where E_2 and E_4 can be expressed in power series

$$E_2(\tau) = 1 - 24 \sum_{n=1}^{\infty} \sigma(n)q^n \quad (2.16)$$

$$E_4(\tau) = 1 + 240 \sum_{n=1}^{\infty} \sigma_3(n)q^n \quad (2.17)$$

and $\sigma(n)$ and $\sigma_3(n)$ are, respectively, the sum of all divisors of n (including 1 and n) and the sum of their cubes.

For the range of values of σ , R and L that we shall study in this paper the terms proportional to q can be systematically neglected and we end up with the following expressions in the two limits:

$$V(R, L) = \sigma R - \frac{\pi}{24R} - \frac{\pi^2 L}{1152\sigma R^3} + \dots \quad \text{for } R \ll L \quad (2.18)$$

$$V(R, L) = \sigma R + \frac{1}{2L} \log \frac{2R}{L} - \frac{\pi R}{6L^2} - \frac{\pi^2 R}{72\sigma L^4} + \dots \quad \text{for } R \gg L \quad (2.19)$$

It is important to stress the role in these calculations of the zeta function regularization, which eliminates all the "bulk", divergent terms of the above partition functions.

2.1.2 Width

The second observable of interest in this work is the width of the string, given by

$$\sigma w^2(R) = \langle \phi^2(\xi_0, R/2) \rangle. \quad (2.20)$$

The only known analytical expression for the width in EST is a perturbative calculation up to the second order in $1/\sigma$, see refs. [95–97]. In the low-temperature regime $L \gg R$ we have that

$$\sigma w^2(\sigma, L, R) = \frac{1}{2\pi} \log \frac{R}{R_c} \left(1 - \frac{\pi}{4\sigma R^2} \right) + \frac{5}{96} \frac{1}{\sigma R^2} + \dots \quad (2.21)$$

where R_c is a new scale of the model, which emerges from the regularization of the correlator defining the string width. This scale sets a threshold below which the EST picture cannot be trusted anymore and in LGTs applications it is associated to the so called "intrinsic width" of the flux tube (see ref. [104] for a discussion of this issue). We will comment further on the role of this scale in our case later in this contribution.

In the high-temperature regime $R \gg L$ the expected behaviour is

$$\sigma w^2(\sigma, L, R) = \frac{1}{2\pi} \log \frac{L}{L_c} + \frac{R}{4L} + \frac{\pi}{24} \frac{R}{\sigma L^3} + \dots \quad (2.22)$$

where the new scale L_c has the same origin of R_c but, interestingly, it does not affect the R dependence of the string width.

In the high-temperature limit it has been conjectured, using the Svetitsky-Yaffe mapping of the LGT into a suitable two-dimensional spin model, that the terms multiplying the $R/4L$ factor can be resummed [98]:

$$\left(1 + \frac{\pi}{6\sigma L^2} + \dots \right) = \frac{1}{\sqrt{1 - \frac{\pi}{3\sigma L^2}}} = \frac{\sigma}{\sigma(L)}.$$

Thus, the behaviour of the linear term in R of the Nambu-Gotō string width is expected to be:

$$w^2(\sigma, L, R) = \frac{1}{\sigma(L)} \frac{R}{4L} + \dots \quad (2.23)$$

A reliable numerical test of this conjecture is one of the main goals of this thesis.

Finally, to avoid confusion, let us stress that the string width that we discuss in this section is a purely quantum effect, as at the classical level the Nambu-Gotō string has zero width. This is one of the reasons behind the difficulty in obtaining analytical results for this quantity.

2.1.3 The Nambu-Gotō string as a $T\bar{T}$ perturbation

As we show in the previous sections, in the physical gauge the Nambu-Gotō action in $D = 3$ can be seen simply as an ordinary two-dimensional QFT of a single bosonic degree of freedom. What is less obvious, but important for what follows, is that the Nambu-Gotō action is actually a very peculiar case of a QFT. In fact it can be shown [16, 18] that it describes the irrelevant perturbation of the free two-dimensional bosonic action, driven by the $T\bar{T}$ operator. The perturbing parameter is exactly the inverse of the string tension and for large values of σ the action flows towards the theory of a Free Boson. This is easy to see if one performs a $1/\sigma$ expansion (which for dimensional reasons corresponds to a large distance expansion) of eq. (2.32).

Following ref. [18], we can define the component of the energy-

momentum tensor in Euclidean and complex coordinates as:

$$\begin{aligned} T_{\xi^0\xi^0} &= -\frac{1}{2\pi} \left(\bar{T} + T - 2\Theta \right) \\ T_{\xi^1\xi^1} &= \frac{1}{2\pi} \left(\bar{T} + T + 2\Theta \right) \\ T_{\xi^0\xi^1} &= \frac{i}{2\pi} \left(\bar{T} - T \right) \end{aligned} \quad (2.24)$$

where:

$$\mathcal{T} = \frac{T}{\pi} = -\frac{\partial\mathcal{L}}{\partial(\bar{\partial}\varphi)}\partial\varphi, \quad \bar{\mathcal{T}} = \frac{\bar{T}}{\pi} = -\frac{\partial\mathcal{L}}{\partial(\partial\varphi)}\bar{\partial}\varphi, \quad (2.25)$$

and

$$\theta = \frac{\Theta}{\pi} = \frac{1}{2} \left(\frac{\partial\mathcal{L}}{\partial(\partial\varphi)}\partial\varphi + \frac{\partial\mathcal{L}}{\partial(\bar{\partial}\varphi)}\bar{\partial}\varphi - 2\mathcal{L} \right), \quad (2.26)$$

with \mathcal{L} is the Lagrangian of the theory to perturb. The perturbation along the "time" t of the energy spectrum of a conformal field theory leads to an equation for the free energy of the theory:

$$\partial_t \left(\beta F(R, L, t) \right) = -\left\langle \int_0^R d\xi^0 \int_0^L d\xi^1 \left(\mathcal{T}\bar{\mathcal{T}}(z, \bar{z}) \right) \right\rangle \quad (2.27)$$

where (R, L) are the characteristic lengths of the perturbed 2-dimensional QFT, $z = \xi^0 + i\xi^1$, and:

$$\mathcal{T}\bar{\mathcal{T}}(z, \bar{z}) = \mathcal{T}(z, \bar{z})\bar{\mathcal{T}}(z, \bar{z}) - \theta^2(z, \bar{z}). \quad (2.28)$$

The dynamics of the Lagrangian under the perturbation is:

$$\partial_t \mathcal{L} = -\mathcal{T}\bar{\mathcal{T}}(z, \bar{z}); \quad (2.29)$$

in the case of the free bosonic CFT:

$$\begin{aligned} S_{\text{CFT}} &= \frac{1}{4} \int d\xi^0 d\xi^1 (\partial_{\xi^0}\varphi)^2 + (\partial_{\xi^1}\varphi)^2 = \frac{1}{2} \int dz d\bar{z} \partial_z \varphi \partial_{\bar{z}} \varphi \\ &= \int d\xi^0 d\xi^1 \partial\varphi \bar{\partial}\varphi, \end{aligned} \quad (2.30)$$

solving (2.29) leads to:

$$\mathcal{L}(z, \bar{z}, t) = \frac{1}{2t} \left(-1 + \sqrt{1 + 4t \partial\varphi \bar{\partial}\varphi} \right) \quad (2.31)$$

By rescaling the fields $2\varphi = \phi$, we found that the perturbed action is equal to:

$$\mathcal{L}(x, y, t) = \frac{1}{2t} \left(-1 + \sqrt{1 + t(\partial_{\xi^0}\phi)^2 + (\partial_{\xi^1}\phi)^2} \right)$$

that is equivalent to the Nambu-Goto action (2.33) with $t = 1/(2\sigma)$.

We shall make use of this framework in the following, to construct one of the algorithms used in this work by mimicking the $T\bar{T}$ flow and leveraging the fact that the Free Boson action can be sampled as easily as a Gaussian distribution.

2.2 Lattice Nambu-Gotō theory

Leveraging the physical gauge discussed in the previous section, for $D = 3$, the Nambu-Gotō action can be rewritten as:

$$S_{\text{NG}}[X] = \sigma \int_0^L d\xi^0 \int_0^R d\xi^1 \sqrt{1 + (\partial_{\xi^0} X)^2 + (\partial_{\xi^1} X)^2}, \quad (2.32)$$

The latter can be regularized on the lattice by discretizing the worldsheet on a two-dimensional square lattice Λ with size $L \times R$, step size $a = 1$, and index $x = (\tau, \epsilon)$. Following the usual conventions of the finite difference for the discrete derivative we obtain (see ref. [105] for further details):

$$S_{\text{NG}}(\phi) = \sigma \sum_{x \in \Lambda} \left(\sqrt{1 + (\partial_{\mu}\phi(\tau, \epsilon))^2/\sigma} - 1 \right) \quad (2.33)$$

where

$$(\partial_{\mu}\phi(x))^2 = (\phi(\tau, \epsilon) - \phi(\tau - 1, \epsilon))^2 + (\phi(\tau, \epsilon) - \phi(\tau, \epsilon - 1))^2 \quad (2.34)$$

and $\phi(x) = \sqrt{\sigma}X(x)$ is a real scalar field representing the transverse degree of freedom of the string; in the context of lattice gauge theories this quantity corresponds to the density of chromoelectric flux.

While this lattice discretization holds in principle for any type of lattice, in view of the application to the description of Polyakov loop correlators in LGTs we are interested in particular on a cylinder geometry: we fix Dirichlet boundary conditions along the ϵ direction (the lattice points at $\epsilon = R$ are considered part of the lattice volume) and periodic boundary conditions along the τ direction; see fig. 2.2 for a schematic representation of a lattice configuration.

As in the previous section we can expand eq. (2.33) in powers of $1/\sigma$:

$$S_{\text{NG}} \sim S_{\text{FB}} + O(\sigma^{-1}) \quad (2.35)$$

where

$$S_{\text{FB}}[\phi] = \frac{1}{2} \sum_{x \in \Lambda} (\partial_\mu \phi(x))^2 \quad (2.36)$$

is the lattice discretization of the CFT free boson theory.

The lattice discretization has two main effects: the first is the appearance of a set of "bulk" constants, which diverge in the continuum limit², the second is the appearance of finite size corrections which are proportional to negative powers of the lattice sizes R and L and vanish in the continuum limit³. Both sets of terms are not universal, but depend on the details of the lattice regularization and are automatically eliminated by the zeta function regularization which selects only the adimensional terms in the expansion. These adimensional terms are the only ones which are "universal" in the renormalization group sense

²These constants are proportional to L (a "bulk" boundary term due to the Polyakov loops) or to RL (area term). Due to the periodicity in the time direction, no term proportional to R can appear.

³Since we are keeping the lattice spacing fixed to $a = 1$ the continuum limit corresponds to the $R, L \rightarrow \infty$ limit.




Figure 2.2: Schematic representation of a lattice configuration: cyan and magenta sites represent the active volume of the lattice, while yellow sites represent the Dirichlet boundaries where the field is fixed to 0; the boundary in $\epsilon = R$ is considered as a part of the lattice. The width σw^2 of eq. (2.45) is computed averaging only on the magenta sites.

and survive in the continuum limit. One of the main goals in the following will be the comparison of our numerical results with the zeta function predictions for these universal terms.

While the non-universal terms are in principle undetermined, with a few simple choices for the lattice regularization, the contribution to the partition function of S_{FB} can be evaluated exactly also on a finite lattice, thus allowing to use them too as benchmarks of our calculations.

2.2.1 Finite size CFT and non-universal terms

In order to obtain the finite size solution of the partition for the discretized CFT free boson theory we can diagonalize the action S_{FB} using the transformation:

$$\phi(\tau, \epsilon) = \sum_{m=1, n=1}^{L, R-1} \Psi(\tau, \epsilon, m, n) y(m, n) \quad (2.37)$$

with orthonormal eigenfunctions:

$$\Psi(\tau, \epsilon, m, n) = \frac{2}{\sqrt{2LR}} \left(\cos\left(\frac{2m\pi\tau}{L}\right) + \sin\left(\frac{2m\pi\tau}{L}\right) \right) \sin\left(\frac{n\pi\epsilon}{R}\right) \quad (2.38)$$

Under this redefinition of the fields, the action can be written as:

$$S_{\text{FB}}(y) = \sum_{m=1, n=1}^{L, R-1} \lambda_{m, n} y^2(m, n) \quad (2.39)$$

where eigenvalues are found to be:

$$\lambda_k \equiv \lambda_{m, n} = 4 \sin^2\left(\frac{m\pi}{L}\right) + 4 \sin^2\left(\frac{n\pi}{2R}\right). \quad (2.40)$$

(see Appendix A.1 for a rigorous derivation of λ). Performing the Gaussian integrations we obtain

$$Z_{\text{FB}} = \int D\phi e^{-S_{\text{FB}}[\phi]} = \prod_{m=1, n=1}^{L, R-1} \sqrt{\frac{2\pi}{\lambda_{m, n}}} \quad (2.41)$$

or equivalently

$$-\log(Z_{\text{FB}}) = -\frac{1}{2} \sum_{m=1, n=1}^{L, R-1} \log\left(\frac{2\pi}{\lambda_{m, n}}\right). \quad (2.42)$$

This sum is divergent and is usually regularized with the zeta function. However, as mentioned above, in the present case we

are interested also in these non-universal terms, because they will appear in the simulation of the whole Nambu-Gotō action, and we can use the exact solution of the free bosonic part to fix them.

Indeed the sum in eq. (2.42) can be evaluated numerically with any precision and besides the Dedekind function one finds two divergent contributions, which are proportional respectively to the area RL and to the length of the boundary L :

$$-\log Z_{\text{FB}} = A_{\text{FB}} RL + C_{\text{FB}} L + \log \eta(\tau) \quad (2.43)$$

with

$$A_{\text{FB}} = -0.3358177\dots, \quad C_{\text{FB}} = 0.478252\dots \quad (2.44)$$

being the constants which we shall use in Section 5 to benchmark our numerical results.

2.2.2 Width and Shape

In the lattice version of EST, the width can be computed as

$$\sigma w^2(R) = \langle \phi^2(\tau, R/2) \rangle_\tau. \quad (2.45)$$

where the $\langle \dots \rangle_\tau$ expectation value takes the average only over the temporal extension, and the scalar fields ϕ are fixed to be at the half-distance $R/2$ between the color sources, see fig. 2.2.

The numerical approaches described in the next chapters allows in principle for the calculation of any observable over the probability density of the EST model under study. Thus, we focus our attention also on a new observable defined as a combination of the quartic and quadratic moments of the flux density in the midpoint $\phi(\tau, R/2)$:

$$U = 1 - \frac{\langle \phi^4(\tau, R/2) \rangle_\tau}{3 \langle \phi^2(\tau, R/2) \rangle_\tau^2}. \quad (2.46)$$

As in eq. (2.45), the expectation values $\langle \dots \rangle_\tau$ are computed fixing the spatial coordinate to $R/2$. In the literature (albeit in different contexts) the quantity U is called the Binder cumulant and we shall denote it as such in the following. It is identically zero if $\phi(\tau, R/2)$ follows a Gaussian distribution: in this sense, it can be used as a probe of the non-Gaussianity of the flux density. In the LGT framework it is used to describe the "shape" of the flux tube, see ref. [106].

Although the numerical approach presented in this thesis offers advantageous generality, as discussed in the previous section, the study of finite-size EST requires accounting for non-universal terms. However, these corrections have a smaller impact when analyzing observables, as they are normalized by partition functions. Furthermore, since the Binder cumulant is expressed as a ratio of observables, it helps mitigate issues related to divergent terms.

2.3 Higher-order terms beyond the Nambu-Gotō action

The study of the higher-order corrections beyond the Nambu-Gotō action represents one of the most interesting open problems in EST. These corrections must respect Poincaré invariance in the target space and thus the first few allowed terms for a three-dimensional target space are⁴

$$S_{\text{EST}} = \int_{\Sigma} d^2\xi \sqrt{g} [\sigma + \gamma_1 \mathcal{R} + \gamma_2 \mathcal{K}^2 + \gamma_3 \mathcal{K}^4 \dots] \quad (2.47)$$

where γ_i are coupling constants, \mathcal{R} represent the Ricci scalar, and \mathcal{K} is the extrinsic curvature. This expansion is strongly constrained by the so called "low energy universality" which

⁴In four dimensions, due to the presence of two transverse fields two different invariants can be constructed at level four.

states that the first two terms can be neglected when studying ordinary non-Abelian LGTs [4]. In fact \mathcal{R} is a topological invariant in two dimensions and it can be neglected. Furthermore, the \mathcal{K}^2 term is proportional to the equation of motion of the Nambu-Gotō action and can be eliminated by a suitable field redefinition. There are however a few exceptions to the last statement: the most interesting one for our purposes is represented by the Polchinski-Yang solution of the rigid string [107], in which the Nambu-Gotō term is treated as a small perturbation of the \mathcal{K}^2 one. This particular case seems to be realized in the so called "reconfined phase" of trace-deformed Lattice Gauge Theories [108]. For this reason we studied (separately) both the inclusion of the \mathcal{K}^2 and of the \mathcal{K}^4 term in the EST action looking at the following two "Beyond Nambu-Gotō" (BNG) actions, namely:

$$S_{\text{BNG}}^1 = S_{\text{NG}} + S_{\mathcal{K}^2} \quad (2.48)$$

$$S_{\text{BNG}}^2 = S_{\text{NG}} + S_{\mathcal{K}^4}. \quad (2.49)$$

It is interesting to notice that the first of these two actions has indeed a long history. Originally introduced to describe the physics of fluid membranes [109–111], it was later proposed by Polyakov and by Kleinert as a way to stabilize the Nambu-Gotō action [112, 113]. It is often denoted as "rigid string" since the \mathcal{K}^2 term is expected to increase the stiffness of the string. As we shall see below, our results on the string width confirm this intuition. To the best of our knowledge, this is the first time that an explicit quantitative evidence of the effect of the rigidity term in the width of the rigid string is reported.

Following ref. [19] we approximate \mathcal{K}^2 to the first order:

$$\mathcal{K}^2 \sim ((\partial_0 \partial_0 \phi(x))^2 + (\partial_1 \partial_1 \phi(x))^2 + 2(\partial_1 \partial_0 \phi(x))^2). \quad (2.50)$$

Thus, we can write the discretized version of the two terms as

follows:

$$S_{\mathcal{K}^2}(\phi) = \gamma_2 \sum_{x \in \Lambda} \mathcal{L}_{\mathcal{K}^2}(\phi(x)) \quad (2.51)$$

$$S_{\mathcal{K}^4}(\phi) = \gamma_3 \sum_{x \in \Lambda} (\mathcal{L}_{\mathcal{K}^2}(\phi(x)))^2 \quad (2.52)$$

where:

$$\begin{aligned} \mathcal{L}_{\mathcal{K}^2}(\phi(\tau, \epsilon)) &= (\phi(\tau + 1, \epsilon) - 2\phi(\tau, \epsilon) + \phi(\tau - 1, \epsilon))^2 + \\ &+ (\phi(\tau, \epsilon + 1) - 2\phi(\tau, \epsilon) + \phi(\tau, \epsilon - 1))^2 + \\ &+ \frac{1}{8} (\phi(\tau + 1, \epsilon + 1) + \phi(\tau - 1, \epsilon - 1) + \\ &- \phi(\tau + 1, \epsilon - 1) - \phi(\tau - 1, \epsilon + 1))^2 \end{aligned} \quad (2.53)$$

Some special care has to be taken on the boundary $\epsilon = R$: the regularization must take into account the Dirichlet boundary conditions. The non-vanishing terms are:

$$\begin{aligned} \mathcal{L}_{\mathcal{K}^2}(\phi(\tau, R)) &= (\phi(\tau, \epsilon - 2) - 2\phi(\tau, \epsilon - 1))^2 + \\ &+ \frac{1}{2} (\phi(\tau - 1, \epsilon - 1) - \phi(\tau + 1, \epsilon - 1))^2 \end{aligned} \quad (2.54)$$

One of the main results of this thesis is the numerical study of the actions of eqs. (2.51) and (2.52). Our goal is to understand how the shape and thickness of flux tube change as a consequence of the addition of the \mathcal{K}^2 or \mathcal{K}^4 terms.

2.4 Lattice gauge theory in the EST picture

We shall conclude this chapter with some final remarks on the comparison between LGT simulations and EST predictions. As mentioned in Section 2.1.1, the EST partition function provides a prediction for the interquark potential computed through




Figure 2.3: Ground state of SU(2), SU(3) and SU(6) as a function of the physical temperature T . The lines represent the EST predictions, the solid black line is the Nambu-Gotō prediction while the dotted lines include also the higher-order corrections.

correlators of Polyakov loops. In fig. 2.3, results for the interquark potential in SU(2), SU(3), and SU(6) LGT are presented and compared with the Nambu-Gotō and BNG predictions [20,25]. In the infrared regime, all the simulations converge to the Nambu-Gotō prediction. This behavior is not surprising and it is stated by the low energy universality. However, as we approach the deconfinement transition, higher-order corrections become fundamental for discriminating between different gauge groups.

Beyond the interquark potential, two important observables for the comparison between EST and LGT and for the study of the confinement mechanism are the shape and the width of the chromoelectric flux tube. In LGT theory, the main ingredient




Figure 2.4: Width of the SU(2) LGT, for two values of β , as a function of the interquark distance R , the dotted line represents the conjectured behavior from eq. (2.23).

needed to compute the shape is the three-point function:

$$\rho(L, R, y) = \frac{\langle P(0)P^\dagger(R)U(y) \rangle}{\langle P(0)P^\dagger(R) \rangle} - \langle U(y) \rangle \quad (2.55)$$

where P are Polyakov loops at a distance R , and U is a plaquette situated at transversal distance y from the center of the plane drawn by the Polyakov loops. In this setup, we refer to $\rho(y)$ as the shape of the string, which, in lattice EST, corresponds to the distribution of the central field configurations $\phi(\tau, R/2)$. The width and the Binder cumulant are then associated with the second and fourth moments of the distribution defined by the shape of the string. In figure 2.4, the width for SU(2) is plotted in the high-temperature limit and compared with the conjectured solution for the Nambu-Gotō EST (see Section 2.1.2). As evident from the plot, the EST conjectured predictions match the LGT results with great agreement. We anticipate that a rigorous numerical validation of the high-temperature Nambu-

Gotō EST width is one of the main results presented in this thesis. On the other hand, in LGT, the shape of the flux tube is characterized by a Gaussian peak with an exponentially decaying tail. In particular, outside the Gaussian regime, the shape follows an exponential form: $\rho(y) \propto e^{-\lambda|y|}$. As a consequence, a new scale, λ , known as the "intrinsic width" [106], emerges and can be analyzed by examining the Binder cumulant. Since this observable measures the Gaussianity of a theory, in LGT, the Binder cumulant is found to deviate significantly from zero. However, as we shall show in Chapter 7, the shape of the EST is Gaussian, confirming the assumption that the EST is widthless and does not include the intrinsic width, at least in its simplest form.

Chapter 3

Markov Chain Monte Carlo

The main goal when studying a lattice QFT is to evaluate the vacuum expectation values (v.e.v.) of a general observable \mathcal{O} over a distribution p :

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \int d\phi p(\phi) \mathcal{O}(\phi) \quad (3.1)$$

where p follows a Boltzmann distribution weighted by the action S of the target theory:

$$p(\phi) = \frac{1}{Z} e^{-S(\phi)} \quad (3.2)$$

In general, v.e.v. are evaluated numerically by sampling N equilibrium configurations $\{\phi_i\}_{i=0}^N$ from p and using the Monte Carlo estimator:

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} \simeq \frac{1}{N} \sum_{i=1}^N \mathcal{O}(\phi_i) \quad (3.3)$$

The crucial challenge in this approach is the sampling procedure for the configurations ϕ_i . The primary approach in lattice QFT is based on Markov Chain Monte Carlo (MCMC) methods. Over the last few decades, numerical Monte Carlo simulations of lattice-regularized QFT have become the primary non-perturbative tool for studying strongly-interacting theories from first principles. In the context of lattice QCD, MCMC

simulations have provided a robust framework for high-precision theoretical calculations that are routinely compared with experimental results (see, for example, refs. [114, 115]).

Despite the great success, lattice MCMC calculations still suffer from considerable shortcomings: for example, one such issue emerges when simulating the theory close to a critical point, where the *autocorrelation* of a sequence of field configurations sampled via MCMC diverges, giving rise to the so-called *critical slowing down* [62]. In order to state this issue quantitatively, one introduces the integrated autocorrelation time $\tau_{\text{int}}(\mathcal{O})$, which estimates the number of effectively independent samples over which the observable of interest \mathcal{O} is measured. The formal definition of the autocorrelation time is:

$$\tau_{\text{int}}(\mathcal{O}) = \frac{1}{2} + \sum_{t=1}^{\infty} \rho_{\mathcal{O}}(t) \quad (3.4)$$

where the function $\rho_{\mathcal{O}}(t)$ is given by

$$\rho_{\mathcal{O}}(t) = \frac{\langle \mathcal{O}(\phi_i) \mathcal{O}(\phi_{i+t}) \rangle_{\phi \sim p} - \langle \mathcal{O} \rangle_{\phi \sim p}^2}{\langle \mathcal{O}^2 \rangle_{\phi \sim p} - \langle \mathcal{O} \rangle_{\phi \sim p}^2} \quad (3.5)$$

and measures the autocorrelation between the samples ϕ_i generated by the MCMC. When approaching criticality, $\tau_{\text{int}}(\mathcal{O})$ diverges with the correlation length ξ of the system, and one typically observes that $\tau_{\text{int}}(\mathcal{O}) \sim \xi^z$. Thus, the cost of generating independent samples can increase unsustainably as the system approaches the critical point $\xi \rightarrow \infty$, causing the critical slowing down. The value of the exponent z depends on the MCMC algorithm and the observable \mathcal{O} .

Another issue with standard MCMC methods is that, in general, it is extremely hard to obtain reliable estimates of the free energy, $F = -\log Z$, with this approach. A partial solution to this problem is that MCMC approaches can more efficiently compute free energy differences ΔF . It is well known, however,

that typically this involves significant computational costs: commonly used methods to perform such calculations are based on the numerical integration of a derivative with respect to some parameter [116] or on reweighting the field configurations of a simulated ensemble to a target ensemble, specified by different parameter values [117]. The former of these methods, however, introduces a systematic uncertainty due to the discretization of the integration interval; the latter, on the other hand, is often hampered by the fact that the overlap between the most typical configurations in the simulated and in the target ensemble becomes exceedingly small in the thermodynamic limit.

The challenging nature of these two problems continues to motivate the search for alternative techniques to standard MCMC methods. As we will see, these issues create significant obstacles for the numerical simulation of EST on the lattice. On the one hand, the primary observable in EST is the free energy; on the other hand, the large correlation length, due to the strong non-linearity of the Nambu-Gotō action, significantly increases $\tau_{\text{int}}(\mathcal{O})$ for the width and Binder cumulant of the effective string as showed in figure 3.1, where we computed the τ_{int} for the width σw^2 .

In the following, we briefly outline the main characteristics of standard MCMC approaches and introduce the Hybrid Monte Carlo method; readers can refer to the textbooks [105, 118] for a more detailed introduction from the perspectives of lattice field theory, and Bayesian inference, respectively.

Furthermore, the following chapters introduce two potential approaches, one variational and the other based on non-equilibrium thermodynamics, that, by rethinking the calculation of v.e.v. from the ground up, provide a robust alternative to the MCMC methods discussed in this section.

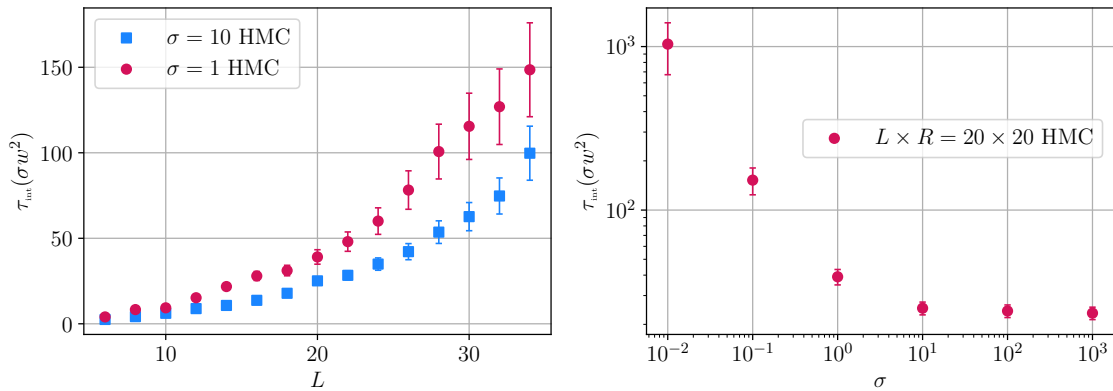


Figure 3.1: τ_{int} for the width computed with standard MCMC methods (see Section 5.5 for more details on the numerical experiments) at fixed σ as a function of the volume for $L \times L$ (left plot) and fixed volume 20×20 and different σ (right plot).

3.1 Hybrid Monte Carlo

An MCMC is a sequence of random variables $\{\phi\}_{i=0}^{\infty}$, where, in the case of a scalar field theory, each $\phi_i \in \mathbb{R}^d$ is generated by updating from the previous sample ϕ_{i-1} . The update is a stochastic process with a specified transition probability $P[\phi \rightarrow \phi']$ that must, of course, satisfy:

$$\int d\phi' P[\phi \rightarrow \phi'] = 1 \quad (3.6)$$

Accordingly, the equilibrium density p of the new sample ϕ' is:

$$p(\phi') = \int d\phi P[\phi \rightarrow \phi'] p(\phi) \quad (3.7)$$

with

$$\int d\phi p(\phi) = 1. \quad (3.8)$$

By definition, a stochastic process satisfying these conditions is called a Markov process, and the resulting sequence of configurations is a Markov Chain. Given an appropriately chosen

initial state ensemble drawn from the density p_0 , applying a long enough sequence of updates leads to the canonical equilibrium ensemble p :

$$\lim_{k \rightarrow \infty} P^k p_0 = p \quad (3.9)$$

Thus, the target density is a fixed point of the transition probability P :

$$Pp = p \quad (3.10)$$

To ensure that this fixed point of the MCMC is achieved, it is sufficient (though not necessary) to impose the *detailed balance* condition:

$$P[\phi \rightarrow \phi']p(\phi) = P[\phi' \rightarrow \phi]p(\phi') \quad (3.11)$$

which implies:

$$\int d\phi P[\phi \rightarrow \phi']p(\phi) = \int d\phi P[\phi' \rightarrow \phi]p(\phi') = p(\phi') \quad (3.12)$$

matching the result in eq. (3.10). With a suitable update procedure, a “reasonable” starting point for the MCMC is typically achieved through a process known as *thermalization*. This process involves generating a sequence of configurations from an arbitrary initial distribution p_0 until a configuration $\phi_j \sim p$ is obtained, thanks to the MCMC properties. At this point, sampling begins to build the desired equilibrium ensemble. Thermalization is one of the essential steps in MCMC because, when p is a highly complex distribution, obtaining the initial equilibrium configuration ϕ_j can be challenging.

One of the most popular update methods for building an MCMC is the so-called Hybrid Monte Carlo (HMC) [119], which combines molecular dynamics (MD) with Metropolis-Hastings (MH) accept-reject steps. The first step in HMC is to extend our target density p by introducing some fictitious Gaussian momenta π :

$$p(\phi, \pi) = \frac{1}{Z} e^{-H_{\text{HMC}}(\phi, \pi)} \quad (3.13)$$

where the new Hamiltonian is given by:

$$H_{\text{HMC}}(\phi, \pi) = S(\phi) + \frac{1}{2} \sum_{x \in \Lambda} \pi^2(x) \quad (3.14)$$

In this augmented space, we can obtain a proposal ϕ' from an initial configuration ϕ by integrating Hamilton's equations of motion:

$$\frac{d\phi}{dt} = \frac{\partial H}{\partial \pi} \quad \frac{d\pi}{dt} = -\frac{\partial H}{\partial \phi} \quad (3.15)$$

where we have introduced the Monte Carlo time t . Since the energy H is conserved during this integration, it constitutes a micro-canonical update. However, because the Hamilton's equations of motion cannot be solved exactly numerically, proposal configurations ϕ' may be biased. To correct this, a Metropolis accept-reject step is applied at the end of the MD integration. Given the previous state of the chain (ϕ_i, π_i) and the proposal (ϕ', π') obtained by integrating eqs.(3.15), the new sample is accepted, i.e., $\phi_{i+1} = \phi'$, with probability:

$$A(\phi \rightarrow \phi') = \min \left[1, \frac{p(\phi', \pi')}{p(\phi, \pi)} \right] \quad (3.16)$$

where

$$\frac{p(\phi', \pi')}{p(\phi, \pi)} = e^{-\Delta H} \quad ; \quad \Delta H = H(\phi', \pi') - H(\phi, \pi) \quad (3.17)$$

If the sample is rejected, then $\phi_{i+1} = \phi_i$. To prove the detailed balance for this accept-reject step, we can first define the proposal of the MD as the kernel $q(\phi \rightarrow \phi')$; observe that in general:

$$p(\phi)q(\phi \rightarrow \phi') \neq p(\phi')q(\phi' \rightarrow \phi) \quad (3.18)$$

since the discretized MD dynamics don't satisfy the detailed balance due to finite size errors. The effect of the accept-reject step with probability A restores the detailed balance:

$$p(\phi)q(\phi \rightarrow \phi')A(\phi \rightarrow \phi') = p(\phi')q(\phi' \rightarrow \phi) \quad (3.19)$$

where A must be fixed such that:

$$A(\phi \rightarrow \phi') = \frac{p(\phi')p(\phi'|\phi)}{p(\phi)p(\phi|\phi')} \quad (3.20)$$

where the $\min [1, A]$ is taken to normalize the kernel A . In the case of HMC, the kernels $q(\phi \rightarrow \phi')$ and $q(\phi' \rightarrow \phi)$ are equal and thus:

$$A(\phi \rightarrow \phi') = \frac{p(\phi')}{p(\phi)} \quad (3.21)$$

We will return to this algorithm in the following sections.

In summary, the HMC algorithm proceeds as follows: given a sample ϕ_j from the MCMC, a new proposal is generated by integrating the MD equations of motion from eq.(3.15) and is then accepted or rejected according to the MH step in eq.(3.16). Generally, after each update, the momenta π are refreshed; however, alternative implementations may partially resample using a Gaussian noise.

3.1.1 Leap-frog integrator

Among the various integrators used to solve the Hamilton's equation of motion (3.15), most of them rely on the repetition of two elementary steps:

$$\begin{aligned} \mathcal{I}_1(\varepsilon) : (\phi, \pi) &\rightarrow (\phi + \varepsilon \nabla_{\pi} H(\phi, \pi), \pi) \\ \mathcal{I}_2(\varepsilon) : (\phi, \pi) &\rightarrow (\phi, \pi - \varepsilon \nabla_{\phi} S(\phi)) \end{aligned} \quad (3.22)$$

where ε is the time shifts of the discretize equation of motion's (3.15). One of the most simple integrator is the Störmer-Verlet integrator, generally referred to as the leap-frog, which corresponds to the composition of:

$$\mathcal{J}_{\varepsilon}(\varpi) = [\mathcal{I}_1(\varepsilon/2)\mathcal{I}_2(\varepsilon)\mathcal{I}_1(\varepsilon/2)]^{N_s} \quad (3.23)$$

where $\varpi = N_s\varepsilon$ is the length of a MD trajectory.

Chapter 4

Neural Generative Samplers

In the previous chapter, we addressed the problem of precisely estimating vacuum expectation values (v.e.v.):

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \int d\phi p(\phi) \mathcal{O}(\phi) \quad (4.1)$$

by sampling equilibrium configurations from the distribution p . However, sampling directly from the desired distribution is not always necessary. Notably, using a reweighting technique called Importance Sampling (IS), we can compute the v.e.v. by drawing configurations from an auxiliary distribution q :

$$\int d\phi p(\phi) \mathcal{O}(\phi) = \int d\phi q(\phi) \frac{p(\phi)}{q(\phi)} \mathcal{O}(\phi) = \left\langle \frac{p(\phi)}{q(\phi)} \mathcal{O}(\phi) \right\rangle_{\phi \sim q} \quad (4.2)$$

The main drawbacks of IS calculations are as follows: q must have good support on p to achieve precise estimations of v.e.v.; additionally, the *importance weight* $w(\phi) = \frac{p(\phi)}{q(\phi)}$ must be computable, which requires $q(\phi)$ and $p(\phi)$ to be evaluated at least in the unnormalized form. Is thus obvious that q must be chosen carefully.

In recent years, thanks to the rapid progress of deep learning, new techniques have been developed to provide robust and expressive variational *Ansatz* for the auxiliary distribution q . In

particular, by using deep neural network architectures, it is possible to build computable parametric densities q_θ that are *trained* to approximate the target p . Two prominent examples in physical applications are Variational Autoregressive Networks [55] and Normalizing Flows [56–58], both of which belong to the field of *variational inference* [120]. The former has been widely investigated in the statistical mechanics community, primarily for sampling Boltzmann distributions on discrete domains (e.g., the Ising model, Potts model, spin glass etc.) [55, 63–65], while the latter has gained significant popularity in lattice field theory for sampling configurations with continuous degrees of freedom [35]. These kind of application are generally referred to as Neural Generative Samplers [59, 121], and provide a new, promising route to extend standard MCMCM methods.

In this thesis, we focus on Normalizing Flows: in the following sections, we discuss this class of generative models, explore specific implementations, outline strategies for training, and demonstrate how to obtain unbiased estimators for the target distribution p .

4.1 Normalizing Flows

Normalizing Flows (NFs) [56–58] are a class of deep generative models able to build variational auxiliary distribution q_θ . The key idea behind NFs is to provide a trainable and expressive variational *Ansatz* for the auxiliary distribution q_θ by leveraging parametric diffeomorphisms to establish an exact mapping between q_θ and a simpler distribution q_0 . Specifically, an NF g_θ is a differentiable and invertible function that enables the exact computation of the density of the transformed samples $\phi = g_\theta(\phi_0)$ using the *change of variables theorem*. Given a sample ϕ_0 drawn from a prior distribution q_0 , the density of the

output ϕ is:

$$q_\theta(g_\theta(\phi_0)) = q_0(\phi_0) |\det J_{g_\theta}(\phi_0)|^{-1} \quad (4.3)$$

where J is the Jacobian matrix of g_θ .

The idea of leveraging diffeomorphisms to obtain exact densities given a base distribution that is easy to sample from has a long history: the first example is the Box-Muller transformation [122] from 1958, used to obtain Gaussian-distributed samples from uniform noise. In the last years, thanks to the explosion of deep learning, neural networks have been exploited to obtain flexible and expressive parametric diffeomorphisms, i.e., NFs. The main idea of NFs is that the parameters, generally denoted by θ , are optimized so that the *learned* variational distribution q_θ closely approximates the target distribution. It is important to emphasize that q_θ is only an approximation of the target. Nonetheless, as we will discuss below, methods like IS allow us to obtain asymptotically unbiased estimations of v.e.v. despite this approximation.

4.1.1 Coupling Layers

One of the most popular ways to implement the diffeomorphisms g is through the so-called coupling layers [123, 124], a class of transformations that, by leveraging specific masking patterns, provide tractable triangular Jacobian determinants. A coupling layer g_θ^l keeps a subset of the sample fixed (e.g., the fields corresponding to the even sites of the lattice) while transforming the remaining degrees of freedom using parametric, invertible transformations. This approach results in a triangular Jacobian, allowing the logarithm of the determinant of J to be computed as the trace of $\log J$.

A notable example is the RealNVP [124] layer, where the parametric functions are affine transformations whose parameters are determined by a neural network. Readers can refer

to the textbooks [120, 125] for a general introduction to neural networks.

Given a sample ϕ_l and a mask b , a coupling layer $\phi_{l+1} = g_\theta^l(\phi_l)$ acts as:

$$\phi_{l+1} = b\phi_l + (1 - b)(\phi_l e^{-s_\theta(b\phi_l)} + t_\theta(b\phi_l)) \quad (4.4)$$

where s and t are the outputs of a neural network that takes as input the unchanged, “frozen,” partition $b\phi_l$.

Normalizing Flows g_θ can generally be constructed by interleaving N coupling layers g_θ^l with alternating masks:

$$g_\theta(\phi_0) = (g_\theta^N \circ \dots \circ g_\theta^1)(\phi_0), \quad (4.5)$$

Because the masking pattern ensures that the Jacobian of each coupling layer is triangular, and the determinant of a composition $\det(g^{l+1} \circ g^l)$ is equal to the product of the determinants $\det(g^{l+1})\det(g^l)$, the output density for the RealNVP NFs of eq. (4.5) is:

$$\log q_\theta(g_\theta\phi_0) = \log q_0(\phi_0) - \log |\det J_{g_\theta}(\phi_0)| \quad (4.6)$$

with:

$$\log |\det J_{g_\theta}(\phi_0)| = - \sum_{i=1}^N s_\theta^i(b_i\phi_{i-1}) \quad (4.7)$$

where $\phi_i = g_\theta^i(\phi_{i-1})$, and b_i is the mask for layer i .

4.1.2 Continuous Normalizing Flows

An intriguing alternative to NFs based on coupling layers are Continuous Normalizing Flows (CNFs). In CNFs, the diffeomorphism g_θ is defined as the solution of a Neural Ordinary Differential Equation (NODE) [126] over time $t \in [0, T]$:

$$\frac{d\phi(t)_x}{dt} = v_\theta(\phi(t), t)_x \quad \text{with} \quad z \equiv \phi(0) \quad \text{and} \quad \phi \equiv \phi(T) \quad (4.8)$$

where x is a site of the lattice Λ , and the vector field v_θ is parameterized by a neural network. The concept of NODE originates from the observation that the Euler discretization of the NODE in eq.(4.8):

$$\phi(t + \delta t)_x = \phi(t)_x + \delta t v_\theta(\phi(t), t)_x \quad (4.9)$$

resembles the so-called residual connections, a widely popular neural network layer [127].

NODEs can be leveraged to construct a NF by noting that the total change in the density of the samples $\phi(t)$ can be determined through the ODE:

$$\frac{d \log q_\theta(\phi(t))}{dt} = - \text{Tr } J_{v_\theta}(\phi(t), t) \quad (4.10)$$

where the trace of the Jacobian corresponds to the divergence of the neural vector field:

$$\text{Tr } J_{v_\theta}(\phi(t)) = (\nabla_{\phi(t)} v_\theta)(\phi(t), t) \quad (4.11)$$

see ref. [126] for a rigorous derivation of this result.

The key advantage of CNFs, as evident from eq.(4.10), is that they eliminate the need to compute Jacobian determinants (a primary limitation of discrete NFs). However, efficient implementation of CNFs depends on the use of high-order integrators to solve the NODE in eq.(4.8), which can be computationally expensive, and the outputs may be biased due to discretization errors in the ODE.

4.2 Training

The key aspect of NFs is that it is possible to optimize the parameters θ of the models to approximate a target distribution p . In the case where p is chosen to be a Boltzmann distribution, $p = \exp(-S)/Z$, we can train a NF g_θ by minimizing, with

respect to the parameters θ , the Kullback-Leibler (KL) divergence:

$$D_{\text{KL}}(q_\theta||p) = \int d\phi q_\theta \log \frac{q_\theta(\phi)}{p(\phi)} \quad (4.12)$$

The KL divergence is a popular metric used to estimate similarity between distributions, and, by definition is always positive definite: $D_{\text{KL}}(q_\theta||p) \geq 0$, while being equal to 0 when q_θ and p are equal. Observe that, since Z does not depend on any parameters, we can minimize the KL divergence by studying the so-called *variational free energy*:

$$\mathcal{F}(\theta) = \int d\phi q_\theta (S(\phi) + \log q_\theta(\phi)) \quad (4.13)$$

Let us emphasize some important points about the KL divergence introduced above. First, note that the expectation value of the variational free energy is calculated over the self-generated samples from the NFs. Consequently, this training procedure, unlike standard machine learning methods, does not require any “training data” but only the functional form of the unnormalized target distribution (e.g. the target action). Another crucial aspect of the KL divergence is that it is, in general, not symmetric under the exchange $q \leftrightarrow p$ (trivially, the expectation value is computed over different distributions). It is well known that optimization using $D_{\text{KL}}(q_\theta||p)$ can suffer from mode collapse [120, 128], meaning that if the target p is multimodal, the learned variational distribution q_θ may collapse onto a single mode of p . This issue can be mitigated by enforcing symmetries that reflect the multimodal nature of p into the NFs. Nevertheless, since lattice EST belongs to a unimodal class of theories, in this thesis, we did not focus our attention on symmetries for NFs.

The variational free energy \mathcal{F} is just the first step in obtaining a reliable loss function to train NFs. Although in this thesis we

trained all the models using \mathcal{F} , it is important to mention that a whole family of loss functions can be derived by studying semi-analytically the gradients of \mathcal{F} . Two noteworthy examples are the reinforce loss with control variates [55, 129], implemented for NFs by [130], and the path gradient [131–133].

4.3 Estimators

Let us now revisit the Importance Sampling (IS) method that we introduced at the beginning of this chapter. Given a NFs g_θ , with density q_θ trained to approximate the target $p = \exp(-S)/Z$, it is possible to obtain unbiased v.e.v. using the reweighting formula:

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \int dp(\phi) \phi \mathcal{O}(\phi) = \int d\phi q_\theta(\phi) \frac{p(\phi)}{q_\theta(\phi)} \mathcal{O}(\phi) \quad (4.14)$$

In order to evaluate the latter, we can now define the *unnormalized importance weight*:

$$\tilde{w}(\phi) = \frac{e^{-S(\phi)}}{q_\theta(\phi)} \quad (4.15)$$

Observe that:

$$1 = \frac{1}{Z} \int d\phi q_\theta(\phi) \tilde{w}(\phi), \quad (4.16)$$

thus, IS also provides an estimator for the partition function itself:

$$Z = \langle \tilde{w}(\phi) \rangle_{\phi \sim q_\theta} \quad (4.17)$$

Knowing this, we can compute the v.e.v. as:

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \frac{\langle \mathcal{O}(\phi) \tilde{w}(\phi) \rangle_{\phi \sim q_\theta}}{\langle \tilde{w}(\phi) \rangle_{\phi \sim q_\theta}} \quad (4.18)$$

The estimators for the partition function and general observables are asymptotically unbiased [59, 61]; however, the quality

of the corresponding Monte Carlo estimators depends on the overlap that the variational density q_θ has on p . Observe that since NFs provide maps between distributions, the autocorrelation of the IS observables matches that of the prior samples. Therefore, NFs can naturally mitigate critical slowing down by carefully choosing the prior distributions.

A popular metric used in the flow-based sampling community to test the quality of the IS (and thus of the training of q_θ) is the so-called Effective Sample Size (ESS), defined as:

$$\text{ESS}(q_\theta) = \frac{\langle \tilde{w}(\phi) \rangle_{\phi \sim q_\theta}^2}{\langle \tilde{w}(\phi)^2 \rangle_{\phi \sim q_\theta}} \quad (4.19)$$

The ESS is bounded in $(0, 1)$: $1 \geq \text{ESS} \geq 0$, and it is equal to 1 when there is a perfect overlap between q_θ and p .

An alternative way to correct the variational distribution q_θ is by using an independent Metropolis-Hastings (iMH) algorithm [60]. In practice, a standard MCMC is run; however, the proposal ϕ' doesn't depend on the previous sample ϕ_i but is sampled independently from q_θ and accepted with probability:

$$A(\phi_i \rightarrow \phi') = \min \left[1, \frac{p(\phi')q_\theta(\phi)}{p(\phi_i)q_\theta(\phi')} \right] = \min \left[1, \frac{\tilde{w}(\phi')}{\tilde{w}(\phi_i)} \right] \quad (4.20)$$

The latter satisfies detailed balance; the proof is the same as for the Metropolis-Hastings steps of the HMC from Section 3.1. Since the proposals are independent, we can rewrite the proposal kernel as $q(\phi_i \rightarrow \phi') = q_\theta(\phi')$, where, in general:

$$p(\phi_i)q_\theta(\phi') \neq p(\phi')q_\theta(\phi_i). \quad (4.21)$$

Thus, to impose detailed balance, we choose A such that:

$$p(\phi_i)q_\theta(\phi')A(\phi_i \rightarrow \phi') = p(\phi')q_\theta(\phi_i). \quad (4.22)$$

with:

$$A(\phi \rightarrow \phi') = \frac{p(\phi')q_\theta(\phi)}{p(\phi)q_\theta(\phi')} = \frac{\tilde{w}(\phi')}{\tilde{w}(\phi)} \quad (4.23)$$

where the min is taken to normalize the kernel A .

The iMH is a well-known algorithm that is closely related to IS: the MCMC tends to select configurations with high importance weight while discarding configurations with low values. In the asymptotic limit of eq. (3.9) it will produce samples from the target p . However, samples drawn with this procedure are affected by autocorrelation. Nevertheless, since the proposals are independent, the autocorrelation of the chain is generated only when a sample is rejected. As in the case of IS, the proposal variational density q_θ must have good support on p in order to obtain high acceptance and thus less autocorrelation. For this reason, iMH is just an alternative to IS, the performances are practically equal and can be tracked through on the ESS.

4.4 Normalizing Flows in lattice QFT simulations

In recent years, the effectiveness of sampler techniques based on NFs, generally called flow-based samplers, has been widely investigated using various approaches across numerous lattice models. These range from scalar theories [26, 60, 61, 66–71], gauge theories [72–78], and fermionic theories [134–136], to QCD [35, 79, 80]. NFs are also widely popular in quantum chemistry, where they are generally referred to as Boltzmann Generators [137]. In this field, they address challenges similar to those in lattice field theory, such as sampling complex energy landscapes [138] and computing free energies [139]. Other notable applications of NFs, which yield nontrivial results for the study of the fine details of confinement in gauge theories, are the works discussed extensively in this thesis [26, 94].

Despite the demonstrated benefits for lower-dimensional theories and their potential to mitigate critical slowing down more

generally, flow-based samplers currently suffer from poor scaling when applied to large lattice volumes, as noted in refs. [66, 76, 81, 82].

Although the precise reasons for this poor scaling remain unclear, it is evident to the flow-based sampling community that a *physics-informed* design—where *a priori* physical knowledge is incorporated into the model design—leads to improved performance [81]. A notable example is provided in ref. [71], where the entanglement entropy is computed for a scalar lattice field theory using the replica trick. In this context, the key observables are free energy differences between replica ensembles, where replicas are partially glued through a defect. The central innovation of ref. [71] is the introduction of a *defect coupling layer*, a coupling layer that operates locally on the defect. At present, this work is one of the few examples where pure NFs have been employed to compute physically meaningful observables, reaching state-of-the-art performances.

Another intriguing and promising approach to scale up NFs is to combine them with Non-Equilibrium MCMC methods, based on out-of-equilibrium statistical mechanics, to obtain Stochastic Normalizing Flows [78, 86, 87]. In the following chapter, we present a proof-of-concept of the application of flow-based samplers to compute interesting quantities in EST. Afterward, in Chapter 6 we introduce Non-Equilibrium MCMC, emphasizing how non-equilibrium thermodynamics can be harnessed to go beyond the standard equilibrium approach for MCMC. In Section 6.3, we will combine these NFs and non-equilibrium methods to develop Stochastic Normalizing Flows.

Chapter 5

A Proof-of-Concept: simulating EST with Continuous Normalizing flows

In the following, we present a proof-of-concept for the application of flow-based samplers to study EST on the lattice. We first introduce our architecture, a CNF, and show that it can be successfully applied to numerically study EST. The focus of our simulations is the computation of both the partition function and the string width in two different regimes of the Nambu-Goto theory: high-temperature (HT) ($R \gg L$) and low-temperature (LT) ($L \ll R$). Afterward, we show that flow-based samplers can clearly outperform the standard MCMC method in the calculation of the width of the string.

In general, observables are computed through the IS procedure described in Section 4.3 using 10^6 configurations sampled with trained CNFs at fixed σ , L , and R .

5.1 Architecture

For the proof-of-concept presented in this thesis, we consider a CNF architecture inspired by [67]. The vector field g_θ that we

use is defined as:

$$g_{\theta}(\phi(t), t)_x = \sum_{y,f} K(t)_f W_{d,x,y} \phi(t)_y \quad (5.1)$$

where the time kernel $K(t) \in \mathbb{R}^F$ is given by the first F coefficients of a Fourier expansion, while $W \in \mathbb{R}^{F \times V \times V}$ is a tensor of learnable weights with $V = L \times (R - 1)$ being the volume of the active variables of the configurations $\phi(t)$. The divergence of the g_{θ} defined in eq. (5.1) is trivial and is found to be

$$(\nabla \cdot g_{\theta})(\phi(t), t) = \text{Tr} \left[\sum_f K(t)_f W_f \right]. \quad (5.2)$$

This model corresponds to a continuous-time extension of the classical linear neuron $y = W\phi$.

For all the models, we use $F = 3$ for the temporal kernel and integrate the NODE over 13 steps with $T = 1$.

5.2 Simulation setup

The performances of the CNFs have been tested using as a metric the ESS. We generally use models with $\text{ESS} \geq 0.1$. Due to the limits of the architectures used in this work, we report results for $\sigma > 5.0$ in the HT regime and for $\sigma > 10.0$ in the LT regime. We initialize the elements of the vector field to 0 (identity initialization) and train the CNFs for 1000 Adam [140] iterations with 10000 samples, 0.0005 initial learning rate, $\beta_1 = 0.8$, $\beta_2 = 0.9$, and a cosine annealing scheduler. To reduce the computational cost of the simulations, we trained each combination of L and R for $\sigma = 5.0, 10.0, 25.0, 100.0$ ($\sigma = 5.0$ only for the HT regime); we then used these models as the initialization for CNFs with different string tensions. When performing this transfer learning procedure between models, we train the new CNFs with

100 Adam iterations with 0.00001 as the initial learning rate. The error is estimated using a jackknife procedure. The code is based on the PyTorch library [141] and a simple implementation can be found in [142]. We ran the computation on Tesla V100 GPUs, with a total cost of approximately 650 GPU hours.

For the high-temperature regime we considered 1804 combinations of L , R and the string tension σ . More precisely, we performed simulations for 11 even values of R between 50 and 100, 9 values of L between 4 and 12, 21 values of σ between 5 and 300 for $L < 8$ and 16 values between 10 and 300 for $L \geq 8$.

In fig. 5.1, the ESS is reported as a function of R , for different σ and $L = 10$. This metric indicates that the overlap between the learned and target distribution is very good, even if the scaling with the volume and with σ can be a limiting factor for the architectures under consideration in this study.




Figure 5.1: Effective Sample Size for CNFs as a function of R for fixed $L = 10$ and various σ . Error bars are not visible due to the very small statistical errors.

In the low-temperature regime, we performed simulations with 1008 different combinations of L , R and the string ten-

sion σ . We extracted results for eight values of R between 8 and 22, six values of L between 84 and 94 and 21 values of σ between 10.0 and 300.0. In fig. 5.2 we report the ESS as a function of R for different σ and $L = 90$. Due to the larger volumes under consideration in this regime the quality of the flows is generally lower than in the HT regime, but it guarantees an efficient sampling of the target distribution nonetheless.




Figure 5.2: Effective Sample Size for CNFs as a function of R for $L = 90$ and different σ .

5.3 Partition Function

As mentioned above we benchmark our approach using results for the partition function. Since we focus on the study of the large- σ region, we used only data for $\sigma \geq 40$. The most direct check is in the low-temperature regime.

R	${}^{pc}a_{LT}^{(0)}(R)$	${}^{pc}a_{LT}^{(1)}(R)$	${}^{pc}a_{LT}^{(2)}(R)$	${}^{pc}a_{LT}^{(3)}(R)$	$\chi^2/d.o.f.$
8	-2.224701(3)	-1.750(1)	0.5(1)	2(2)	1.29
10	-2.893064(5)	-2.280(1)	0.6(1)	7(3)	1.18
12	-3.562534(5)	-2.810(2)	0.5(1)	12(4)	1.15
14	-4.232608(5)	-3.344(2)	0.9(1)	9(4)	0.74
16	-4.903077(6)	-3.878(2)	1.1(1)	9(4)	0.86
18	-5.573818(8)	-4.410(3)	1.4(2)	7(6)	1.28
20	-6.244732(8)	-4.943(3)	1.7(2)	4(6)	1.00
22	-6.91576(1)	-5.480(3)	2.2(3)	-1(7)	0.96

Table 5.1: Results for the coefficients of the fit of $\log Z$ of eq. (5.3) in the LT regime.

5.3.1 Low-temperature regime ($L \gg R$)

Following the discussion of Section 2.1.1, we fitted our data for the partition function for each value of R , assuming a linear dependence on L which keeps into account both the area and the perimeter terms; then we looked at the first few terms of the $1/\sigma$ expansion of this linear behaviour

$$\begin{aligned}
 -\log Z(\sigma, L, R) = & \left({}^{pc}a_{LT}^{(0)}(R) + \right. \\
 & \left. + \frac{{}^{pc}a_{LT}^{(1)}(R)}{\sigma} + \frac{{}^{pc}a_{LT}^{(2)}(R)}{\sigma^2} + \frac{{}^{pc}a_{LT}^{(3)}(R)}{\sigma^3} \right) L .
 \end{aligned} \tag{5.3}$$

where the left superscript pc refers to Proof-of-Concept. Results of these fits are reported in table 5.1.

Then we used the ${}^{pc}a_{LT}^{(0)}(R)$ coefficient to further validate our results, since its behavior should coincide with that of the lattice discretization of the free bosonic action (i.e. the $\sigma \rightarrow \infty$ limit of the Nambu-Goto action). Fitting these values according to

$${}^{pc}a_{LT}^{(0)}(R) = {}^{pc}A_{LT}^{(0)}R + \frac{{}^{pc}B_{LT}^{(0)}}{R} + {}^{pc}C_{LT}^{(0)} \tag{5.4}$$

	${}^{pc}A_{LT}^{(0)}$	${}^{pc}B_{LT}^{(0)}$	${}^{pc}C_{LT}^{(0)}$	$\chi^2/d.o.f.$
	-0.335820(2)	-0.1309(2)	0.47822(4)	0.93
True	-0.3358177...	-0.13089969...	0.478252...	

	${}^{pc}A_{LT}^{(1)}$	${}^{pc}B_{LT}^{(1)}$	${}^{pc}C_{LT}^{(1)}$	$\chi^2/d.o.f.$
	-0.2669(2)	-5(1)	0.393(4)	0.56

Table 5.2: Results for the fit of ${}^{pc}a_{LT}^{(0)}(R)$ of eq. (5.4) (upper table) and of ${}^{pc}a_{LT}^{(1)}(R)$ of eq. (5.5) (lower table) in the LT regime.

led to values for the constants ${}^{pc}A_{LT}^{(0)}$ and ${}^{pc}C_{LT}^{(0)}$ which are compatible with those reported in eq. (2.44) (see table 5.2).

Moreover, the coefficient ${}^{pc}B_{LT}^{(0)}$ coefficient is compatible with the coefficient of the Lüscher term $-\frac{\pi}{24} = -0.13089969\dots$, which is the only remnant in this limit of the Dedekind function. The agreement with the expectation is clearly visible in fig. 5.3, where we plotted the quantity ${}^{pc}a_{LT}^{(0)}(R) - {}^{pc}A_{LT}^{(0)}R - {}^{pc}C_{LT}^{(0)}$ as a function of R and compared it to the Lüscher term $-\frac{\pi}{24R}$.

A similar analysis can be performed also for the higher order terms of the $1/\sigma$ expansion. We report for completeness in the second line of tab. 5.2 the results for ${}^{pc}a_{LT}^{(1)}(R)$, fitted with the expression:

$${}^{pc}a_{LT}^{(1)}(R) = {}^{pc}A_{LT}^{(1)}R + \frac{{}^{pc}B_{LT}^{(1)}}{R^3} + {}^{pc}C_{LT}^{(1)}. \quad (5.5)$$

5.3.2 High-temperature regime ($R \gg L$)

A non trivial aspect of this regime is that, due to the resummation of the infinite set of exponentials in the Dedekind function in eq. (2.11), a logarithmic term appears in the partition function: thus, the analogous of the Lüscher term has a different coefficient moving from $-\frac{\pi}{24R}$ to $-\frac{\pi}{6L^2}$ (see eq. (2.13)). We first performed a set of preliminary fits keeping the coefficient of the logarithmic term as a free parameter: since the result was




Figure 5.3: Plot of ${}^{pc}a_{LT}^{(0)}(L) - {}^{pc}A_{LT}^{(0)}R - {}^{pc}C_{LT}^{(0)}$ of eq. (5.4) as a function of R compared to the Lüscher term $-\frac{\pi}{24R}$ (solid line).

compatible with very high precision to $1/2$, we then fixed the coefficient to this value, subtracted the log term from the data and fitted again with the following function:

$$\begin{aligned}
 & -\log Z(\sigma, L, R) - \frac{1}{2} \log\left(\frac{2R}{L}\right) = \\
 & = \left({}^{pc}a_{HT}^{(0)}(L) + \frac{{}^{pc}a_{HT}^{(1)}(L)}{\sigma} + \frac{{}^{pc}a_{HT}^{(2)}(L)}{\sigma^2} + \frac{{}^{pc}a_{HT}^{(3)}(L)}{\sigma^3} \right) R + \\
 & + {}^{pc}c_{HT}^{(0)}(L) + \frac{{}^{pc}c_{HT}^{(1)}(L)}{\sigma}.
 \end{aligned} \tag{5.6}$$

Results are listed in table 5.3.

Then, in analogy to the low-temperature regime, we fitted the ${}^{pc}a_{HT}^{(0)}(L)$ coefficient with

$${}^{pc}a_{HT}^{(0)} = {}^{pc}A_{HT}^{(0)}L + \frac{{}^{pc}B_{HT}^{(0)}}{L} + \frac{{}^{pc}B1_{HT}^{(0)}}{L^3} + \frac{{}^{pc}B2_{HT}^{(0)}}{L^5} \tag{5.7}$$

and the corresponding results are reported in table 5.4. The value of ${}^{pc}A_{HT}^{(0)}$ is again compatible within two standard devia-

L	${}^{pc}a_{HT}^{(0)}(L)$	${}^{pc}a_{HT}^{(1)}(L)$	${}^{pc}a_{HT}^{(2)}(L)$	${}^{pc}a_{HT}^{(3)}(L)$	${}^{pc}c_{HT}^{(0)}(L)$	${}^{pc}c_{HT}^{(1)}(L)$	$\chi^2/d.o.f.$
4	-1.477442(2)	-1.0654(5)	0.37(4)	1(1)	1.9137(1)	1.61(1)	0.98
5	-1.785398(3)	-1.3318(6)	0.52(5)	0(1)	2.3921(1)	1.95(1)	1.14
6	-2.103064(3)	-1.5973(8)	0.52(6)	2(1)	2.8699(2)	2.34(2)	1.17
7	-2.426093(4)	-1.8627(9)	0.58(6)	3(2)	3.3488(2)	2.67(2)	1.05
8	-2.752383(4)	-2.1280(9)	0.64(7)	4(2)	4.3051(2)	3.43(2)	1.11
9	-3.080823(4)	-2.3952(9)	0.77(7)	4(2)	4.3051(2)	3.43(2)	0.86
10	-3.410762(5)	-2.661(1)	0.87(9)	3(2)	4.7835(2)	3.79(3)	1.08
11	-3.741772(5)	-2.928(1)	1.0(1)	4(3)	5.2612(3)	4.21(3)	1.29
12	-4.073607(5)	-3.194(1)	1.0(1)	5(2)	5.7400(3)	4.56(3)	0.99

Table 5.3: Results for the coefficients of the fit of $\log Z$ of eq. (5.6) in the HT regime.

	${}^{pc}A_{HT}^{(0)}$	${}^{pc}B_{HT}^{(0)}$	${}^{pc}B1_{HT}^{(0)}$	${}^{pc}B2_{HT}^{(0)}$	$\chi^2/d.o.f.$
	-0.335823(2)	-0.5234(2)	-0.179(7)	-0.51(7)	1.89
True	-0.3358177...	-0.523598...			

	${}^{pc}C_{HT}^{(0)}$	${}^{pc}D_{HT}^{(0)}$	$\chi^2/d.o.f.$
	0.47827(3)	0.0007(2)	1.50
True	0.478252...		

	${}^{pc}A_{HT}^{(1)}$	${}^{pc}B_{HT}^{(1)}$	$\chi^2/d.o.f.$
	-0.26612(3)	-0.07(2)	0.43

Table 5.4: Results for the coefficients of the fit of eq. (5.7) (upper table), for the coefficients of eq. (5.8) (middle table) and for the coefficients of eq. (5.9) (lower table).

tions with A_{FB} and ${}^{pc}B_{HT}^{(0)}$ is compatible with the expected value $-\frac{\pi}{6} = -0.523598\dots$. The quality of this agreement is visible in fig. 5.4 where ${}^{pc}a_{HT}^{(0)}(L) - {}^{pc}A_{HT}^{(0)}L$ is plotted as a function of L and compared to the Dedekind function prediction $-\frac{\pi}{6L}$.




Figure 5.4: Plot of ${}^{pc}a_{HT}^{(0)}(L) - {}^{pc}A_{HT}^{(0)}L$ of eq. (5.7) as a function of L compared to the Dedekind function prediction $-\frac{\pi}{6L}$ (solid line).

In a similar way we also studied the L dependence of the "perimeter" term ${}^{pc}c_{HT}^{(0)}(L)$ by fitting:

$${}^{pc}c_{HT}^{(0)}(L) = {}^{pc}C_{HT}^{(0)}L + {}^{pc}D_{HT}^{(0)}. \quad (5.8)$$

Results are listed in table 5.3, where we see that the dominant term ${}^{pc}C_{HT}^{(0)}$ is in excellent agreement with the expected value C_{FB} . We report for completeness also the results of the fit to the $1/\sigma$ term ${}^{pc}a_{HT}^{(1)}(L)$. Using

$${}^{pc}a_{HT}^{(1)}(L) = {}^{pc}A_{HT}^{(1)}L + \frac{{}^{pc}B_{HT}^{(1)}}{L^3} \quad (5.9)$$

we found the values reported in table 5.4.

5.4 Width of the flux tube

The main results of this proof-of-concept is to show that with a NF-based simulation it is possible to compute numerically

the next-to-leading correction to the flux tube width, which has never been measured before in LGT simulations. At the same time we shall use the agreement of our results with the leading order predictions for the flux tube width as a further sanity check of our approach. For the analysis of the flux tube width we used all the values of σ , R and L at our disposal. In the low-temperature regime the next-to-leading corrections are too small and will be below our resolution; however, in the high-temperature regime they are larger and within the precision of our numerical approach.

5.4.1 Low-temperature regime ($L \gg R$)

Following the discussion of Section 2.1.2 we fitted our results with

$$\begin{aligned} \sigma w^2 = & \left(1 + \frac{{}^{pc}e_{LT}^{(0)}}{\sigma} + \frac{{}^{pc}e_{LT}^{(1)}}{\sigma R^2} \right) \left({}^{pc}f_{LT} \log(R) + {}^{pc}g_{LT} \right) + \\ & + \frac{{}^{pc}h_{LT}^{(0)}}{R^2} + \frac{{}^{pc}h_{LT}^{(1)}}{\sigma R^2} \end{aligned} \quad (5.10)$$

where we expect (in agreement with eq. 2.21) ${}^{pc}f_{LT} = \frac{1}{2\pi} = 0.159155\dots$ for the leading correction, ${}^{pc}e_{LT}^{(1)} = -\frac{\pi}{4} = -0.785398\dots$ and ${}^{pc}h_{LT}^{(1)} = \frac{5}{96} = 0.05208\dots$ for the next-to-leading ones. The results for the various coefficients are reported in table 5.5. For the leading term we found an excellent agreement with the expected analytical values while, as anticipated, for the next-to-leading ones the statistical uncertainty of the fits is of the same order of the corrections we aim to detect. In fig. 5.5 we report values of σw^2 as a function of R for different L at fixed $\sigma = 100.0$, with the corresponding curve from the fit.

${}^{pc}e_{LT}^{(0)}$	${}^{pc}e_{LT}^{(1)}$	${}^{pc}f_{LT}$	g_{LT}	${}^{pc}h_{LT}^0$	${}^{pc}h_{LT}^1$	$\chi^2/d.o.f.$
1.016(3)	-4(2)	0.15919(9)	0.1854(3)	-0.023(6)	1(1)	1.02
True	0.785398...	0.159155...			0.05208...	

Table 5.5: Results for the coefficients of the fit of the string thickness of eq. (5.10) in the low-temperature regime.




Figure 5.5: σw^2 in the LT regime as a function of R with $\sigma = 100.0$ and different L . Dotted line represents the best fit of the measures. The error bars are narrower than the points.

5.4.2 High-temperature regime ($R \gg L$)

The situation is more interesting in the high-temperature regime. In agreement with eq. (2.22) we fitted our data with

$$\sigma w^2(\sigma, L, R) = \left(1 + \frac{{}^{pc}i_{HT}^{(0)}}{\sigma} + \frac{{}^{pc}i_{HT}^{(1)}}{\sigma L^2} \right) \cdot \left({}^{pc}j_{HT} \frac{R}{L} + {}^{pc}k_{HT} + {}^{pc}l_{HT} \log(L) \right) \quad (5.11)$$

where we expect for the leading term ${}^{pc}j_{HT} = \frac{1}{4}$, ${}^{pc}l_{HT} = \frac{1}{2\pi} = 0.159155\dots$, and for the next-to-leading correction ${}^{pc}i_{HT}^{(1)} = \frac{\pi}{6} =$

0.523598...; results are reported in table 5.6. We see an excellent agreement with the expected values for all the constants, with a relative error on the next-to-leading correction which is less than 10%. As we mentioned in Section 2.1.2, a non-trivial behaviour of the flux tube width in the HT regime is its linear increase with the interquark distance R . This feature is clearly visible in fig. 5.6, where we plotted the results of the simulations, together with the best fit curves for a few selected values of L as a function of R . In the same figure it is also possible to appreciate the fact that, as L increases, the angular coefficient of the linear term decreases with $1/L$.

Then, in order to isolate the next-to-leading contribution, we looked at the following quantity:

$$\langle \sigma w_{NLO}^2 \rangle_R(\sigma, L) = \left\langle \frac{\sigma w^2(\sigma, L, R)}{{}^{pc}j_{HT} \frac{R}{L} + {}^{pc}k_{HT} + {}^{pc}l_{HT} \log(L)} - 1 - \frac{{}^{pc}i_{HT}^{(0)}}{\sigma} \right\rangle_R, \quad (5.12)$$

where we used the best-fit values for the coefficients ${}^{pc}j_{HT}$, ${}^{pc}k_{HT}$, ${}^{pc}l_{HT}$ and ${}^{pc}i_{HT}^{(0)}$ and we took the average over different values of R (as represented by the $\langle \dots \rangle_R$ parentheses). Namely, we are assuming to have encoded the whole dependence on R of the width in the R/L term at the denominator. Thus, $\langle \sigma w_{NLO}^2 \rangle_R$ should behave as $\frac{\pi}{6\sigma L^2}$ [95–97]: in fig. 5.7 we plot it as a function of σ for three values of L . The same behaviour for $L = 4$ in the $\sigma < 50$ region is shown in fig. 5.8, together with the expected behaviour of the two-loop calculation [95–97]. The agreement with the theoretical expectation can be easily appreciated in the plots.

	${}^{pc}i_{HT}^{(0)}$	${}^{pc}i_{HT}^{(1)}$	${}^{pc}j_{HT}$	${}^{pc}k_{HT}$	${}^{pc}l_{HT}$	$\chi^2/d.o.f.$
	0.991(2)	0.55(5)	0.25007(4)	-0.032(1)	0.1579(5)	1.02
True		0.523598...	0.25		0.159155...	

Table 5.6: Results for the coefficients of the fit of the string thickness of eq. (5.11) in the HT regime.




Figure 5.6: Plot of σw^2 in the HT regime as a function of R with $\sigma = 100.0$ for several values of L . The dotted lines represent the best fit; the error bars are narrower than the points.




Figure 5.7: Plot of $\langle \sigma w_{NLO}^2 \rangle_R$ as a function of σ for three values of L compared to the expected analytical solution $\frac{\pi}{6\sigma L^2}$




Figure 5.8: Plot of $\log \langle \sigma w_{NLO}^2 \rangle_R$ as a function of $\log \sigma$ for $L = 4$ compared to $\log \frac{\pi}{6\sigma L^2}$.

5.5 Comparison with Hybrid Monte Carlo




Figure 5.9: Bias on σw^2 between CNFs and HMC from simulations on $L \times R = 20 \times 20$ volumes.

In order to further validate the efficiency of our approach, we compared the results obtained with flows with simulations performed using the Hybrid Monte Carlo (HMC) algorithm [119]. We performed five simulations at fixed volume $L \times R = 20 \times 20$ and $\sigma = 10, 25, 50, 75, 100$ in order to measure the width σw^2 . We trained and evaluated the CNFs using the same procedure described at the beginning of this section, whereas for the HMC we generated 100000 thermalized configurations for each value of σ . For each molecular dynamics trajectory we fixed $\epsilon = 0.1$ and integrated over 10 steps. We compute errors and autocorrelations for results generated with the HMC algorithm using the Γ -method [143, 144] implementation by [145].

In fig. 5.9, we report the bias between the CNFs and the HMC simulations: the results are fully compatible well within the statistical error. Furthermore, in fig. 5.10 we show the error of both methods multiplied by the square root of the sampling time for a transparent comparison between the cost for the measurements in the two cases. The trained flows appear to be more precise than the HMC simulations by as much as a factor 4. This strongly indicates that in the numerical simulation of the Nambu-Goto string CNFs possess a clear advantage with respect to the more traditional Monte Carlo approach. One of the reasons for the effectiveness of flow-based samplers is, as shown in fig. 3.1, that standard MCMC methods suffer from large autocorrelation times in large volumes or small coupling regimes. This problem is avoided by NFs, as they generate uncorrelated samples, eliminating critical slowing down.

Furthermore, we also mention the fact that in traditional Monte Carlo simulations only ratios of partition functions are directly accessible and their computations via integration methods are notoriously cumbersome. On the other hand, Normalizing Flows offer the possibility (at least in relatively simple models) to compute directly the partition function itself. In light of

the results for the flux tube, where the partition function itself has the important physical meaning of the interquark potential, CNFs clearly represent the more efficient option. However, since a careful analysis of the efficiency of the two approaches in computing $\log Z$ is beyond the scope of this work, we refer to [61] for a rigorous comparison (albeit with a different NF architecture).




Figure 5.10: Error on the thickness $\Delta\sigma w^2$ multiplied by the square root of sampling time.

5.6 Summary

Let us conclude this chapter with a brief summary of the results obtained. Using a simple CNF architecture, we studied the lattice-regularized Nambu-Gotō EST in the large coupling regime. In particular:

- We numerically computed the leading-order contribution to the Nambu-Gotō partition function in both the high- and low-temperature regimes, benchmarking the physical and non-physical terms.

- We analyzed the width of the Nambu-Gotō EST, recovering the well-known leading-order results. In the high-temperature regime, we also observed the next-to-leading correction computed in refs. [95–97].
- We demonstrated that CNFs can outperform standard MCMC methods for equilibrium calculations of v.e.v..

Despite these robust findings, the CNFs used in this chapter still suffer from poor scaling with the string tension σ . In the following, we address the scaling issues of flow-based by exploring how v.e.v. at equilibrium can be generalized using non-equilibrium thermodynamics.

Chapter 6

Non-equilibrium methods and Stochastic Normalizing Flows

In Chapter 4 we discussed an approach for the calculations of v.e.v. that leverages the sampling over an auxiliary variational distribution q_θ ; in this section, we will rethink the v.e.v.:

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \int d\phi p(\phi) \mathcal{O}(\phi) \quad (6.1)$$

using a stochastic, out-of-equilibrium approach, indeed, going beyond the equilibrium canonical ensembles. This problem has been faced in non-equilibrium thermodynamics that, thanks to the so-called *fluctuation relations* [146], has provided a framework to compute equilibrium quantities from non-equilibrium experiments. Non-equilibrium statistical mechanics was introduced in 1993 by Evans, Cohen, and Morris with the aim of studying how macroscopic, irreversible phenomena can arise from microscopic, fully reversible systems, and to generalize the Second Law of thermodynamics [147]. By doing so, Evans et al. introduced a new kind of ensemble, the *transient ensemble*, that generalizes the standard canonical ensembles [148–152]. Some years later, Jarzynski and Crooks provided an estimator and a fluctuation relation that can be leveraged to obtain equilibrium quantities by means of transient ensembles [83, 153].

Consider a general thermodynamic transformation starting from a distribution $q_0(\phi_0) = \exp(-S_0(\phi_0))/Z_0$ and ending in the distribution $p(\phi) = \exp(-S(\phi))/Z$. Crooks' fluctuation theorem states that the ratio of the forward and reverse probability densities characterizing the thermodynamic transformation is equal to the exponential of the (dimensionless) *dissipated work* done on the system during the evolution. Thus:

$$\frac{q_0(\phi_0)\mathcal{P}_f[\phi_0 \rightarrow \phi]}{p(\phi)\mathcal{P}_r[\phi \rightarrow \phi_0]} = e^{W_d(\phi_0, \phi)} \quad (6.2)$$

where \mathcal{P}_f and \mathcal{P}_r are the forward and reverse transition probability densities of the thermodynamic transformations, W_d is called dissipated work and it is equal to $W_d(\phi_0, \phi) = W(\phi_0, \phi) - \Delta F$, with ΔF being the free energy difference between the prior and target, and W is the (dimensionless) work done during the transformation.

One immediate and intriguing consequence of the Crooks' theorem is that, if we average over all the possible forward realizations described by \mathcal{P}_f , the inverse of Crooks' theorem yields:

$$\begin{aligned} 1 &= \int d\phi_0 d\phi q_0(\phi_0)\mathcal{P}_f[\phi_0 \rightarrow \phi] \left(\frac{p(\phi)\mathcal{P}_r[\phi \rightarrow \phi_0]}{q_0(\phi_0)\mathcal{P}_f[\phi_0 \rightarrow \phi]} \right) = \\ &= \langle e^{-W_d(\phi_0, \phi)} \rangle_{\phi_0, \phi \sim q\mathcal{P}_f} \end{aligned} \quad (6.3)$$

It is straightforward to observe that:

$$e^{-\Delta F} = \langle e^{-W(\phi_0, \phi)} \rangle_{\phi_0, \phi \sim q\mathcal{P}_f} \quad (6.4)$$

which corresponds to Jarzynski's equality [83]. Let us stress that the average $\langle \dots \rangle_{\phi_0, \phi \sim q\mathcal{P}_f}$ is general and no assumption about the equilibrium of samples ϕ has been made. In general, as we will prove in the following, we can also compute v.e.v. using Jarzynski's equality as:

$$\langle \mathcal{O}(\phi) \rangle_{\phi \sim p} = \langle \mathcal{O}(\phi) e^{-W_d(\phi_0, \phi)} \rangle_{\phi_0, \phi \sim q\mathcal{P}_f} \quad (6.5)$$

which provides a powerful estimator for equilibrium quantities over non-equilibrium transient ensembles.

In the following, we outline how to design general \mathcal{P}_f and \mathcal{P}_r with an MCMC method, discuss the thermodynamic aspects of this approach, and explore its application to lattice QFT. Afterward, we will combine NFs and NE-MCMC to obtain Stochastic Normalizing Flows, a class of hybrid neural variational and out-of-equilibrium algorithms for which a fluctuation theorem can be written and leveraged to obtain equilibrium quantities.

6.1 Non-Equilibrium Markov Chain Monte Carlo

In the previous paragraph, we discussed the Crooks fluctuation theorem for general thermodynamics transformations. These transformations can be implemented, and leveraged to compute equilibrium quantities, using an algorithm called Non-Equilibrium Markov Chain Monte Carlo (NE-MCMC). A NE-MCMC simulation is composed of an ensemble of out-of-equilibrium evolutions from a prior distribution $q_0 = \exp(-S_0)/Z_0$ to the target distribution $p = \exp(-S)/Z$ that we wish to sample from. Each of these evolutions is generated by drawing a sequence of n_{step} configurations ϕ_n using a varying transition probability. The starting configuration ϕ_0 of this sequence is sampled from the prior q_0 and then driven towards the target distribution p through a sequence of Monte Carlo updates that define the transition probabilities $P_{c(n)}$. These transition probabilities depend on a set of parameters $c(n)$, known as the *protocol*, which controls the intermediate steps of the system's evolution, and drive the configurations sequence from the initial state toward the target distribution, i.e.

$$\phi_0 \xrightarrow{P_{c(1)}} \phi_1 \xrightarrow{P_{c(2)}} \phi_2 \xrightarrow{P_{c(3)}} \dots \xrightarrow{P_{c(n_{\text{step}})}} \phi_{n_{\text{step}}} \equiv \phi \quad (6.6)$$

where we note that the last transition probability is fixed to be the one of the target distribution, i.e. $P_{c(n_{\text{step}})} \equiv P_p$.

Formally, after fixing n_{step} and the protocol $c(n)$, we can introduce the forward \mathcal{P}_f probability distribution of each evolution $[\phi_0, \phi_1, \dots, \phi] \equiv \Phi$, which can be written as:

$$\mathcal{P}_f[\Phi] = \prod_{n=1}^{n_{\text{step}}} P_{c(n)}(\phi_{n-1} \rightarrow \phi_n); \quad (6.7)$$

similarly, the reverse probability distribution \mathcal{P}_r takes the form of:

$$\mathcal{P}_r[\Phi] = \prod_{n=1}^{n_{\text{step}}} P_{c(n)}(\phi_n \rightarrow \phi_{n-1}). \quad (6.8)$$

Provided that the transition probabilities $P_{c(n)}$ satisfy detailed balance, using these definitions we can state Crooks' theorem [153], which relates the probability distributions of forward and reverse evolutions to the dissipation of the evolution Φ :

$$\frac{q_0(\phi_0)\mathcal{P}_f[\Phi]}{p(\phi)\mathcal{P}_r[\Phi]} = \exp(W(\Phi) - \Delta F). \quad (6.9)$$

In this expression, $\Delta F = \log Z/Z_0$ is the free energy difference between the states described by q_0 and p , and W represents the dimensionless work done on the system during its transformation from the initial to the final state:

$$W(\Phi) = \sum_{n=0}^{n_{\text{step}}-1} \{S_{c(n+1)}[\phi_n] - S_{c(n)}[\phi_n]\}. \quad (6.10)$$

We also introduce the dissipated work $W_d(\Phi) = W(\Phi) - \Delta F$, which serves as a measure of the dissipation of the transformation between q_0 and p uniquely defined by n_{step} and the protocol $c(n)$. Interestingly, we can rewrite the work as

$$W(\Phi) = S(\phi) - S_0(\phi_0) - Q(\Phi), \quad (6.11)$$

where Q denotes the pseudo heat exchanged during the evolution, defined as:

$$Q(\Phi) = \sum_{n=1}^{n_{\text{step}}} \{S_{c(n)}[\phi_n] - S_{c(n)}[\phi_{n-1}]\}. \quad (6.12)$$

The quantity Q tracks the energy difference at each step as the system is updated at each step n with the transition probability defined by the protocol $c(n)$. It is the natural quantity that arises when deriving eq. (6.9) from the definition of forward and reverse probabilities; in particular we remark that the detail balance condition must be satisfied by the Monte Carlo updates:

$$\frac{P_{c(n)}(\phi_{n-1} \rightarrow \phi_n)}{P_{c(n)}(\phi_n \rightarrow \phi_{n-1})} = \frac{p_{c(n)}(\phi_n)}{p_{c(n)}(\phi_{n-1})} \quad (6.13)$$

where $p_{c(n)} \propto \exp(-S_{c(n)})$ is the distribution fixed in the Markov kernels.

The NE-MCMC algorithm introduced above is equivalent to, and derived independently from, the so-called Annealed Importance Sampling (AIS) [84]. In the following, we show how NE-MCMC can be leveraged to obtain equilibrium quantities from non-equilibrium measures by exploiting Jarzynski's equality.

6.2 Jarzynski's equality and Non-Equilibrium iMH

Let us now focus on how to use NE-MCMC to compute the expectation value of an observable \mathcal{O} over the probability distribution p . First, as done at the beginning of the section, we introduce the average $\langle \dots \rangle_{\Phi \sim q_0 \mathcal{P}_f}$ over all the possible forward evolutions drawn accordingly to $q_0 \mathcal{P}_f$, that we naturally define

as:

$$\begin{aligned} \langle \dots \rangle_{\Phi \sim q_0 \mathcal{P}_f} &= \int d\phi_0 \dots d\phi q_0(\phi_0) \mathcal{P}_f[\phi_0, \dots, \phi] \dots \\ &= \int d\Phi q_0(\phi_0) \mathcal{P}_f[\Phi] \dots \end{aligned} \quad (6.14)$$

where we used the shorthand $d\Phi = d\phi_0 \dots d\phi$ for the integration over all possible configurations. Since the reverse sequence ϕ, \dots, ϕ_0 starts from configurations drawn accordingly to the target p , computing $\mathcal{O}(\phi)$ over p or \mathcal{P}_r is equivalent:

$$\begin{aligned} \langle \mathcal{O} \rangle_p &= \int d\phi p(\phi) \mathcal{O}(\phi) \\ &= \int d\Phi p(\phi) \mathcal{P}_r[\phi, \dots, d\phi_0] \mathcal{O}(\phi) \\ &= \int d\Phi q_0(\phi_0) \mathcal{P}_f[\Phi] e^{-W_d(\Phi)} \mathcal{O}(\Phi) \end{aligned} \quad (6.15)$$

where in the last step we used Crooks' theorem (6.9). We finally obtain the reweighting-like formula

$$\langle \mathcal{O}(\phi) \rangle_p = \langle \mathcal{O}(\Phi) e^{-W_d(\Phi)} \rangle_{\Phi \sim q_0 \mathcal{P}_f} \quad (6.16)$$

where we correct for the fact that the probability distribution generated at the end of the evolution does not correspond exactly to p . In order to determine ΔF we just set $\mathcal{O} = 1$, obtaining

$$1 = \langle e^{-W_d(\Phi)} \rangle_{\Phi \sim q_0 \mathcal{P}_f} \quad (6.17)$$

that in turn leads to Jarzynski's equality [83]:

$$\langle e^{-W(\Phi)} \rangle_{\Phi \sim q_0 \mathcal{P}_f} = e^{-\Delta F}. \quad (6.18)$$

The computation of free energy differences between two distributions using Monte Carlo simulations is notoriously cumbersome, as it generally requires a large number of samples from intermediate distributions *at equilibrium*. Thus, eq. (6.18) represents a

powerful estimator of free energy differences, as the intermediate distributions can be sampled without ever letting the system thermalize.

Note that no assumptions regarding thermalization were made during the construction of the NE-MCMC evolutions: the only exception is that if the prior distribution q_0 is sampled with a MCMC, it must be at equilibrium. While the starting and ending states of the transformations are fixed, there are in principle no restrictions on the protocol, either in the velocity of the evolution (i.e. n_{step}) or in the functional form of $c(n)$; thus, the intermediate and final configurations ϕ_n , with $0 < n \leq n_{\text{step}}$, are in principle allowed to lie arbitrarily far from equilibrium.

In practice, however, the reliability of NE-MCMC estimators is deeply connected to how far from equilibrium evolutions defined by a given protocol are. Namely, it is of the utmost importance to understand how the design of a NE-MCMC (in particular, different choices in terms of n_{step} and $c(n)$) can affect its performance.

An intriguing consequence of this non-equilibrium framework is that the average dissipated work W_d over all the possible evolutions of the NE-MCMC, is equivalent to the Kullback-Leibler (KL) divergence between the forward and reverse probability densities \mathcal{P}_f and \mathcal{P}_r , i.e.:

$$\begin{aligned} \langle W_d(\Phi) \rangle_{\Phi \sim q_0 \mathcal{P}_f} &= \left\langle \log \frac{q_0(\phi_0) \mathcal{P}_f(\Phi)}{p(\phi) \mathcal{P}_r(\Phi)} \right\rangle_{\Phi \sim q_0 \mathcal{P}_f} \\ &= D_{\text{KL}}(q_0(\phi_0) \mathcal{P}_f \| p(\phi) \mathcal{P}_r) \geq 0 \end{aligned} \quad (6.19)$$

The KL divergence of eq. (6.19) measures the similarity between two distributions and, in this non-equilibrium context, provides a way to study the reversibility of a stochastic thermodynamical transformation. Large dissipation, which physically corresponds to strong irreversibility, leads to a large KL divergence; conversely, when the evolution is fully reversible we have for-

ward and reverse evolution are indistinguishable, resulting in $D_{\text{KL}}(q_0\mathcal{P}_f\|p\mathcal{P}_r) = 0$. Intriguingly, since the KL divergence is positive-definite eq. (6.19) implies the Second Law of thermodynamics for these non-equilibrium transformations, i.e.

$$\langle W(\Phi) \rangle_{\Phi \sim q_0\mathcal{P}_f} \geq \Delta F \quad (6.20)$$

providing a natural interpretation in terms of the thermodynamics of Markov Chain Monte Carlo.

The KL divergence of eq. (6.19) represents a crucial metric to determine the reliability of NE-MCMC when sampling a finite number of evolutions. To understand why, we start by looking at the probability distribution of the last configuration $U \equiv U_{n_{\text{step}}}$, i.e. the one generated by a given protocol at the end of the evolution. We can write it as:

$$q(\phi) = \int d\phi_0 \dots d\phi_{n_{\text{step}}-1} q_0(\phi_0) \mathcal{P}_f(\Phi) \quad (6.21)$$

and it is in general intractable unless we let the system thermalize. In order to quantitatively describe the overlap between q and the target p , we would like to estimate the KL divergence

$$D_{\text{KL}}(q\|p) = \int d\Phi q(\phi) \log \left(\frac{q(\phi)}{p(\phi)} \right) = \langle \log \frac{q}{p} \rangle_{\Phi \sim q_0\mathcal{P}_f}. \quad (6.22)$$

It is easy to prove that the divergence of eq. (6.19) represents an upper bound for $D_{\text{KL}}(q\|p)$:

$$\begin{aligned} D_{\text{KL}}(q_0\mathcal{P}_f\|p\mathcal{P}_r) - D_{\text{KL}}(q\|p) &= \langle \log \frac{q_0\mathcal{P}_f}{p\mathcal{P}_r} - \log \frac{q}{p} \rangle_{\Phi \sim q_0\mathcal{P}_f} \\ &= \langle \log \frac{q_0\mathcal{P}_f}{q\mathcal{P}_r} \rangle_{\Phi \sim q_0\mathcal{P}_f} \geq 0 \end{aligned} \quad (6.23)$$

where the last inequality holds since $D_{\text{KL}}(q_0\mathcal{P}_f\|q\mathcal{P}_r)$ is itself a KL divergence. It is then clear that studying the reversibility of a NE-MCMC with $D_{\text{KL}}(q_0\mathcal{P}_f\|p\mathcal{P}_r)$ provides a reliable metric to

evaluate the quality of the sampling of the target distribution. In particular, protocols far from equilibrium lead to estimators of eqs. (6.18) and (6.16) with poor overlap with p , whereas quasi-reversible evolutions are guaranteed to sample it efficiently.

One can understand this by observing that irreversible evolutions characterized by $W(\Phi) \gg \Delta F$ will have a vanishing weight in eqs. (6.16) and (6.18): conversely, transformations with negative dissipation, which violate the Second Law ($W(\Phi) < \Delta F$), provide the largest weights; a good sampling of such evolutions is instrumental for a reliable estimation of the non-equilibrium averages.

As in the case of NFs, another metric that is useful to evaluate the performance of non-equilibrium algorithms is the ESS, that in the context of NE-MCMC is defined as

$$\text{ESS} = \frac{\langle \exp(-W) \rangle_{\Phi \sim q_0 \mathcal{P}_f}^2}{\langle \exp(-2W) \rangle_{\Phi \sim q_0 \mathcal{P}_f}} = \frac{1}{\langle \exp(-2W_d) \rangle_{\Phi \sim q_0 \mathcal{P}_f}}. \quad (6.24)$$

This quantity is an approximate estimator of the ratio

$$\frac{\text{Var}(\mathcal{O})_p}{\text{Var}(\mathcal{O})_{\text{NE}}} = \text{ESS}$$

where the variance at the numerator is the one of an observable \mathcal{O} computed directly on the target distribution p , whereas at the denominator we find the variance of the estimator of eq. (6.16). The interpretation in terms of NE-MCMC evolutions is straightforward: it is indeed easy to see that

$$\text{Var}(\exp(-W)) = \left(\frac{1}{\text{ESS}} - 1 \right) \exp(-2\Delta F) \geq 0. \quad (6.25)$$

The positivity of the variance directly implies $0 \leq \text{ESS} \leq 1$. More generally, this quantity gives us a quantitative description of the distribution of the weights of $\exp(-W(\Phi))$ across different evolutions: indeed, investigating this distribution represents an effective way to monitor the overlap with the target distribution.

One simple and effective approach to either reduce the KL divergence of eq. (6.19) or to increase the ESS of eq. (6.24) is to increase the number of steps in the NE-MCMC protocols, i.e. going towards the $n_{\text{step}} \rightarrow \infty$ limit: smoother, “quasi-static” transformations keep the thermodynamic transformations close to equilibrium, thereby maintaining small dissipation. A more elaborate and potentially much more effective strategy is to optimize the functional form of the protocol $c(n)$ itself: this approach remains one of the main challenges for implementing more efficient NE-MCMC algorithms. We refer to ref. [154] for a recent review on the matter.

Since NE-MCMC effectively builds maps between samples, the autocorrelation on samples obtained from the estimators of eqs. (6.16) and (6.18) cannot be larger than the one of the starting distribution q_0 . In this sense this approach can be exploited in a novel strategy to mitigate critical slowing down: for instance, configurations at coarser lattice spacing can be sampled with reasonable autocorrelation and each of them can then be used as starting point for a NE-MCMC. Eq. (6.16) can be used to compute unbiased expectation values at finer lattice spacing, where we expect critical slowing down to be more severe.

An alternative unbiased way to build estimator over p , is by leveraging a Non-Equilibrium iMH (NE-iMH) Markov Chain that sample from $p\mathcal{P}_r$ using as proposal $q\mathcal{P}_f$. The derivation is the same of the iMH of Section (4.3), with acceptance:

$$\begin{aligned} A(\phi^{(j-1)} \rightarrow \phi') &= \\ &= \min\left(1, \frac{p(\phi')\mathcal{P}_r[\Phi']}{p(\phi^{(j-1)})\mathcal{P}_r[\Phi^{(j-1)}]} \frac{q(\phi_0^{(j-1)})\mathcal{P}_f[\Phi^{(j-1)}]}{q(\phi_0')\mathcal{P}_f[\Phi']}\right) \end{aligned} \quad (6.26)$$

Thanks to the Crooks’ fluctuation relation we can compute the

non-trivial term as:

$$\frac{p(\phi')\mathcal{P}_r[\Phi']}{q(\phi'_0)\mathcal{P}_f[\Phi']} \frac{q(\phi_0^{(j-1)})\mathcal{P}_f[\Phi^{(j-1)}]}{p(\phi^{(j-1)})\mathcal{P}_r[\Phi^{(j-1)}]} = e^{-\left(W(\Phi')-W(\Phi^{(j-1)})\right)} \quad (6.27)$$

A qualitative explanation of the latter is that we want to "filter" the equilibrium configurations from the non-equilibrium distribution $q_0\mathcal{P}_f$. As in the case of the iMH for NFs, the autocorrelation of the chain is refreshed every time a new configuration is accepted and the performance strictly depends on the support that $q_0\mathcal{P}_f$ has on $p\mathcal{P}_r$. Intriguingly, with this NE-iMH we are collecting the samples that either violate the forward or reverse Second Law:

$$\langle W \rangle_{\Phi \sim q_0\mathcal{P}_f} \geq \Delta F \geq -\langle W \rangle_{\Phi \sim p\mathcal{P}_r} \quad (6.28)$$

thus the samples for which the dissipated work W_d (either forward or reverse) is smaller than 0: $W_d \leq 0$.

6.2.1 Jarzynski's equality in lattice QFT simulations

In the context of lattice gauge theories Jarzynski's equality has seen a wide range of applications in the last few years, see refs. [85, 88–93]. In ref. [85] out-of-equilibrium Monte Carlo simulations were applied for the first time in lattice field theory to compute the interface free energy in the \mathbb{Z}_2 gauge theory. Over the last decade these methods have been employed with great effectiveness for high-precision lattice calculations across a variety of domains, including the equation of state of SU(3) Yang-Mills theory [88], the running coupling in SU(N) gauge theories [89], the entanglement entropy in gauge theories [90, 91] and topological freezing [92, 93]. We stress that the implementation of the protocol is not limited to changes in the coupling of the theory (controlling the lattice spacing); it can also involve variations in the geometry of the lattice configurations. In ref. [85] the

free energy of an interface was computed changing the boundary conditions. In refs. [90, 91], the entanglement entropy of lattice quantum field theories was computed using evolutions that connect replica spaces with different geometries. Additionally, ref. [92] addresses directly topological freezing by leveraging non-equilibrium evolutions between a system with open boundary conditions and one with periodic boundary conditions.

6.3 Stochastic Normalizing Flows

NE-MCMC evolutions often require large number of steps to provide a reliable sampling tool, especially in the case of non-optimal protocols. One intriguing way to tackle this problem is to augment NE-MCMC with NFs to obtain Stochastic Normalizing Flows (SNFs) [86], thus combining variational and non-equilibrium methods. In particular, our proposal is to interleave deterministic, NFs functions g_θ^n , with the Monte Carlo updates of the NE-MCMC. In practice, the sequence of eq. (6.6) becomes

$$\phi_0 \xrightarrow{g_1} g_\theta^1(\phi_0) \xrightarrow{P_{c(1)}} \phi_1 \xrightarrow{g_\theta^2} g_\theta^2(\phi_1) \xrightarrow{P_{c(2)}} \dots \quad (6.29)$$

The idea of combining stochastic and deterministic transformations to enhance NE-MCMC-like methods was already investigated in refs. [155–157]. The Crooks’ theorem for this algorithm can be written defining the transition probability for the g_n as δ function:

$$P_{g_\theta^n}(\phi_{n-1} \rightarrow \phi_n) = \delta(\phi_n - g_\theta^n(\phi_{n-1})) \quad (6.30)$$

and modifying \mathcal{P}_f and \mathcal{P}_r appropriately. Eq. (6.9) then still holds if we naturally generalize the work W of eq. 6.11 as:

$$W(\Phi) = S(\phi) - S_0(\phi_0) - Q(\Phi) - \log J_\theta(\Phi), \quad (6.31)$$

where the new term:

$$\log J_{\theta}(\Phi) = \sum_{n=1}^{n_{\text{step}}} \log J_{g_{\theta}^n}(\phi_{n-1}) \quad (6.32)$$

corresponds to the sum of the logarithms of the Jacobian determinant of the NF layers and comes from the use of the change of variables theorem. In particular we have more explicitly that

$$W(\Phi) = \sum_{n=0}^{n_{\text{step}}-1} S_{c(n+1)} [g_{\theta}^{n+1}(\phi_n)] - S_{c(n)} [\phi_n] + \quad (6.33)$$

$$- \log J_{g_{\theta}^{n+1}}(\phi_n).$$

Setting g^n to the identity we immediately recover NE-MCMC. Naturally, we want to tune the parameters $\{\theta^{(n)}\}$ of the functions g_{θ}^n to improve the performances of NE-MCMC. In a very natural way, we optimize the SNF by computing the appropriate generalization of the KL divergence of eq. 6.19 and minimizing it:

$$\min_{\{\theta\}} D_{\text{KL}}(q_0 \mathcal{P}_f \| p \mathcal{P}_r) = \min_{\{\theta\}} \langle W_d(\Phi) \rangle_f \quad (6.34)$$

Numerically, this optimization tunes the parameters of the coupling layers in such a way to minimize the dissipated work W_d of each trajectory as much as possible. From a thermodynamic perspective, the optimized SNF generates a more reversible set of non-equilibrium evolutions compared to a NE-MCMC with the same protocol. SNFs thus offer a way to optimize NE-MCMC, forcing the evolutions to be as reversible as possible through deterministic maps, while inheriting the scalability of the stochastic updates.

A similar approach to SNFs is the Continual Repeated Annealed Transport Flow (CRAFT) [158] method, in which Sequential Monte Carlo [159], an evolution of AIS, is combined with NFs. In this approach, each NF layer of the variational

NE-MCMC is trained separately by minimizing the KL divergence $D_{\text{KL}}(q_n|p_{c(n)})$, where $\phi_n = g_\theta^n(\phi_{n-1})$ and $\phi_{n-1} \sim p_{c(n-1)}$, to optimize the transition between the preceding and successive Markov kernel of the deterministic transformation. Alternative approaches that combine machine learning models and non-equilibrium identities explore the application of Langevin dynamics for sampling and computing observable from Boltzmann distributions [160, 161].

6.3.1 SNFs in lattice QFT

In lattice field theory, SNFs have been first applied to the ϕ^4 scalar field theory in two dimensions [87], and then to more challenging theoretical setups such as the sampling of Effective String Theories on the lattice [94] (this thesis) and the computation of entanglement entropy in scalar field theories in two and three dimensions [71]. Recently, SNFs have establish as the most promising class of flow-based samplers in lattice QFT; namely, in ref. [78] an SNF for pure SU(3) lattice gauge theory have been provided, showing a promising soft polynomial scaling with with the degrees of freedom of the system.

Chapter 7

EST results with Stochastic Normalizing Flows

In this chapter, we present the main numerical results obtained using flow-based samplers to study lattice regularized EST. In particular, after introducing the SNFs used in our numerical calculations and showing that they outperform all existing baselines, we test the algorithm on the well-known results for the Nambu-Gotō partition function. Afterward, we present new results for the width and shape of the EST, studying both the Nambu-Gotō theory and its higher-order corrections.

7.1 Architecture and simulation details

In this thesis, we took inspiration from the $T\bar{T}$ irrelevant integrable perturbations [17, 18] (see Section 2.1.3) to design the SNF used in our simulations. Specifically, we used as a prior distribution a massless free field regularized on the lattice:

$$S_{c(0)}[\phi] = S_0[\phi] = \frac{1}{2} \sum_{x \in \Lambda} (\partial_\mu \phi)^2, \quad (7.1)$$

while we chose the intermediate actions for the non-equilibrium Monte Carlo updates to correspond simply to the Nambu-Gotō

action with a string tension σ_n , i.e.

$$\begin{aligned} q_{c(n)}[\phi] &= \exp(-S_{c(n)}[\phi])/Z_n \\ S_{c(n)}[\phi] &= \sigma_n \sum_{x \in \Lambda} \left(\sqrt{1 + (\partial_\mu \phi(x))^2 / \sigma_n} - 1 \right). \end{aligned} \quad (7.2)$$

The specific protocol for σ_n that we have followed in this work is a linear interpolation in the inverse of the string tension, starting from $1/\sigma = 0$ (i.e. the action of eq. (2.36)) and ending at the inverse of the target value of the string tension. In the case of the higher-order corrections we use as a prior distribution the Nambu-Gotō theory (again obtained with a SNF) and then the couplings γ_2 or γ_3 are switched on using a linear protocol. The prior $q_{c(0)}$ can be sampled exactly using a procedure similar to the one used in ref. [81]: we first sample a Gaussian distribution with identity covariance, and then we rescale the samples using the eigenvalues of the theory to obtain free field configurations in the momentum space. Finally, we transform the samples using the eigenfunctions of the analytical solution. See Appendix (A.2) for further details on this procedure.

The SNFs used in this work are made up by n_{step} blocks, each of which is composed of two affine coupling layers with even-odd masks (see Section 4.1.1) and one HMC update that follows the $T\bar{T}$ protocol described above (see Section 3.1). The two coupling layers of each block share the same convolutional neural network. The number of channels and layers depends on the specific set of simulations (generally between 8 and 32 channels and 1 and 4 layers), balancing performance and computational costs. The same reasoning has been followed for the total number of blocks. The HMC update uses a leapfrog integrator with 10 steps and a stepsize of 0.1.

To train the models, we strongly rely on transferring the parameters of the networks between volumes. We first trained a single PI-SNF for a 40×40 lattice at fixed target coupling for




Figure 7.1: ESS as a function of $1/\sigma$ on 40×40 lattices for CNFs and SNFs.

5000 iterations of the Adam algorithm [140] with batch size 32 and learning rate 0.0001; we then used the weights of this flow as the initialization for the other lattice volumes and trained for 100 iterations with learning rate 0.00001.

The partition function and the expectation value of the width are computed using the reweighting procedure described in sec. 6.2. In order to study the Binder cumulant we used the output of the PI-SNFs as a proposal for the NE-iMH algorithm described in Section 6.2: the expectation values of eq. (2.46) are then computed on the configurations of the new Markov Chain. The code is based on the PyTorch library [141] and the statistical analysis is performed using pyerrors [145].

7.2 Comparison between SNFs, NFs, and NE-MCMC

In order to compare different architectures, we used the ESS estimator introduced in the previous chapters. As a first step, in fig. 7.1, we show a comparison between SNFs and CNFs used in Chapter 5 at a fixed volume of 40×40 : as can be observed from the plot, SNFs outperform CNFs and can target far smaller

values of σ (at least two orders of magnitude). In fig. 7.2, we present a comparison between SNFs, and the underlying NFs and NE-MCMC as a function of the number of blocks divided by the target volume, at a fixed $\sigma = 1$, and volumes 40×40 and 20×20 . It is noteworthy to observe that both pure NFs and NE-MCMC exhibit poor performance in this σ regime with the given architecture. However, both algorithms show a positive trend when increasing the number of layers. This behavior is theoretically expected for the latter, while for the former, it is more surprising since, in general, NFs in standard setups do not exhibit a clear scaling. Furthermore, the results for each algorithms roughly collapse on the same curve, meaning a soft polynomial scaling with the degree of freedom as showed in ref. [78]. Altogether, SNFs clearly outperform the underlying algorithms at a fixed number of steps, inheriting a clear scaling from NE-MCMC while boosting the performance of the underlying NFs. We also stress that the CNFs used in Chapter 5 outperform standard MCMC methods in the calculation of observables.

7.3 Nambu-Gotō: partition function

We start our simulations by using analytical results for the Nambu-Gotō partition function as a benchmark of our approach. In particular, our goal is to test the ability to clearly distinguish the prediction of the NG action from that of the free bosonic action, i.e if we are able to see the effect of the higher order terms in the NG action beyond the first order approximation.

In order to pursue this program, we had to reach small values of the string tension σ while keeping the lattice sizes reasonably large. This is a particularly demanding setting even for flow-based approaches, as the CNF used in Chapter 5 could sample only values of $\sigma \sim \mathcal{O}(10)$. The SNF architecture used in this work represents a substantial improvement in the scaling with




Figure 7.2: ESS as a function of the number of steps divided by the target volume for SNFs and the underlying NFs and NE-MCMC. Colors represent target lattice volumes.

the string tension, as values of $\sigma \sim \mathcal{O}(10^{-2})$ are now accessible.

In order to make the analysis more straightforward, we set our simulations in the high-temperature regime in which the expected behaviour of the NG free energy is given by eq. (2.9). We computed the partition function of the NG model for $\sigma = 1/30$ and $\sigma = 1/50$ and for lattices with $R \in [60, 100]$ and $L \in [10, 15]$. We first fitted the results for $-\log Z$ in R with an expression inspired by eq. (2.9), namely

$$-\log Z = a(L)\sigma R + b(L) + c(L) \log(R). \quad (7.3)$$

The results for the various coefficients are reported in table A.1.

Next, we take a closer look at the coefficient $a(L)$ by using a functional form reminiscent of the NG behaviour of eq. (2.8).




Figure 7.3: Results for the coefficient $a(L)$ of the partition function of the lattice Nambu-Gotō model, for various values of L and for $\sigma = 1/30$ (left panel) and $\sigma = 1/50$ (right panel). The expected behaviour for the Free Boson and the Nambu-Gotō actions are also shown.

In particular, we fit $a(L)$ in L using

$$a(L) = \left(\sqrt{1 - \frac{a_0}{\sigma L^2}} + a_1 \right) L \quad (7.4)$$

where a_1 is an undetermined bulk constant, while according to eq. (2.9) we expect $a_0 = \frac{\pi}{3} = 1.047\dots$. Results of the fits are reported in table 7.1 and the corresponding curves are displayed in fig. 7.3.

σ	a_0	a_1	χ_{red}^2
1/30	1.03(2)	-52.517(3)	0.33
1/50	1.04(7)	-98.99(1)	1.56
True	1.047...		

Table 7.1: Results for the coefficients of the fit of $a(L)$ of eq. (7.4).

We find in both fits an excellent agreement of the parameter a_0 with the expected result $\frac{\pi}{3} = 1.047\dots$. Crucially, the difference between the whole square root behaviour of eq. (7.4), a clear fingerprint of the NG action, and the free bosonic action is clearly visible in our data (see fig. 7.3). This provides

compelling evidence that other observables can be reliably and efficiently addressed using a SNF architecture.

7.4 Nambu-Gotō: string width

The natural next step in this program is to use numerical simulations to compute observables that lack an analytical description. In particular, as we mentioned above, we are interested in the string width σw^2 (see eq. (2.45)): we performed independent simulations with $R \in [30, 100]$ and $L \in [5, 20]$ at fixed $\sigma = 1/10$.

We focus our analysis on the high temperature ($R \gg L$) setting; mimicking the behaviour of eq. (2.23) we fitted the R dependence of our results with the following expression:

$$\sigma w^2(L, R) = f(L)R + g(L); \quad (7.5)$$

the results of the two fit parameters are reported in table A.2. Then, we fit again in L the coefficient $f(L)$ using

$$f(L) = \frac{1}{4L} \left(\frac{1}{\sqrt{1 - \frac{f_0}{\sigma L^2}}} + f_1 \right) \quad (7.6)$$

where f_1 is an undetermined bulk constant and f_0 , according to the conjecture mentioned in Section 2.1.2, should take the value $\pi/3$. Results of the fit are reported in table 7.2.

	f_0	f_1	χ_{red}^2
	1.01(9)	12.35(1)	0.47
True	1.047...		

Table 7.2: Results for the coefficients of the fit of $f(L)$ of eq. (7.6).

Looking at table 7.2 we see that our result $f_0 = 1.01(9)$ remarkably agrees with the conjectured coefficient $\pi/3 = 1.047\dots$. This can be appreciated in fig. 7.4 as well, where we compare




Figure 7.4: Results for the numerical prediction of the term $\sigma/\sigma(L)$, which corresponds to the coefficient $f(L)$ of the string width for $\sigma = 1/10$, after removing the divergent term f_1 and multiplying by $4L$. The solid curve corresponds to the true value from eq. (2.8).

our data for the combination $4Lf(L) - f_1$, which corresponds to the numerical prediction for the term $\sigma/\sigma(L)$, with the true value reported in eq. (2.8). Notice that the free bosonic solution would correspond to a flat horizontal line at the value of 1, which is clearly excluded by the data.

7.5 Beyond Nambu-Gotō: the string width in presence of the \mathcal{K}^2 and \mathcal{K}^4 terms

We performed numerical simulations in the presence of a \mathcal{K}^2 term both in the high- ($L \ll R$) and low- ($L \gg R$) temperature limits, where the theoretical expectations for the free bosonic string have simple analytical expressions, see eqs. (2.21) and (2.22). We set $\gamma_2 \in [-0.05, 0.1]$, we fixed $\sigma = 100$: all results

for the width are plotted in fig. 7.5 for both regimes.

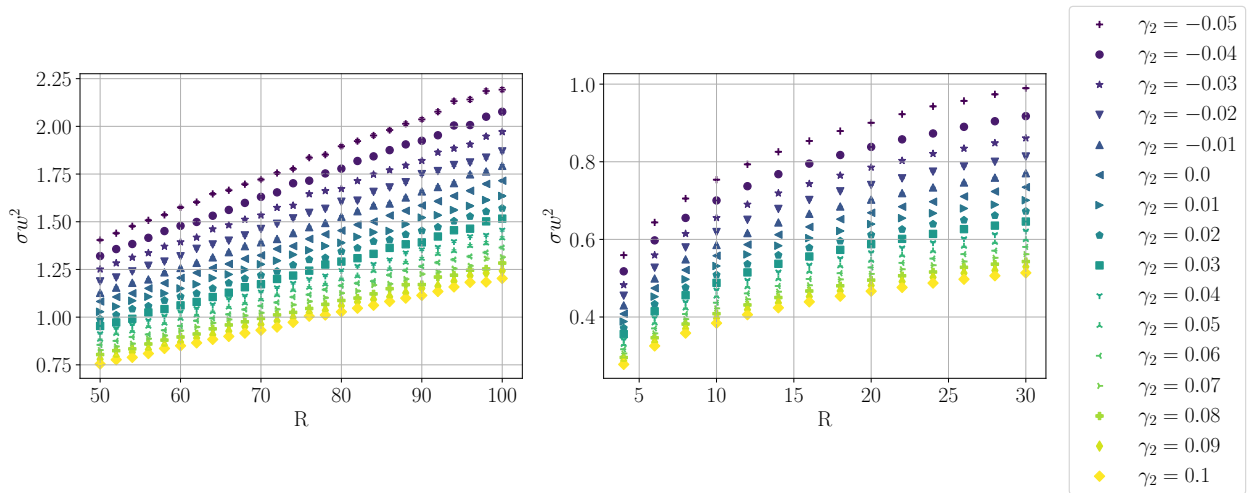


Figure 7.5: Results for the width $\sigma w^2(R, L, \gamma_2)$ at high temperature ($R \gg L = 20$, left panel) and at low temperature ($R \ll L = 80$, right panel). The $\gamma_2 = 0$ case is equivalent to the Nambu-Gotō action.

As it can be seen from both panels of fig. 7.5 the rigidity term has a strong effect on the string width. As γ_2 increases the width of the string decreases; its dependence on R changes as well and it becomes flatter. This is consistent with the qualitative picture in which the rigidity term "squeezes" the quantum fluctuations of the string: yet, the magnitude of this change is surprising. We made an attempt to analyse the contribution of the rigidity term by fitting our data with a functional form inspired by the free boson behaviour ($\sigma \rightarrow \infty$) of eqs. (2.21) and (2.22).

Studying the high-temperature regime first, we fixed $L = 20$ and chose values of the spatial extent so that $R \gg L$; then we fitted the R dependence of our results using the following functional form

$$\sigma w_{(\text{HT})}^2(R, L, \gamma_2) = a^{(\text{HT})}(\gamma_2) \frac{R}{4L} + b^{(\text{HT})}(\gamma_2) \quad (7.7)$$

This two-parameter fit works extremely well for all values of γ_2 : the results are reported in table A.3 and the parameter $a^{(\text{HT})}(\gamma_2)$

is displayed in fig. 7.6 as well. The $\gamma_2 = 0$ result for $a^{(\text{HT})}$ is quite close to 1, which is the value expected for the Free Boson action and consistent with the large value of σ . More interestingly, deviations from this value at $\gamma_2 \neq 0$ are quite remarkable and clearly visible.

The same happens in the low-temperature regime, which we realized by fixing $L = 80$ and choosing this time $R \ll L$. In this case we fitted the data with the following functional form inspired by eq. (2.21):

$$\sigma w_{(\text{LT})}^2(R, L, \gamma_2) = a^{(\text{LT})}(\gamma_2) \frac{\log R}{2\pi} + b^{(\text{LT})}(\gamma_2) + c^{(\text{LT})}(\gamma_2) \frac{1}{R^2} \quad (7.8)$$

Results for the various coefficients are reported in table A.4, with $a^{(\text{LT})}(\gamma_2)$ being plotted in the left panel of fig. 7.6 as well. Interestingly, the behaviour of the coefficient of the logarithm looks qualitatively similar to that of $a^{(\text{HT})}(\gamma_2)$ in the high-temperature limit.

Finally, we can relate the $R_c(\gamma_2)$ term as it appears in eq. (2.21) with the $b^{(\text{LT})}(\gamma_2)$ coefficient simply assuming

$$R_c(\gamma_2) = \exp(-2\pi b^{(\text{LT})}(\gamma_2)); \quad (7.9)$$

we show the values of this term in the right panel of fig. 7.6.

Interestingly, all the values of R_c (and in particular the one at $\gamma_2 = 0$ which corresponds to the pure Nambu-Gotō model) are below one lattice spacing, providing evidence that we can trust our model down to very small values of R or L . Furthermore, a value of R_c smaller than one lattice spacing is perfectly consistent with the absence in our model of an intrinsic width of the flux tube. Conversely, the value of R_c measured in non-Abelian LGTs is usually rather large, of the order of the inverse of the glueball mass. Thus, these results suggest that the value of R_c computed in non-Abelian LGTs is mostly if not entirely due to the intrinsic width, with essentially no contribution from the quantum fluctuations of the EST.

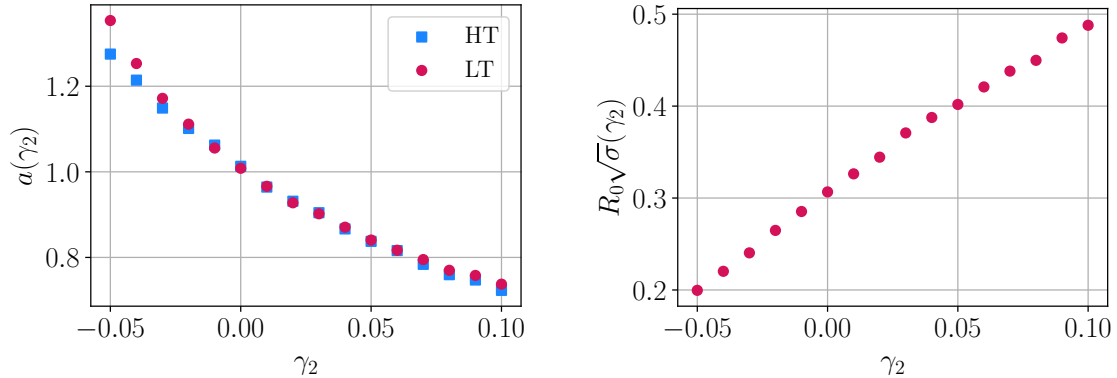


Figure 7.6: Results for $a^{(\text{HT})}(\gamma_2)$ and $a^{(\text{LT})}(\gamma_2)$ (left panel), and for $R_c(\gamma_2)$ (right panel).

Modeling the γ_2 dependence of our coefficients is beyond the scope of this work. In the spirit of our approach, we hope that our numerical results, reported in tables A.3, A.4, A.5, A.6 could be used to inspire and possibly benchmark future analytical studies of the string width in presence of the \mathcal{K}^2 correction.

We turn now to the study of the EST model described by eq. (2.49), i.e. with inclusion of the \mathcal{K}^4 correction. We performed a similar analysis, both in the high- and low-temperature limits, choosing the same values of σ , R and L and $\gamma_3 \in [-0.0005, 0.1]$. On a qualitative level, we found very similar results for the string width, which are fully reported in fig. 7.7.

As before, we fitted our data with the same functional forms discussed above, namely eqs. (7.7), (7.8) and (7.9), and we found the results reported in tables A.5 and A.6 and plotted in figure 7.8 as a function of the γ_3 coefficient.

We conclude this section by stressing a few important consequences of our analysis. The most interesting outcome of our simulations is that the effect of both the \mathcal{K}^2 or the \mathcal{K}^4 term on the string width is very strong. Conversely, in ordinary non-Abelian LGTs the behaviour of the string width is very well described

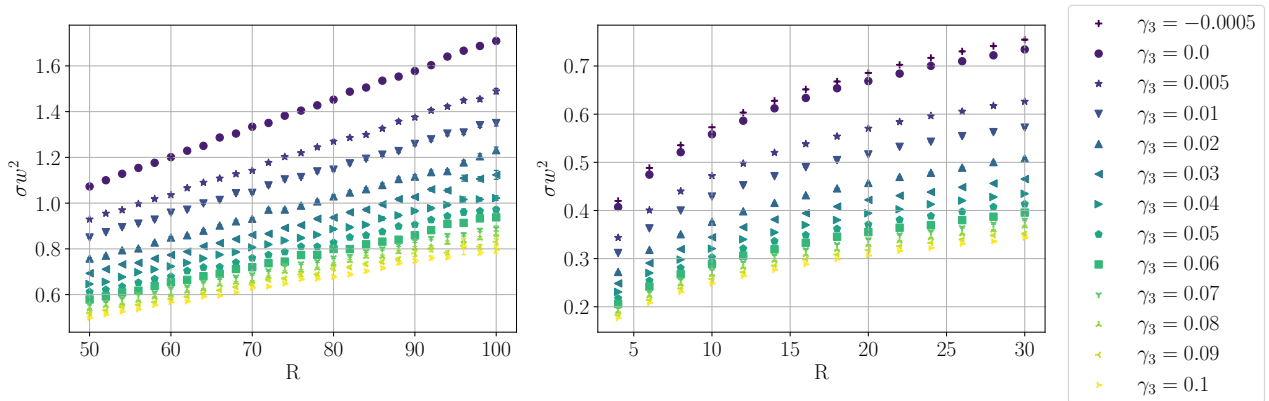


Figure 7.7: Results for the width $\sigma w^2(R, L, \gamma_3)$ at high temperature ($R \gg L = 20$, left panel) and at low temperature ($R \ll L = 80$, right panel). The $\gamma_3 = 0$ case is equivalent to the Nambu-Gotō action with no \mathcal{K}^4 term.

by the pure Nambu-Gotō prediction: our results suggest that in the EST describing the gauge theories terms beyond the NG approximation should be either absent (as it is the case for the \mathcal{K}^2 term, due to the low energy universality) or very small (as it has been recently observed for the \mathcal{K}^4 term in refs. [20, 24, 25]). At the same time it is clear that the string width represents a powerful tool to detect the possible presence of a rigidity term in other gauge theories on the lattice, such as the trace-deformed LGTs [108] or the three-dimensional U(1) model [162].

Furthermore, it is particularly interesting to study the behaviour of the critical radius R_c as a function of γ_2 and γ_3 . Looking at figures 7.6 and 7.8 we see that in both cases R_c increases as a function of γ_2 or γ_3 , but with different shapes. Modeling this behaviour might provide a tool to distinguish between the two types of corrections: in particular the critical radius seems to scale linearly with γ_2 in the \mathcal{K}^2 case. In this respect it is interesting to notice that a qualitatively similar behaviour of R_c was recently observed in the three-dimensional U(1) model [162]: in this setting confinement, thanks to the Polyakov solution [163], is indeed expected to be described by a rigid string. More pre-

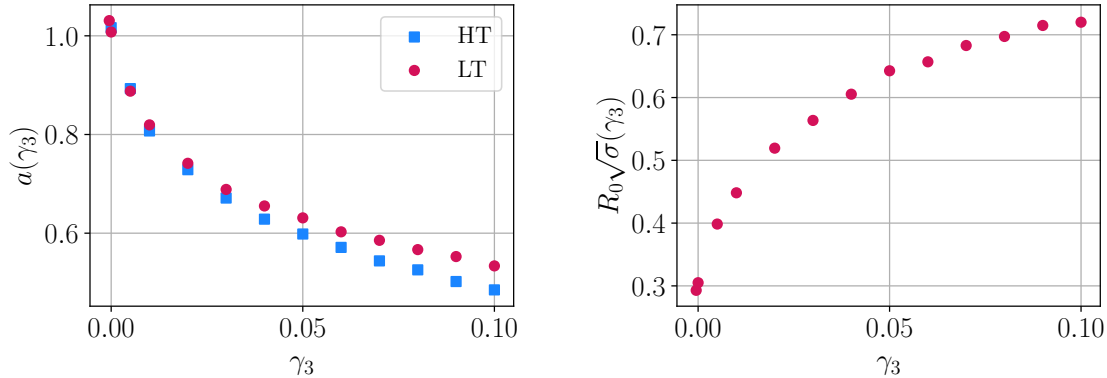


Figure 7.8: Results for $a^{(\text{HT})}(\gamma_3)$ and $a^{(\text{LT})}(\gamma_3)$ (left panel), and $R_c(\gamma_3)$ (right panel).

cisely, in this model the rigidity term is expected to become more and more important as β (the coupling constant of the model) increases and, remarkably, also R_c is observed to increase as β increases [162].

7.6 Binder Cumulant

We compute the Binder cumulant U using eq. (2.46) and we use it as a quantitative probe to understand whether the distribution of the field variable $\phi(\tau, R/2)$ (which contains info about the flux tube profile) is compatible with a Gaussian distribution (for which $U = 0$) or not.

7.6.1 The Nambu-Gotō action

In fig. 7.9 we report the results for U , for square lattices, in the case of the pure Nambu-Gotō action for decreasing values of σ and various values of the volume $L \times L$ of the lattice.

As expected, the cumulant is almost vanishing for large values of σ : in this regime we are very close to the Free Boson




Figure 7.9: Results for the Binder cumulant U as a function of $1/\sigma$ for a 20×20 lattice (right panel) and for various lattice volumes $L \times L$ (left panel) in the Nambu-Gotō lattice model.

limit and we expect the profile of the flux tube to be Gaussian. Interestingly, for decreasing values of the string tension, the cumulant becomes negative before reaching what appears to be a plateau for $\sigma \lesssim 1/3$, see left panel of fig. 7.9. This might suggest that the profile of the flux tube (i.e. our field variable ϕ) is not Gaussian anymore as soon as the Free Boson approximation is not valid. However as we increase the lattice size L this trend becomes weaker and weaker, see the right panel of fig. 7.9, suggesting that a finite-size effect due to Dirichlet boundary conditions might be responsible for a nonzero U .

A possible way to test this issue is to study the NG model in an asymmetric setting, i.e. either in the high- or in the low-temperature limit, in which either R or L is very large. In fig. 7.10 we report the results we obtained for the cumulant U as a function of $\sqrt{\sigma}R$, for several combinations of σ and R .

In the high temperature case we observe values of U different from zero only for very small values of $\sqrt{\sigma}R$. These values have large errors and fluctuate randomly around zero. This behaviour agrees with the idea that the deviations are only due to the Dirichlet boundary conditions which are almost negli-




Figure 7.10: Results for the Binder cumulant U as a function of $\sqrt{\sigma}R$ in the high-temperature limit ($R \gg L = 8$, left plot) and in the low-temperature limit ($L = 80 \gg R$, right plot). Different colors represent different values of the string tension σ .

ble in the high temperature limit. In the low temperature limit, i.e. when L (the direction along which we fix Dirichlet boundary conditions) is large and the Dirichlet b.c. play a more important role, we see a systematic deviation of the Binder cumulant for small values of $\sqrt{\sigma}R$, which is exactly the region of parameters in which the effect of the Dirichlet boundary conditions may propagate more easily in the bulk.

7.6.2 The Binder cumulant in presence of the \mathcal{K}^2 or the \mathcal{K}^4 terms

We performed a similar analysis also for the actions which include the \mathcal{K}^2 or the \mathcal{K}^4 correction terms following a strategy similar to the one used in Section 7.5. We set $\sigma = 100$ and then we varied the coefficients γ_2 or γ_3 and the lattice sizes L and R : results are reported in fig. 7.11 and in fig. 7.12 respectively. In both figures we report in the left panel a scan in the γ_2 or γ_3 coefficients for a 20×20 lattice, while in the right panel a scan in R for a fixed value of γ_2 or γ_3 and selected values of L .

In the case of the \mathcal{K}^2 action (the "rigid" string) we see essen-




Figure 7.11: Results for the Binder cumulant U in the NG model with the inclusion of the \mathcal{K}^2 term, as a function of γ_2 for a 20×20 lattice (left panel) and as a function of R for fixed $\gamma_2 = 0.02$ (right panel); in both plots $\sigma = 100$.

tially no deviations with respect to $U = 0$. On the contrary, in the \mathcal{K}^4 case a small deviation from a vanishing Binder cumulant can be detected even for rather small values of γ_3 , see fig. 7.12. Interestingly, this deviation is positive, i.e. it goes in the opposite direction with respect to what we observed in the NG case. As in the NG case its intensity decreases as R increases and thus seems to be again a boundary effect.

Let us close this section with a final important remark. In ordinary LGTs the Binder cumulant of the flux tube is typically always negative, much larger in magnitude than the values that we have found in the current work; furthermore, it shows a different dependence on R . As discussed in Section 2.4, this behavior arises from the exponential decay of the tails of the LGT flux tube shape. A quantitative explanation of its origin can be found in the intrinsic width of the confining flux tube mentioned above. Standard EST, therefore, can only describe the peak of the shape and model its width. Nonetheless, the prediction for this observable is remarkably precise (see fig. 2.4 from Section 2.4).




Figure 7.12: Results for the Binder cumulant U in the NG model with the inclusion of the \mathcal{K}^4 term, as a function of γ_3 for a 20×20 lattice (left panel) and as a function of R for fixed $\gamma_3 = 0.02$ (right panel); in both plots $\sigma = 100$.

Chapter 8

Conclusions and final remarks

In this thesis, we addressed the problem of sampling field configurations for different Effective String Theory (EST) actions regularized on the lattice. Our aim is to provide a reliable numerical framework that can work in synergy with analytical efforts to solve these models.

As discussed in Chapter 2, EST is a powerful approach to quantitatively describe confinement in pure Yang-Mills theories. However, the study of challenging observables, such as the width and shape of the effective string, imposes significant limitations on investigating LGT from the EST perspective. To tackle this issue, we introduced a new approach that involves regularizing the EST functional on the lattice, enabling the evaluation of general vacuum expectation values using numerical methods.

As shown in Chapter 3, due to the strong non-linearity of the EST model, standard MCMC approaches struggle to provide reliable estimates of relevant EST observables. Furthermore, one of the key observables in EST is the partition function, which is notoriously challenging to compute using standard MCMC methods.

To address these numerical challenges, in Chapter 4, we introduced a numerical method based on neural variational deep generative models, specifically Normalizing Flows (NFs). This

class of algorithms is able to provide a reliable approximation of a given unnormalized Boltzmann distribution (e.g., the EST action) and can be leveraged to obtain asymptotically unbiased estimations over the target ensemble. Additionally, NFs provide an efficient way to compute partition functions.

In Chapter 5, we presented a proof-of-concept application of NFs to compute partition functions and the width of the Nambu-Goto EST using a Continuous NFs. Although the numerical findings replicated the well-known behaviours from previous EST calculations, and the models used could not reach small values of the string tension, we provided the first numerical verification of the next-to-leading correction to the flux tube width, originally computed in [95–97]. Furthermore, we demonstrated that flow-based samplers are significantly more efficient than standard MCMC methods for computing EST-related observables.

In order to address the scaling issues of the flow-based samplers used in Chapter 5, in Chapter 6, we introduced an orthogonal non-equilibrium approach based on Crooks’ fluctuation theorem and Jarzynski’s equality, called Non-Equilibrium MCMC (NE-MCMC). This method can be used to efficiently compute free energy and mitigate critical slowing down. Subsequently, in Section 6.3, we combined NFs with NE-MCMC to develop Stochastic Normalizing Flows (SNFs), a powerful class of hybrid neural out-of-equilibrium algorithms. SNFs could outperform both NFs and NE-MCMC (see Section 6.3) and could afford regimes of the parameters of the Nambu-Goto theory, as well as models with additional terms proportional to quadratic and quartic powers of the extrinsic curvature, which are inaccessible to the other baselines. A natural focus of our numerical efforts was the behavior of the quantum width of the string, for which very few analytical results exist.

Our results have several important implications for LGT stud-

ies of confinement. First, we successfully measured the contribution of the terms included in the Nambu-Goto string to the string width beyond the free bosonic approximation. In particular, we confirmed a conjecture about this behavior proposed a few years ago in ref. [98].

Concerning terms beyond the Nambu-Goto action that are proportional to powers of the extrinsic curvature, we have shown that they have a significant effect on the string width. This observable clearly serves as a natural probe for identifying the role of these higher-order terms. In particular, the fact that in Monte Carlo simulations of ordinary non-Abelian LGTs, the width of the flux tube closely follows the expected Nambu-Goto behavior provides further evidence that for these models, the \mathcal{K}^2 term is absent (in agreement with the so-called “low energy universality”) and that the \mathcal{K}^4 term has a very small coefficient [20, 24, 25].

Furthermore, we have shown that a powerful tool to detect the presence of a rigidity term in the action is the critical radius R_c (the scale below which the EST predictions are not valid), which, in the presence of a $S_{\mathcal{K}^2}$ term in the action, qualitatively grows linearly with the coefficient γ_2 . See eq. (2.51), as also observed in ref. [162] for the U(1) model.

Finally, we have shown that the corrections to the Gaussian behaviour are in general very small and show a dependence on R perfectly compatible with boundary effects. Crucially, the behaviour of the Binder cumulant in LGTs is quite different from the one that we have observed for the EST models analyzed in this work. These results support a picture in which the shape of the flux tube in LGTs cannot be quantitatively described by the EST alone but requires the addition of a further contribution, the so-called “intrinsic width” [104]. In the EST language, this contribution is associated with the critical radius R_c , which identifies the scale below which the EST description is no longer valid. This finding can be clearly appre-




Figure 8.1: Comparison between the shape of the flux tube in lattice SU(2) gauge theory and Nambu-Gotō EST. The shape has been normalized by the width to absorb the dependence of the string tension and rescaled such that the peaks of both shapes coincide. The solid line represents a standard Gaussian with unit variance.

ciated in fig. 8.1, where we compared the shape of the SU(2) gauge flux tube and the Nambu-Gotō model, computed as described in sec. 2.4. The shapes have been normalized with unit variance (to absorb the effect of the string tension) and rescaled such that the peaks coincide. As evident from the plot, the peak of the SU(2) shape is Gaussian and matches the one provided by the EST. This explains why the Nambu-Gotō width prediction describes the width of the SU(2) flux tube with excellent agreement (see fig. 2.4). However, as mentioned above, a new set of ingredients is needed in order to model the intrinsic width.

We think that the same methods could be applied to several other models, both in the context of more complicated EST constructions or in the study of other classes of effective theories.

In particular, a natural candidate for further exploration is the extension of our model to two transverse degrees of freedom, which is the basic requirement for a EST description of four dimensional LGTs. In this context it would be interesting to study the so called "axionic string" which has been proposed in ref. [164] as a possible EST candidate for four dimensional LGTs. Similarly, another candidate for further numerical work would be the Polchinski-Yang limit of the rigid string [107], for which an exact solution of the free energy exists, but no information of the width and shape is available.

Regardless of the EST applications—which we emphasize once more as a rather unusual yet entirely novel and original approach—an additional objective of this thesis was to investigate the fundamentals of these new neural and non-equilibrium methods for numerical calculations in lattice field theory, particularly their combination into SNFs. Our findings indicate that, despite the low dimensionality of the EST models, these theories present a challenging setup for numerical methods. It is thus clear that SNFs provide a promising new route for scaling up lattice simulations, particularly in lattice SU(3) and QCD.

Natural future developments of this thesis include addressing critical slowing down and topological freezing in $d = 3+1$ lattice SU(3) (the former has already been studied in ref. [78]). Subsequently, a natural extension of this work is its application to the same issue in gauge theories with dynamical fermions, such as Lattice QCD itself.

Let us conclude this manuscript with a final remark. Both neural and non-equilibrium methods address the limitations of standard MCMC approaches by rethinking the estimators commonly used to measure vacuum expectation values over a desired target ensemble. In one of its most general and elegant definitions, physics is the science of what can be measured. Therefore, paradigm shifts in modeling and understanding natural phenom-

ena can arise from rethinking the very way we measure. SNFs thus serve as a prime example of this process, demonstrating how the interdisciplinary fusion of machine learning, non-equilibrium thermodynamics, and MCMC methods can lead to significant advancements in lattice field theory.

Bibliography

- [1] K.G. Wilson, *Confinement of Quarks*, *Phys. Rev. D* **10** (1974) 2445.
- [2] D.J. Gross and F. Wilczek, *Ultraviolet behavior of non-abelian gauge theories*, *Phys. Rev. Lett.* **30** (1973) 1343.
- [3] H.D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, *Phys. Rev. Lett.* **30** (1973) 1346.
- [4] O. Aharony and Z. Komargodski, *The Effective Theory of Long Strings*, *JHEP* **1305** (2013) 118 [[1302.6257](#)].
- [5] B.B. Brandt and M. Meineri, *Effective string description of confining flux tubes*, *Int. J. Mod. Phys.* **A31** (2016) 1643001 [[1603.06969](#)].
- [6] M. Caselle, *Effective String Description of the Confining Flux Tube at Finite Temperature*, *Universe* **7** (2021) 170 [[2104.10486](#)].
- [7] H.B. Meyer, *Poincare invariance in effective string theories*, *JHEP* **05** (2006) 066 [[hep-th/0602281](#)].
- [8] M. Luscher and P. Weisz, *String excitation energies in $SU(N)$ gauge theories beyond the free-string approximation*, *JHEP* **07** (2004) 014 [[hep-th/0406205](#)].

- [9] O. Aharony and E. Karzbrun, *On the effective action of confining strings*, *JHEP* **06** (2009) 012 [[0903.1927](#)].
- [10] O. Aharony and M. Dodelson, *Effective String Theory and Nonlinear Lorentz Invariance*, *JHEP* **1202** (2012) 008 [[1111.5758](#)].
- [11] S. Dubovsky, R. Flauger and V. Gorbenko, *Effective String Theory Revisited*, *JHEP* **1209** (2012) 044 [[1203.1054](#)].
- [12] M. Billo, M. Caselle, F. Gliozzi, M. Meineri and R. Pellegrini, *The Lorentz-invariant boundary action of the confining string and its universal contribution to the inter-quark potential*, *JHEP* **05** (2012) 130 [[1202.1984](#)].
- [13] F. Gliozzi and M. Meineri, *Lorentz completion of effective string (and p-brane) action*, *JHEP* **1208** (2012) 056 [[1207.2912](#)].
- [14] Y. Nambu, *Strings, Monopoles and Gauge Fields*, *Phys. Rev.* **D10** (1974) 4262.
- [15] T. Goto, *Relativistic quantum mechanics of one-dimensional mechanical continuum and subsidiary condition of dual resonance model*, *Prog.Theor.Phys.* **46** (1971) 1560.
- [16] M. Caselle, D. Fioravanti, F. Gliozzi and R. Tateo, *Quantisation of the effective string with TBA*, *JHEP* **07** (2013) 071 [[1305.1278](#)].
- [17] F.A. Smirnov and A.B. Zamolodchikov, *On space of integrable quantum field theories*, *Nucl. Phys. B* **915** (2017) 363 [[1608.05499](#)].

- [18] A. Cavaglià, S. Negro, I.M. Szécsényi and R. Tateo, *$T\bar{T}$ -deformed 2D Quantum Field Theories*, *JHEP* **10** (2016) 112 [[1608.05534](#)].
- [19] J. Elias Miró, A.L. Guerrieri, A. Hebbar, J.a. Penedones and P. Vieira, *Flux Tube S-matrix Bootstrap*, *Phys. Rev. Lett.* **123** (2019) 221602 [[1906.08098](#)].
- [20] F. Caristo, M. Caselle, N. Magnoli, A. Nada, M. Panero and A. Smecca, *Fine corrections in the effective string describing $SU(2)$ Yang-Mills theory in three dimensions*, *JHEP* **03** (2022) 115 [[2109.06212](#)].
- [21] A. Athenodorou, B. Bringoltz and M. Teper, *Closed flux tubes and their string description in $D=2+1$ $SU(N)$ gauge theories*, *JHEP* **1105** (2011) 042 [[1103.5854](#)].
- [22] S. Dubovsky, R. Flauger and V. Gorbenko, *Flux Tube Spectra from Approximate Integrability at Low Energies*, *J.Exp.Theor.Phys.* **120** (2015) 399 [[1404.0037](#)].
- [23] C. Chen, P. Conkey, S. Dubovsky and G. Hernández-Chifflet, *Undressing Confining Flux Tubes with $T\bar{T}$* , *Phys. Rev.* **D98** (2018) 114024 [[1808.01339](#)].
- [24] A. Baffigo and M. Caselle, *Ising string beyond the Nambu-Goto action*, *Phys. Rev. D* **109** (2024) 034520 [[2306.06966](#)].
- [25] M. Caselle, N. Magnoli, A. Nada, M. Panero, D. Panfalone and L. Verzichelli, *Confining strings in three-dimensional gauge theories beyond the Nambu-Gotō approximation*, *JHEP* **08** (2024) 198 [[2407.10678](#)].
- [26] M. Caselle, E. Cellini and A. Nada, *Sampling the lattice Nambu-Goto string using Continuous Normalizing Flows*, *JHEP* **02** (2024) 048 [[2307.01107](#)].

- [27] A. Radovic, M. Williams, D. Rousseau, M. Kagan, D. Bonacorsi, A. Himmel et al., *Machine learning at the energy and intensity frontiers of particle physics*, *Nature* **560** (2018) 41.
- [28] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby et al., *Machine learning and the physical sciences*, *Rev. Mod. Phys.* **91** (2019) 045002 [[1903.10563](#)].
- [29] J. Carrasquilla, *Machine Learning for Quantum Matter*, *Adv. Phys. X* **5** (2020) 1797528 [[2003.11040](#)].
- [30] A. Boehnlein et al., *Colloquium: Machine learning in nuclear physics*, *Rev. Mod. Phys.* **94** (2022) 031003 [[2112.02309](#)].
- [31] M.D. Schwartz, *Modern Machine Learning and Particle Physics*, *Harvard Data Science Review* **3** (2021) .
- [32] A. Dawid et al., *Modern applications of machine learning in quantum sciences*, 2022 [[2204.04198](#)].
- [33] D. Boyda et al., *Applications of Machine Learning to Lattice Quantum Field Theory*, in *2022 Snowmass Summer Study*, 2022 [[2202.05838](#)].
- [34] K. Zhou, L. Wang, L.-G. Pang and S. Shi, *Exploring QCD matter in extreme conditions with Machine Learning*, *Prog. Part. Nucl. Phys.* **135** (2024) 104084 [[2303.15136](#)].
- [35] K. Cranmer, G. Kanwar, S. Racanière, D.J. Rezende and P.E. Shanahan, *Advances in machine-learning-based sampling motivated by lattice quantum chromodynamics*, *Nature Rev. Phys.* **5** (2023) 526 [[2309.01156](#)].

- [36] G. Aarts, K. Fukushima, T. Hatsuda, A. Ipp, S. Shi, L. Wang et al., *Physics-driven learning for inverse problems in quantum chromodynamics*, [2501.05580](#).
- [37] A. Tomiya, *Machine learning for lattice QCD*, *Journal of the Physical Society of Japan* **94** (2025) 031006.
- [38] Y. LeCun, Y. Bengio and G. Hinton, *Deep learning*, *Nature* **521** (2015) 436.
- [39] Y. Bengio, A. Courville and P. Vincent, *Representation learning: A review and new perspectives*, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35** (2013) 1798 [[1206.5538](#)].
- [40] A. Krizhevsky, I. Sutskever and G.E. Hinton, *Imagenet classification with deep convolutional neural networks*, *Commun. ACM* **60** (2017) 84–90.
- [41] D. Silver, A. Huang, C.J. Maddison, A. Guez, L. Sifre, G. van den Driessche et al., *Mastering the game of go with deep neural networks and tree search*, *Nature* **529** (2016) 484.
- [42] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger et al., *Highly accurate protein structure prediction with alphafold*, *Nature* **596** (2021) 583.
- [43] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez et al., *Attention Is All You Need*, in *31st International Conference on Neural Information Processing Systems*, 6, 2017 [[1706.03762](#)].
- [44] M. Belkin, D. Hsu, S. Ma and S. Mandal, *Reconciling modern machine-learning practice and the classical bias–variance trade-off*, *Proceedings of the National Academy of Sciences* **116** (2019) 15849 [[1812.11118](#)].

- [45] J.W. Rocks and P. Mehta, *Memorizing without overfitting: Bias, variance, and interpolation in overparameterized models*, *Phys. Rev. Res.* **4** (2022) 013201 [[2010.13933](#)].
- [46] J. Schmidhuber, *Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook*, diploma thesis, Technische Universitat Munchen, Germany, 14 May, 1987.
- [47] L. Kirsch, J. Harrison, J. Sohl-Dickstein and L. Metz, *General-purpose in-context learning by meta-learning transformers*, [2212.04458](#).
- [48] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal et al., *Language models are few-shot learners*, in *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020 [[2005.14165](#)].
- [49] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan and S. Ganguli, *Deep unsupervised learning using nonequilibrium thermodynamics*, in *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37 of *Proceedings of Machine Learning Research*, pp. 2256–2265, PMLR, 2015 [[1503.03585](#)].
- [50] J. Ho, A. Jain and P. Abbeel, *Denoising diffusion probabilistic models*, *Advances in neural information processing systems* **33** (2020) 6840 [[2006.11239](#)].
- [51] L. Wang, G. Aarts and K. Zhou, *Diffusion models as stochastic quantization in lattice field theory*, *JHEP* **05** (2024) 060 [[2309.17082](#)].
- [52] Q. Zhu, G. Aarts, W. Wang, K. Zhou and L. Wang, *Diffusion models for lattice gauge field simulations*, in

- 38th conference on Neural Information Processing Systems*, 10, 2024 [[2410.19602](#)].
- [53] G. Aarts, D.E. Habibi, L. Wang and K. Zhou, *On learning higher-order cumulants in diffusion models*, in *38th conference on Neural Information Processing Systems*, 10, 2024 [[2410.21212](#)].
- [54] Q. Zhu, G. Aarts, W. Wang, K. Zhou and L. Wang, *Physics-Conditioned Diffusion Models for Lattice Gauge Theory*, [2502.05504](#).
- [55] D. Wu, L. Wang and P. Zhang, *Solving statistical mechanics using variational autoregressive networks*, *Physical Review Letters* **122** (2019) .
- [56] E.G. Tabak and E. Vanden-Eijnden, *Density estimation by dual ascent of the log-likelihood*, *Communications in Mathematical Sciences* **8** (2010) 217 .
- [57] E.G. Tabak and C.V. Turner, *A family of nonparametric density estimation algorithms*, *Communications on Pure and Applied Mathematics* **66** (2013) 145.
- [58] D. Rezende and S. Mohamed, *Variational inference with normalizing flows*, in *International conference on machine learning*, pp. 1530–1538, PMLR, 2015 [[1505.05770](#)].
- [59] K.A. Nicoli, S. Nakajima, N. Strodthoff, W. Samek, K.-R. Müller and P. Kessel, *Asymptotically unbiased estimation of physical observables with neural samplers*, *Phys. Rev. E* **101** (2020) 023304 [[1910.13496](#)].
- [60] M.S. Albergo, G. Kanwar and P.E. Shanahan, *Flow-based generative models for markov chain monte carlo in lattice field theory*, *Phys. Rev. D* **100** (2019) 034515.

- [61] K.A. Nicoli, C.J. Anders, L. Funcke, T. Hartung, K. Jansen, P. Kessel et al., *Estimation of thermodynamic observables in lattice field theories with deep generative models*, *Phys. Rev. Lett.* **126** (2021) 032001.
- [62] ALPHA collaboration, *Critical slowing down and error analysis in lattice QCD simulations*, *Nucl. Phys. B* **845** (2011) 93 [[1009.5228](#)].
- [63] P. Białas, P. Czarnota, P. Korcyl and T. Stebel, *Simulating first-order phase transition with hierarchical autoregressive networks*, *Phys. Rev. E* **107** (2023) 054127 [[2212.04955](#)].
- [64] I. Biazzo, D. Wu and G. Carleo, *Sparse autoregressive neural networks for classical spin systems*, *Machine Learning: Science and Technology* **5** (2024) 025074 [[2402.16579](#)].
- [65] A. Singha, E. Cellini, K.A. Nicoli, K. Jansen, S. Kühn and S. Nakajima, *Multilevel generative samplers for investigating critical phenomena*, in *The Thirteenth International Conference on Learning Representations*, 3, 2025 [[2503.08918](#)].
- [66] L. Del Debbio, J. Marsh Rossney and M. Wilson, *Efficient modeling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability*, *Phys. Rev. D* **104** (2021) 094507.
- [67] M. Gerdes, P. de Haan, C. Rainone, R. Bondesan and M.C.N. Cheng, *Learning lattice quantum field theories with equivariant continuous flows*, *SciPost Phys.* **15** (2023) 238.
- [68] A. Singha, D. Chakrabarti and V. Arora, *Conditional normalizing flow for Markov chain Monte Carlo sampling*

- in the critical region of lattice field theory*, *Phys. Rev. D* **107** (2023) 014512 [[2207.00980](#)].
- [69] S. Chen, O. Savchuk, S. Zheng, B. Chen, H. Stoecker, L. Wang et al., *Fourier-flow model generating Feynman paths*, *Phys. Rev. D* **107** (2023) 056001 [[2211.03470](#)].
- [70] D. Albandea, L. Del Debbio, P. Hernández, R. Kenway, J. Marsh Rossney and A. Ramos, *Learning trivializing flows*, *Eur. Phys. J. C* **83** (2023) 676 [[2302.08408](#)].
- [71] A. Bulgarelli, E. Cellini, K. Jansen, S. Kühn, A. Nada, S. Nakajima et al., *Flow-Based Sampling for Entanglement Entropy and the Machine Learning of Defects*, *Phys. Rev. Lett.* **134** (2025) 151601 [[2410.14466](#)].
- [72] G. Kanwar, M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, S. Racanière et al., *Equivariant flow-based sampling for lattice gauge theory*, *Phys. Rev. Lett.* **125** (2020) 121601.
- [73] D. Boyda, G. Kanwar, S. Racanière, D.J. Rezende, M.S. Albergo, K. Cranmer et al., *Sampling using $SU(N)$ gauge equivariant flows*, *Phys. Rev. D* **103** (2021) 074504 [[2008.05456](#)].
- [74] S. Bacchio, P. Kessel, S. Schaefer and L. Vaitl, *Learning trivializing gradient flows for lattice gauge theories*, *Phys. Rev. D* **107** (2023) L051504.
- [75] A. Singha, D. Chakrabarti and V. Arora, *Sampling $U(1)$ gauge theory using a retrainable conditional flow-based model*, *Phys. Rev. D* **108** (2023) 074518 [[2306.00581](#)].
- [76] R. Abbott et al., *Normalizing flows for lattice gauge theory in arbitrary space-time dimension*, [2305.02402](#).

- [77] M. Gerdes, P. de Haan, R. Bondesan and M.C.N. Cheng, *Continuous normalizing flows for lattice gauge theories*, [2410.13161](#).
- [78] A. Bulgarelli, E. Cellini and A. Nada, *Scaling of stochastic normalizing flows in $SU(3)$ lattice gauge theory*, *Phys. Rev. D* **111** (2025) 074517 [[2412.00200](#)].
- [79] R. Abbott, M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, G. Kanwar et al., *Gauge-equivariant flow models for sampling in lattice field theories with pseudofermions*, *Phys. Rev. D* **106** (2022) 074506.
- [80] R. Abbott, A. Botev, D. Boyda, D.C. Hackett, G. Kanwar, S. Racanière et al., *Applications of flow models to the generation of correlated lattice QCD ensembles*, *Phys. Rev. D* **109** (2024) 094514 [[2401.10874](#)].
- [81] R. Abbott et al., *Aspects of scaling and scalability for flow-based sampling of lattice QCD*, *Eur. Phys. J. A* **59** (2023) 257 [[2211.07541](#)].
- [82] J. Komijani and M.K. Marinkovic, *Generative models for scalar field theories: how to deal with poor scaling?*, *PoS LATTICE2022* (2023) 019 [[2301.01504](#)].
- [83] C. Jarzynski, *Nonequilibrium equality for free energy differences*, *Phys. Rev. Lett.* **78** (1997) 2690.
- [84] R.M. Neal, *Annealed importance sampling*, *Statistics and computing* **11** (2001) 125 [[physics/9803008](#)].
- [85] M. Caselle, G. Costagliola, A. Nada, M. Panero and A. Toniato, *Jarzynski's theorem for lattice gauge theory*, *Phys. Rev. D* **94** (2016) 034503.

- [86] H. Wu, J. Köhler and F. Noé, *Stochastic normalizing flows*, *Advances in Neural Information Processing Systems* **33** (2020) 5933 [[2002.06707](#)].
- [87] M. Caselle, E. Cellini, A. Nada and M. Panero, *Stochastic normalizing flows as non-equilibrium transformations*, *JHEP* **07** (2022) 015 [[2201.08862](#)].
- [88] M. Caselle, A. Nada and M. Panero, *QCD thermodynamics from lattice calculations with nonequilibrium methods: The $SU(3)$ equation of state*, *Phys. Rev. D* **98** (2018) 054513.
- [89] O. Francesconi, M. Panero and D. Preti, *Strong coupling from non-equilibrium Monte Carlo simulations*, *JHEP* **07** (2020) 233 [[2003.13734](#)].
- [90] A. Bulgarelli and M. Panero, *Entanglement entropy from non-equilibrium Monte Carlo simulations*, *JHEP* **06** (2023) 030 [[2304.03311](#)].
- [91] A. Bulgarelli and M. Panero, *Duality transformations and the entanglement entropy of gauge theories*, *JHEP* **06** (2024) 041 [[2404.01987](#)].
- [92] C. Bonanno, A. Nada and D. VDACCHINO, *Mitigating topological freezing using out-of-equilibrium simulations*, *JHEP* **04** (2024) 126 [[2402.06561](#)].
- [93] D. VDACCHINO, A. Nada and C. Bonanno, *Topological susceptibility of $SU(3)$ pure-gauge theory from out-of-equilibrium simulations*, *PoS LATTICE2024* (2025) 415 [[2411.00620](#)].
- [94] M. Caselle, E. Cellini and A. Nada, *Numerical determination of the width and shape of the effective*

- string using Stochastic Normalizing Flows*, *JHEP* **02** (2025) 090 [[2409.15937](#)].
- [95] F. Gliozzi, M. Pepe and U.J. Wiese, *Linear Broadening of the Confining String in Yang-Mills Theory at Low Temperature*, *JHEP* **01** (2011) 057 [[1010.1373](#)].
- [96] F. Gliozzi, M. Pepe and U.J. Wiese, *The Width of the Color Flux Tube at 2-Loop Order*, *JHEP* **11** (2010) 053 [[1006.2252](#)].
- [97] F. Gliozzi, M. Pepe and U.J. Wiese, *The Width of the Confining String in Yang-Mills Theory*, *Phys. Rev. Lett.* **104** (2010) 232001 [[1002.4888](#)].
- [98] M. Caselle, *Flux tube delocalization at the deconfinement point*, *JHEP* **08** (2010) 063 [[1004.3875](#)].
- [99] M. Luscher, *Symmetry Breaking Aspects of the Roughening Transition in Gauge Theories*, *Nucl.Phys.* **B180** (1981) 317.
- [100] M. Luscher, K. Symanzik and P. Weisz, *Anomalies of the Free Loop Wave Equation in the WKB Approximation*, *Nucl.Phys.* **B173** (1980) 365.
- [101] J. Polchinski and A. Strominger, *Effective string theory*, *Phys.Rev.Lett.* **67** (1991) 1681.
- [102] M. Billo and M. Caselle, *Polyakov loop correlators from D0-brane interactions in bosonic string theory*, *JHEP* **0507** (2005) 038 [[hep-th/0505201](#)].
- [103] K. Dietz and T. Filk, *On the renormalization of string functionals*, *Phys.Rev.* **D27** (1983) 2944.

- [104] M. Caselle and P. Grinza, *On the intrinsic width of the chromoelectric flux tube in finite temperature LGTs*, *JHEP* **11** (2012) 174 [[1207.6523](#)].
- [105] I. Montvay and G. Munster, *Quantum fields on a lattice*, Cambridge Monographs on Mathematical Physics, Cambridge University Press (3, 1997), [10.1017/CBO9780511470783](#).
- [106] L. Verzichelli, M. Caselle, E. Cellini, A. Nada and D. Panfalone, *Intrinsic width of the flux tube in 2+1 dimensional Yang-Mills theories*, *PoS LATTICE2024* (2025) 403 [[2501.01740](#)].
- [107] J. Polchinski and Z. Yang, *High temperature partition function of the rigid string*, *Phys. Rev. D* **46** (1992) 3667 [[hep-th/9205043](#)].
- [108] M. Unsal and L.G. Yaffe, *Center-stabilized Yang-Mills theory: Confinement and large N volume independence*, *Phys. Rev. D* **78** (2008) 065035 [[0803.0344](#)].
- [109] L. Peliti and S. Leibler, *Effects of Thermal Fluctuations on Systems with Small Surface Tension*, *Phys. Rev. Lett.* **54** (1985) 1690.
- [110] W. Helfrich, *Effect of thermal undulations on the rigidity of fluid membranes and interfaces*, *J. Phys. France* **46** (1985) 1263.
- [111] D. Förster, *On the scale dependence, due to thermal fluctuations, of the elastic properties of membranes*, *Phys. Lett.* **A114** (1986) 115.
- [112] A.M. Polyakov, *Fine Structure of Strings*, *Nucl. Phys.* **B268** (1986) 406.

-
- [113] H. Kleinert, *The Membrane Properties of Condensing Strings*, *Phys. Lett.* **B174** (1986) 335.
- [114] FLAVOUR LATTICE AVERAGING GROUP (FLAG) collaboration, *FLAG Review 2021*, *Eur. Phys. J. C* **82** (2022) 869 [[2111.09849](#)].
- [115] FLAVOUR LATTICE AVERAGING GROUP (FLAG) collaboration, *FLAG Review 2024*, [2411.04268](#).
- [116] J. Engels, J. Fingberg, F. Karsch, D. Miller and M. Weber, *Nonperturbative thermodynamics of $SU(N)$ gauge theories*, *Phys. Lett. B* **252** (1990) 625.
- [117] A.M. Ferrenberg and R.H. Swendsen, *New Monte Carlo Technique for Studying Phase Transitions*, *Phys. Rev. Lett.* **61** (1988) 2635.
- [118] P. Fearnhead, C. Nemeth, C.J. Oates and C. Sherlock, *Scalable monte carlo for bayesian learning*, [2407.12751](#).
- [119] S. Duane, A. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid monte carlo*, *Physics Letters B* **195** (1987) 216.
- [120] C.M. Bishop, *Pattern Recognition and Machine Learning*, Springer (2006).
- [121] K.A. Nicoli, *Deep generative models for thermodynamics of spin systems and field theories*, Ph.D. thesis, Technische Universität Berlin, 2023.
- [122] G.E. Box and M.E. Muller, *A note on the generation of random normal deviates*, *The annals of mathematical statistics* **29** (1958) 610.
- [123] L. Dinh, D. Krueger and Y. Bengio, *Nice: Non-linear independent components estimation*, [1410.8516](#).

- [124] L. Dinh, J. Sohl-Dickstein and S. Bengio, *Density estimation using Real NVP*, in *International Conference on Learning Representations*, 2017 [[1605.08803](#)].
- [125] I. Goodfellow, Y. Bengio and A. Courville, *Deep Learning*, MIT Press (2016).
- [126] R.T. Chen, Y. Rubanova, J. Bettencourt and D.K. Duvenaud, *Neural ordinary differential equations*, *Advances in neural information processing systems* **31** (2018) [[1806.07366](#)].
- [127] K. He, X. Zhang, S. Ren and J. Sun, *Deep residual learning for image recognition*, in *2016 IEEE Conference on Computer Vision and Pattern Recognition*, 2015 [[1512.03385](#)].
- [128] K.A. Nicoli, C.J. Anders, T. Hartung, K. Jansen, P. Kessel and S. Nakajima, *Detecting and mitigating mode-collapse for flow-based sampling of lattice field theories*, *Phys. Rev. D* **108** (2023) 114501.
- [129] A. Mnih and K. Gregor, *Neural variational inference and learning in belief networks*, in *Proceedings of the 31st International Conference on Machine Learning*, pp. 1791–1799, 2014 [[1402.0030](#)].
- [130] P. Bialas, P. Korcyl and T. Stebel, *Gradient Estimators for Normalizing Flows*, *Acta Phys. Polon. B* **55** (2024) 3 [[2202.01314](#)].
- [131] L. Vaitl, K.A. Nicoli, S. Nakajima and P. Kessel, *Path-gradient estimators for continuous normalizing flows*, in *International conference on machine learning*, pp. 21945–21959, PMLR, 2022 [[2206.09016](#)].

- [132] L. Vaitl, K.A. Nicoli, S. Nakajima and P. Kessel, *Gradients should stay on path: better estimators of the reverse-and forward kl divergence for normalizing flows*, *Machine Learning: Science and Technology* **3** (2022) 045006 [[2207.08219](#)].
- [133] L. Vaitl, L. Winkler, L. Richter and P. Kessel, *Fast and unified path gradient estimators for normalizing flows*, in *The Twelfth International Conference on Learning Representations*, 2024 [[2403.15881](#)].
- [134] M.S. Albergo, G. Kanwar, S. Racanière, D.J. Rezende, J.M. Urban, D. Boyda et al., *Flow-based sampling for fermionic lattice field theories*, *Phys. Rev. D* **104** (2021) 114507.
- [135] J. Finkenrath, *Tackling critical slowing down using global correction steps with equivariant flows: the case of the Schwinger model*, [2201.02216](#).
- [136] M.S. Albergo, D. Boyda, K. Cranmer, D.C. Hackett, G. Kanwar, S. Racanière et al., *Flow-based sampling in the lattice Schwinger model at criticality*, *Phys. Rev. D* **106** (2022) 014514 [[2202.11712](#)].
- [137] F. Noé, S. Olsson, J. Köhler and H. Wu, *Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning*, *Science* **365** (2019) eaaw1147 [[1812.01729](#)].
- [138] M. Invernizzi, A. Krämer, C. Clementi and F. Noé, *Skipping the replica exchange ladder with normalizing flows*, *The Journal of Physical Chemistry Letters* **13** (2022) 11643–11649 [[2210.14104](#)].
- [139] P. Wirnsberger, A.J. Ballard, G. Papamakarios, S. Abercrombie, S. Racanière, A. Pritzel et al., *Targeted*

- free energy estimation via learned mappings*, *The Journal of Chemical Physics* **153** (2020) [2002.04913].
- [140] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 1412.6980.
- [141] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan et al., *Pytorch: An imperative style, high-performance deep learning library*, in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035 (2019) [1912.01703].
- [142] E. Cellini and A. Nada, “Continuous Normalizing Flows for the Nambu-Goto String.” <https://github.com/TurinLatticeFieldTheoryGroup/NambuGotoCNF>, 2023.
- [143] U. Wolff, *Monte carlo errors with less errors*, *Computer Physics Communications* **156** (2004) 143.
- [144] A. Ramos, *Automatic differentiation for error analysis of monte carlo data*, *Computer Physics Communications* **238** (2019) 19.
- [145] F. Joswig, S. Kuberski, J.T. Kuhlmann and J. Neuendorf, *pyerrors: A python framework for error analysis of monte carlo data*, *Computer Physics Communications* **288** (2023) 108750.
- [146] E.M. Sevick, R. Prabhakar, S.R. Williams and D.J. Searles, *Fluctuation theorems*, *Annu. Rev. Phys. Chem.* **59** (2008) 603 [0709.3888].
- [147] D.J. Evans, E.G.D. Cohen and G.P. Morriss, *Probability of second law violations in shearing steady states*, *Phys. Rev. Lett.* **71** (1993) 2401.

- [148] G. Gallavotti and E.G.D. Cohen, *Dynamical ensembles in stationary states*, *Journal of Statistical Physics* **80** (1995) 931–970 [[chao-dyn/9501015](#)].
- [149] G. Gallavotti and E.G.D. Cohen, *Dynamical Ensembles in Nonequilibrium Statistical Mechanics*, *Phys. Rev. Lett.* **74** (1995) 2694 [[chao-dyn/9410007](#)].
- [150] G. Gallavotti, *Reversible anosov diffeomorphisms and large deviations*, *Mathematical Physics Electronic Journal* (1995) [[chao-dyn/9501011](#)].
- [151] D.J. Searles and D.J. Evans, *Fluctuation theorem for stochastic systems*, *Physical Review E* **60** (1999) 159 [[cond-mat/9901258](#)].
- [152] G. Gallavotti, *Fluctuation theorem revisited*, [cond-mat/0404699](#).
- [153] G.E. Crooks, *Nonequilibrium Measurements of Free Energy Differences for Microscopically Reversible Markovian Systems*, *Journal of Statistical Physics* **90** (1998) 1481.
- [154] S. Blaber and D.A. Sivak, *Optimal control in stochastic thermodynamics*, *J. Phys. Commun.* **7** (2023) 033001 [[2212.00706](#)].
- [155] S. Vaikuntanathan and C. Jarzynski, *Escorted free energy simulations*, *The Journal of Chemical Physics* **134** (2011) 054107.
- [156] J.P. Nilmeier, G.E. Crooks, D.D.L. Minh and J.D. Chodera, *Nonequilibrium candidate monte carlo is an efficient tool for equilibrium simulation*, *Proceedings of the National Academy of Sciences* **108** (2011) [[1105.2278](#)].

- [157] M. Arbel, A. Matthews and A. Doucet, *Annealed flow transport monte carlo*, in *Proceedings of the 38th International Conference on Machine Learning*, vol. 139, pp. 318–330, PMLR, 2021 [[2102.07501](#)].
- [158] A.G.D.G. Matthews, M. Arbel, D.J. Rezende and A. Doucet, *Continual Repeated Annealed Flow Transport Monte Carlo*, in *International Conference on Machine Learning*, pp. 15196–15219, PMLR, 1, 2022 [[2201.13117](#)].
- [159] C. Dai, J. Heng, P. Jacob and N. Whiteley, *An invitation to sequential monte carlo samplers*, *Journal of the American Statistical Association* **117** (2022) 1 [[2007.11936](#)].
- [160] L. Richter and J. Berner, *Improved sampling via learned diffusions*, in *The Twelfth International Conference on Learning Representations*, 2024 [[2307.01198](#)].
- [161] M.S. Albergo and E. Vanden-Eijnden, *Nets: A non-equilibrium transport sampler*, [2410.02711](#).
- [162] M. Caselle, M. Panero and D. Vadicchino, *Width of the flux tube in compact $U(1)$ gauge theory in three dimensions*, *JHEP* **02** (2016) 180 [[1601.07455](#)].
- [163] A.M. Polyakov, *Quark Confinement and Topology of Gauge Groups*, *Nucl. Phys.* **B120** (1977) 429.
- [164] S. Dubovsky and G. Hernandez-Chifflet, *Yang–Mills Glueballs as Closed Bosonic Strings*, *JHEP* **02** (2017) 022 [[1611.09796](#)].

Appendix A

Appendix

A.1 Free boson analytical solution

Consider the free boson CFT action regularized on a lattice Λ of size $L \times R$:

$$\begin{aligned} S_{\text{FB}}[\phi] &= \frac{1}{2} \sum_{x \in \Lambda} (\partial_\mu \phi(x))^2 = \frac{1}{2} \sum_{x \in \Lambda} \phi(x) \square \phi(x) \\ \square \phi(x) &= \sum_{\mu=1}^2 (2\phi(x) - \phi(x + \hat{\mu}) - \phi(x - \hat{\mu})) \end{aligned} \tag{A.1}$$

where $\phi(x) \in \mathbb{R}$, \square is the discrete d'Alembert operator, $\hat{\mu}$ is a unit vector pointing on the direction μ and $x = (\tau, \epsilon)$. The boundary conditions are fixed to be periodic along temporal extension L : $\phi(\tau, \epsilon) = \phi(\tau + L, \epsilon)$ and Dirichlet along R : $\phi(\tau, 0) = \phi(\tau, R) = 0$. The partition function of the theory is defined as:

$$Z = \int D[\phi] e^{-S_{\text{FB}}[\phi]} \tag{A.2}$$

where $D[\phi] = \prod_{x \in \Lambda'} d[\phi(x)]$ and Λ' is the lattice without the fixed sites of the Dirichlet boundary conditions. The action

(A.1) can be diagonalized using the orthogonal transform:

$$\phi(x) = \sum_{m=1, n=1}^{L, (R-1)} \Psi(\tau, \epsilon, m, n) y(m, n) \quad (\text{A.3})$$

with $m = 1, \dots, L$, $n = 1, \dots, R - 1$ ¹ and:

$$\Psi(\tau, \epsilon, m, n) = \frac{2}{\sqrt{LR}} \left(\cos\left(\frac{2m\pi\tau}{L}\right) + \sin\left(\frac{2m\pi\tau}{L}\right) \right) \sin\left(\frac{n\pi\tau}{R}\right) \quad (\text{A.4})$$

Substituting (A.3) in (A.1) results in:

$$\begin{aligned} S[\phi] &= \frac{1}{2} \sum_{\tau, \epsilon} \sum_{m, n} \Psi(\tau, \epsilon, m, n) y(m, n) \square \sum_{k'} \Psi(\tau, \epsilon, m', n') y(m', n') \\ &= \frac{1}{2} \sum_{m, n} \sum_{m', n'} \lambda_{(m, n), (m', n')} y(m, n) y(m', n') \end{aligned} \quad (\text{A.5})$$

Since $\lambda_{(m, n), (m', n')}$ is the element of a diagonal matrix, we can write S in term of m, n only:

$$S[\phi] = \frac{1}{2} \sum_{m, n} \lambda_{(m, n), (m, n)} y(m, n)^2 \quad (\text{A.6})$$

The eigenvalues can be computed as:

$$\begin{aligned} \lambda_{(m, n), (m, n)} \equiv \lambda_{m, n} &= \sum_{\tau, \epsilon} \Psi(\tau, \epsilon, m, n) \square \Psi(\tau, \epsilon, m, n) \\ \square \Psi(\tau, \epsilon, m, n) &= \sum_{\mu=1}^2 (2\Psi(\tau, \epsilon, m, n) + \\ &\quad - \Psi(x + \hat{\mu}, m, n) - \Psi(x - \hat{\mu}, m, n)). \end{aligned} \quad (\text{A.7})$$

The separability of the transformation (A.4) allows to compute the eigenvalues separately; for the spatial extension, using using

¹Observe that: $\sum_x \psi(x, m, R) = 0$

bisection and prosthaphaeresis:

$$\begin{aligned}
\Box\Psi(\epsilon, n) &= 2c \sin\left(\frac{n\pi\epsilon}{R}\right) - c \sin\left(\frac{n\pi(\epsilon+1)}{R}\right) - c \sin\left(\frac{n\pi(\epsilon-1)}{R}\right) \\
&= 2c \sin\left(\frac{n\pi\tau}{R}\right) - 2c \sin\left(\frac{n\pi\epsilon}{R}\right) \cos\left(\frac{n\pi}{R}\right) \\
&= 2\Psi(\tau, m) \left(1 - \cos\left(\frac{n\pi}{R}\right)\right) \\
&= 4 \sin^2\left(\frac{n\pi}{2R}\right) \Psi(\epsilon, n).
\end{aligned} \tag{A.8}$$

where for simplicity we defined $c = \frac{1}{\sqrt{LR}}$ In the case of the periodic boundary conditions:

$$\begin{aligned}
\Box\Psi(\tau, m) &= 2c \left(\cos\left(\frac{2m\pi\tau}{L}\right) + \sin\left(\frac{2m\pi\tau}{L}\right) \right) + \\
&\quad - c \left(\cos\left(\frac{2m\pi(\tau-1)}{L}\right) + \sin\left(\frac{2m\pi(\tau-1)}{L}\right) \right) + \\
&\quad - \left(\cos\left(\frac{2m\pi(\tau+1)}{L}\right) + \sin\left(\frac{2m\pi(\tau+1)}{L}\right) \right)
\end{aligned} \tag{A.9}$$

thus:

$$\begin{aligned}
&2c \cos\left(\frac{2m\pi\tau}{L}\right) - c \cos\left(\frac{2m\pi(\tau-1)}{L}\right) - c \cos\left(\frac{2m\pi(\tau+1)}{L}\right) \\
&= 2c \cos\left(\frac{2m\pi\tau}{L}\right) - 2c \cos\left(\frac{2m\pi\tau}{L}\right) \cos\left(\frac{2m\pi}{L}\right) \\
&= 2c \cos\left(\frac{2m\pi\tau}{L}\right) \left(1 - \cos\left(\frac{2m\pi}{L}\right)\right) \\
&= 4 \sin^2\left(\frac{m\pi}{L}\right) c \cos\left(\frac{2m\pi\tau}{L}\right)
\end{aligned} \tag{A.10}$$

and:

$$\begin{aligned}
& 2c \sin\left(\frac{2m\pi\tau}{L}\right) - c \sin\left(\frac{2m\pi(\tau-1)}{L}\right) - c \sin\left(\frac{2m\pi(\tau+1)}{L}\right) \\
&= 2c \sin\left(\frac{2m\pi\tau}{L}\right) - 2c \sin\left(\frac{2m\pi\tau}{L}\right) \cos\left(\frac{2m\pi}{L}\right) \\
&= 2c \sin\left(\frac{2m\pi\tau}{L}\right) \left(1 - \cos\left(\frac{2m\pi}{L}\right)\right) \\
&= 4 \sin^2\left(\frac{m\pi}{L}\right) c \sin\left(\frac{2m\pi\tau}{L}\right)
\end{aligned} \tag{A.11}$$

Finally:

$$\begin{aligned}
\Box\Psi(\tau, m) &= 4 \sin^2\left(\frac{m\pi}{L}\right) c \cos\left(\frac{2m\pi\tau}{L}\right) + 4 \sin^2\left(\frac{m\pi}{L}\right) c \sin\left(\frac{2m\pi\tau}{L}\right) \\
&= 4 \sin^2\left(\frac{m\pi}{L}\right) \Psi(\tau, m)
\end{aligned} \tag{A.12}$$

The explicit expression for the eigenvalues is found to be:

$$\lambda_{m,n} = 4 \sin^2\left(\frac{m\pi}{L}\right) + 4 \sin^2\left(\frac{n\pi}{2R}\right) \tag{A.13}$$

using Gaussian integral, the partition function can be computed as:

$$Z = \prod_{m=1, n=1}^{L, R-1} \sqrt{\frac{2\pi}{\lambda_{m,n}}}. \tag{A.14}$$

A.2 Exact sampling of the Free Boson

Given the exact eigenfunctions and eigenvalues for the free boson theory:

$$S_{\text{FB}}[\phi] = \frac{1}{2} \sum_{x \in \Lambda} (\partial_\mu \phi(x))^2 = \frac{1}{2} \sum_{x \in \Lambda} \phi(x) \square \phi(x) \quad (\text{A.15})$$

$$\square \phi(x) = \sum_{\mu=1}^2 (2\phi(x) - \phi(x + \hat{\mu}) - \phi(x - \hat{\mu}))$$

see Section 2.2.1 and Appendix A.1, it is possible to exploit them to provide an exact sampler for this theory. The sampling procedure follows the steps:

- Configurations z are sampled from a multivariate Gaussian with covariance identity, zero mean and density q_0
- Free boson samples y in the space of momenta are obtained by re-scaling the Gaussian configurations with the eigenvalues:

$$y_k = \sqrt{\frac{2}{\lambda_k}} z_k; \quad (\text{A.16})$$

the density q_1 of the new samples can be easily computed as:

$$\ln q_1(y) = \ln q_0 + \frac{1}{2} \sum_k \ln\left(\frac{\lambda_k}{2}\right) \quad (\text{A.17})$$

- Free boson samples ϕ in the configuration space are obtained using the orthonormal transformation:

$$\phi(x) = \sum_k \Psi(x, k) y(k), \quad (\text{A.18})$$

by definition of the eigenfunctions:

$$\ln q(\phi) = \ln q_1(y) \quad (\text{A.19})$$

The latter procedure have been leveraged to obtain the prior configurations for the $T\bar{T}$ -inspired SNF used in this work. Observe that this way of sampling can be imagined as an exact, non-parametric, NF.

A.3 Fits of numerical results

σ	L	$a(L)$	$b(L)$	$c(L)$	χ_{red}^2
30	10	-516.87(3)	19.8(2)	0.78(8)	0.63
30	11	-568.17(4)	22.5(3)	0.6(1)	0.89
30	12	-619.60(5)	24.4(4)	0.6(1)	1.17
30	13	-670.95(4)	27.1(4)	0.5(1)	0.90
30	14	-722.38(4)	29.3(4)	0.5(1)	0.98
30	15	-773.85(5)	31.1(4)	0.5(1)	1.05
50	11	-1080.5(1)	25.2(5)	0.6(2)	1.87
50	12	-1178.26(7)	27.6(3)	0.6(1)	0.65
50	13	-1276.04(6)	30.2(4)	0.5(1)	0.69
50	14	-1373.93(7)	32.4(3)	0.6(1)	0.58
50	15	-1471.57(7)	36.1(4)	0.2(1)	0.6

Table A.1: Results for the coefficients of the fit of $\log Z$ of eq. (7.3).

L	$f(L)$	$g(L)$	χ_{red}^2
5	0.683(3)	2.6(1)	1.12
6	0.562(3)	3.1(1)	0.85
7	0.480(2)	3.5(1)	1.78
8	0.411(2)	3.79(7)	1.12
9	0.372(1)	4.12(5)	0.57
10	0.335(1)	4.33(6)	1.35
11	0.3052(7)	4.55(3)	0.462
12	0.2779(9)	4.81(4)	0.80
13	0.2584(9)	4.89(4)	0.93
14	0.2387(7)	5.17(3)	0.64
15	0.2231(8)	5.29(4)	0.91
16	0.2092(6)	5.46(3)	0.71
17	0.1961(7)	5.60(3)	1.01
18	0.1852(8)	5.76(4)	1.16
19	0.1754(7)	5.88(3)	1.17
20	0.1669(6)	5.98(3)	0.98

Table A.2: Results for the coefficients of the fit of σw^2 of eq. (7.5).

γ_2	$a^{(\text{HT})}(\gamma_2)$	$b^{(\text{HT})}(\gamma_2)$	χ_{red}^2
-0.05	1.275(8)	0.613(7)	1.13
-0.04	1.214(6)	0.565(5)	0.90
-0.03	1.149(7)	0.535(6)	1.34
-0.02	1.101(6)	0.496(5)	1.22
-0.01	1.062(5)	0.462(4)	1.09
0.0	1.013(4)	0.445(4)	1.01
0.01	0.964(5)	0.429(4)	1.19
0.02	0.931(5)	0.408(4)	1.14
0.03	0.905(5)	0.385(4)	1.29
0.04	0.867(4)	0.375(3)	0.80
0.05	0.838(4)	0.362(4)	0.98
0.06	0.816(4)	0.344(4)	1.05
0.07	0.784(4)	0.338(4)	1.01
0.08	0.760(5)	0.327(4)	1.33
0.09	0.747(4)	0.310(3)	0.73
0.1	0.723(5)	0.304(4)	1.31

Table A.3: Results for the coefficients of the fit of eq. (7.7) in the high-temperature limit.

γ_2	$a^{(\text{LT})}(\gamma_2)$	$b^{(\text{LT})}(\gamma_2)$	$c^{(\text{LT})}(\gamma_2)$	χ_{red}^2
-0.05	1.354(5)	0.256(2)	0.07(2)	1.17
-0.04	1.253(5)	0.241(2)	0.0(2)	1.63
-0.03	1.172(3)	0.227(1)	-0.04(1)	0.77
-0.02	1.111(4)	0.211(2)	-0.03(2)	1.55
-0.01	1.056(3)	0.200(1)	-0.04(1)	0.80
0.0	1.008(3)	0.188(1)	-0.04(1)	1.16
0.01	0.966(4)	0.178(2)	-0.048(2)	2.29
0.02	0.928(2)	0.170(1)	-0.04(1)	0.74
0.03	0.902(2)	0.1579(9)	-0.015(8)	0.59
0.04	0.871(2)	0.1508(7)	-0.021(7)	0.46
0.05	0.841(3)	0.145(1)	-0.03(1)	1.28
0.06	0.817(2)	0.138(1)	-0.02(1)	1.10
0.07	0.795(2)	0.1314(8)	-0.010(7)	0.62
0.08	0.770(3)	0.127(2)	-0.02(1)	2.25
0.09	0.758(3)	0.119(1)	0.01(1)	1.66
0.1	0.738(3)	0.114(1)	0.02(1)	1.35

Table A.4: Results for the coefficients of the fit of eq. (7.8) in the low-temperature limit.

γ_3	$a^{(\text{HT})}(\gamma_3)$	$b^{(\text{HT})}(\gamma_3)$	χ_{red}^2
0.0	1.016(4)	0.440(4)	0.82
0.005	0.893(5)	0.371(4)	0.79
0.01	0.807(6)	0.348(5)	0.88
0.02	0.729(7)	0.299(5)	1.14
0.03	0.671(6)	0.272(5)	1.35
0.04	0.628(4)	0.253(4)	0.69
0.05	0.598(5)	0.234(4)	1.10
0.06	0.571(5)	0.222(4)	1.09
0.07	0.544(4)	0.217(3)	0.72
0.08	0.526(4)	0.206(3)	0.77
0.09	0.502(4)	0.205(4)	0.57
0.1	0.485(5)	0.198(4)	0.92

Table A.5: Results for the coefficients of the fit of eq. (7.7) in the high-temperature limit for a model including both the NG and \mathcal{K}^4 terms in the action.

γ_3	$a^{(\text{LT})}(\gamma_3)$	$b^{(\text{LT})}(\gamma_3)$	$c^{(\text{LT})}(\gamma_3)$	χ_{red}^2
-0.0005	1.031(4)	0.195(2)	-0.05(2)	1.08
0.0	1.008(4)	0.189(2)	-0.06(2)	1.35
0.005	0.889(3)	0.146(1)	0.02(1)	0.71
0.01	0.820(4)	0.128(2)	0.05(2)	1.02
0.02	0.742(4)	0.104(2)	0.7(1)	0.91
0.03	0.689(5)	0.091(2)	0.09(2)	2.03
0.04	0.655(4)	0.080(2)	0.11(2)	1.43
0.05	0.631(3)	0.070(1)	0.12(1)	0.90
0.06	0.603(4)	0.067(2)	0.10(2)	1.66
0.07	0.586(3)	0.061(1)	0.121(9)	0.58
0.08	0.567(3)	0.057(1)	0.107(9)	0.69
0.09	0.553(3)	0.053(1)	0.11(1)	0.91
0.10	0.534(3)	0.052(1)	0.10(1)	0.77

Table A.6: Results for the coefficients of the fit of eq. (7.8) in the low-temperature limit for a model including both the NG and \mathcal{K}^4 terms in the action.