PROCEEDINGS OF SCIENCE

Diffusion models and stochastic quantisation in lattice field theory

Gert Aarts,[a,*] **Lingxiao Wang**[b] **and Kai Zhou**[c,d]

[a] *Department of Physics, Swansea University, Swansea, SA2 8PP, United Kingdom*
[b] *Interdisciplinary Theoretical and Mathematical Sciences Program (iTHEMS), RIKEN, Wako, Saitama 351-0198, Japan*
[c] *School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), Guangdong, 518172, China*
[d] *Frankfurt Institute for Advanced Studies, Ruth Moufang Strasse 1, D-60438, Frankfurt am Main, Germany*

E-mail: g.aarts@swansea.ac.uk, lingxiao.wang@riken.jp, zhoukai@cuhk.edu.cn

Diffusion models are currently the leading generative AI approach used for image generation in e.g. DALL-E and Stable Diffusion. In this talk we relate diffusion models to stochastic quantisation in field theory and employ it to generate configurations for scalar fields on a two-dimensional lattice. We end with some speculations on possible applications.

The 41st International Symposium on Lattice Field Theory (LATTICE2024)
28 July - 3 August 2024
Liverpool, UK

*Speaker

PoS(LATTICE2024)037

1. Introduction

In recent years, a rich programme has been developed to apply methods of artificial intelligence and machine learning (AI/ML) to lattice field theories (LTFs), see e.g. Refs. [1–3]. One particular direction is the use of ML to generate LFT configurations, going beyond standard approaches, such as hybrid Monte Carlo (HMC) [4]. One reason is the notion that a well-trained ML model will generate new configurations fast, with reduced auto-correlations and possibly not suffering from critical slowing down. Evidence for this can be found in e.g. Refs. [5, 6] and can intuitively be understood as follows: in a trained model each configuration is generated starting from a fresh initial configuration, which is sampled from a simple prior, rather than from a long chain of configurations with potentially lingering correlations. Besides this important promise, it is noted that LFT is an ideal playground to learn and develop ML approaches in the context of theoretical physics.

Generally speaking, there are two schemes to devise ML algorithms to generate configurations:

- Generate configurations by approximating the (unnormalised) target distribution, $p(\phi) \sim e^{-S(\phi)}$, directly, as is done in e.g. normalising flow [5–12] and variations thereof, such as continuous normalizing flow [13–16] and stochastic normalizing flow [17, 18];
- Approximate the underlying distribution by learning from data, i.e. previously generated ensembles, as is done in e.g. GANs [19] and diffusion models, discussed here.

Recently we have introduced diffusion models in the context of LFTs. We have explored the relation between diffusion models and stochastic quantisation in scalar field theory in Refs. [20, 21] and extended this to U(1) gauge theories in Ref. [22]. In Ref. [23] we studied in detail the evolution of higher-order cumulants, encoding the interactions in field theory, and compared two popular — variance-exploding and variance-preserving or DDPM — schemes. At this conference, we also showed first results of the application of diffusion models for theories with a complex action in which configurations are generated using complex Langevin dynamics [24]. Further connections between diffusion models and field theory are pointed out in Refs. [25, 26]. In this contribution we give a high-level overview, referring to the references above for further detail.

2. Diffusion models and stochastic quantisation

Diffusion models are an extremely popular approach in Generative AI, used by e.g. DALLE-E [27] and Stable Diffusion [28]. Interestingly, the method is based on concepts in non-equilibrium physics, with one of the pioneering papers called *Deep unsupervised learning using non-equilibrium thermodynamics* [29]. Some obvious questions are:

- Can one use diffusion models in lattice field theory?
- Is there a physics connection with existing methods?
- Is the method competitive with other approaches?

The first two questions are answered positively in Ref. [20], which also contains encouraging indications for the third one. As mentioned, more details can be found in Refs. [21–24].

PoS(LATTICE2024)037

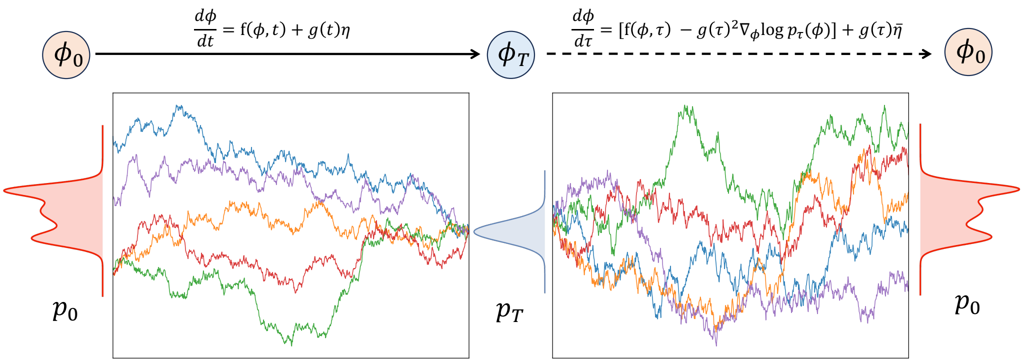

Figure 1: Sketch of the forward and backward processes in a diffusion model. The forward process, with $t = 0 \ldots T$, evolves configurations representing the target distribution p_0 to ones forming a simpler distribution p_T, while the backward process, with $\tau = T \ldots 0$, reverses this to generate new configurations ("denoising"). The additional term in the backward process, $\nabla_\phi \log p_\tau(\phi)$, is the score, which is approximated by a neural network. From Ref. [20].

Diffusion models work in combination with a previously obtained set of images or configurations, representing the target distribution $p_0(\phi)$. During the forward process, these images are made blurry or noisy, using a stochastic process. During the backward process, this is reversed and new images or configurations are created ("denoising"), starting from a normal distribution. This setup is illustrated in Fig. 1. The crucial difference between the forward and backward process is the presence of the so-called score, $\nabla_\phi \log p_\tau(\phi)$, which controls the convergence of the backward process. The score is approximated by a neural network and learnt during the forward process.

In the simplest case, with no drift applied during the forward process (i.e. $f(\phi, t) = 0$ in Fig. 1), the stochastic equations read

$$\text{forward:} \qquad \partial_t \phi(x,t) = g(t)\eta(x,t), \tag{1}$$

$$\text{backward:} \qquad \partial_\tau \phi(x,\tau) = g^2(T-\tau)\nabla_\phi \log p(\phi, T-\tau) + g(T-\tau)\eta(x,\tau). \tag{2}$$

Here $\eta(x,t) \sim \mathcal{N}(0,1)$ is Gaussian noise with variance 1, applied locally at each pixel or lattice coordinate, and $g(t)$ is the diffusion coefficient, setting the time-dependent noise strength. A common choice is $g(t) = \sigma^{t/T}$, with $\sigma \gg 1$. Compared to Fig. 1, we have redefined time in the backward process, $\tau \to T - \tau$, such that $0 \leq t, \tau \leq T$. Importantly, the time intervals are finite. The scheme with no drift, as in Eqs. (1, 2), is commonly referred to as the variance-exploding scheme, since the variance increases in time as $\mathbb{E}[\phi^2(x,t)] \sim \sigma^{2t/T}$, such that the noise will dominate the signal at the end of the forward process. A detailed analysis of the evolution of the higher-order moments and cumulants can be found in Ref. [23].

If we assume that the score follows from a time-dependent distribution,

$$p(\phi,t) = \frac{1}{Z}\exp[-S(\phi,t)] \qquad \Rightarrow \qquad \nabla_\phi \log p(\phi,t) = -\nabla S(\phi,t), \tag{3}$$

the backward process takes a familiar form

$$\partial_\tau \phi(x,\tau) = -g^2(T-\tau)\nabla S(\phi, T-\tau) + g(T-\tau)\eta(x,\tau). \tag{4}$$

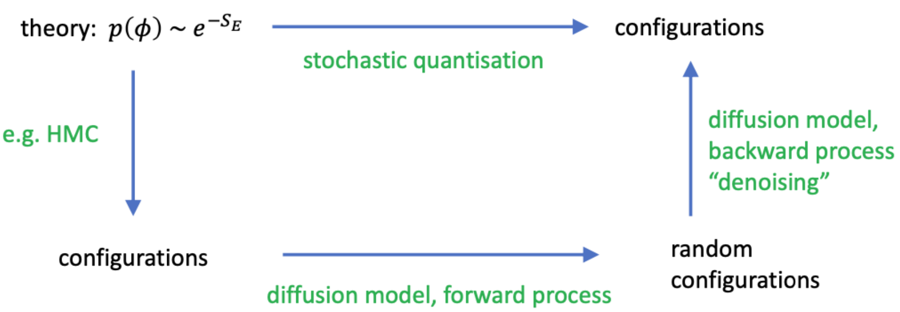

Figure 2: Flow chart indicating the relation between stochastic quantisation and diffusion models in the case of lattice field theory. Starting from the defining theory (top left), new configurations (top right) can be generated via stochastic quantisation or via the application of a diffusion model, trained on pre-existing configurations (bottom left). From Ref. [20].

One cannot help but notice that this is similar to the equation encountered in stochastic quantisation, i.e., path integral quantisation via a stochastic process in a fictitious time [30, 31],

$$\partial_\tau \phi(x,\tau) = -\nabla S(\phi,\tau) + \sqrt{2}\eta(x,\tau). \tag{5}$$

Besides the normalisation of the noise (which can be changed by rescaling the time step), we note the following:

- **stochastic quantisation**:
  – the drift is time-independent and determined by a known action;
  – the noise variance is constant (but this can be generalised using kernels [31][1]);
  – the dynamics consists of a thermalisation stage followed by evolution in equilibrium during which measurements are made.

- **diffusion models**:
  – the drift or score is not known a priori but is learnt from data;
  – the score and diffusion coefficient are time-dependent;
  – the evolution consists of many short runs ($0 \le \tau \le T$), with measurements taken at $\tau = T$;
  – correlations between runs starting from a simple prior should be absent and generated configurations can be used as proposals in a Markov chain with reduced auto-correlation.

A flow chart summarising the relation between stochastic quantisation and diffusion models is given in Fig. 2. If all algorithms are working well, the ensembles generated bottom left and top right are all representative of the target distribution $p(\phi) \sim \exp[-S(\phi)]$.

3. Two-dimensional scalar fields

We have applied the diffusion model in the variance-exploding scheme to a $\lambda\phi^4$ theory on a two-dimensional lattice, with parameter choices in the symmetric and broken phase [20]. The results shown here are obtained on a volume of 32^2. Training data is generated using Hybrid Monte Carlo;

[1] Kernels change the dynamics but leave the stationary solution unchanged; for a recent application, see e.g. Ref. [32].

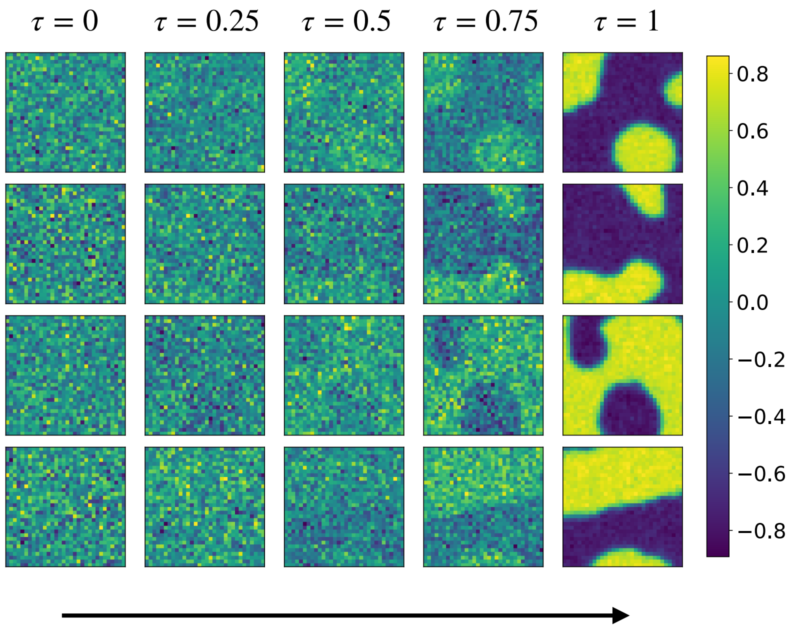

Figure 3: Denoising process in a two-dimensional scalar field theory. Four independent configurations are generated at the end of the backward process ($\tau = 1$), with clusters characteristic of the broken phase. From Ref. [20].

for the results shown here the diffusion model was trained using a U-net architecture on ensembles with 5120 independent configurations. Fig. 3 shows the denoising process in action during the backward process: four independent configurations are generated, with clusters characteristic of the broken phase appearing at the end of the backward process ($\tau = 1$). For a detailed discussion, including comparisons of the susceptibility, the Binder cumulant and higher-order cumulants, as well as correlation times and acceptance rates, we refer to Refs. [20, 23].

To illustrate how the diffusion model interpolates between the prior and the target distribution, we show in Fig. 4 the evolution during the backward process of the drift (top row) and the time-dependent action (middle row) learnt by the diffusion model in the case of a simple model with one degree of freedom, with the action

$$S(\phi) = \frac{1}{2}\mu^2\phi^2 + \frac{1}{4!}g\phi^4, \qquad \mu^2 = \pm 1, g = 0.4. \tag{6}$$

The dashed lines indicate the exact (target) results, and coloured lines show the evolution from $\tau = 0$ (blue) to $\tau = 1$ (ref). The bottom row finally shows samples generated directly from the target distribution and from the trained diffusion model. It is worth pointing out that the diffusion model can only learn where data is available, which explains the deviations seen for larger values of $|\phi|$ in the top and middle rows.

4. Outlook

In this contribution we have only shown the start of a programme to apply diffusion models to generate configurations in lattice field theory and supplement existing ensembles. Indeed, directions

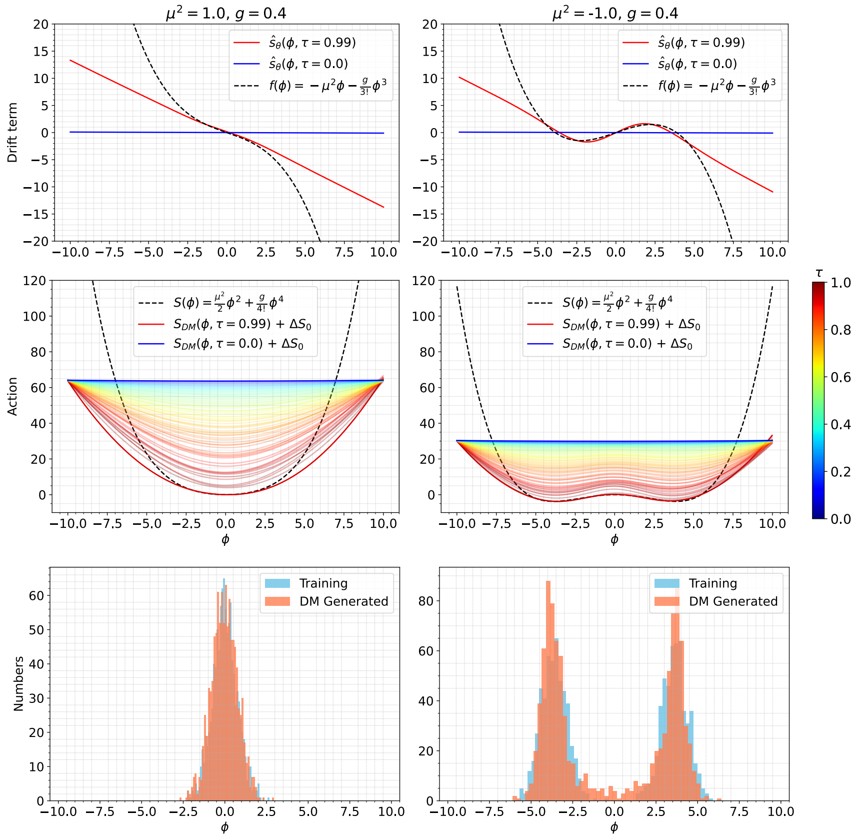

Figure 4: Toy model: Drift terms (upper row) and effective actions (middle row) learned by the diffusion model as a function of ϕ in both single-well (left column) and double-well (right column) actions, for various values of the time τ during the stochastic process. The action is shifted by a constant ΔS_0. The dashed lines indicate the exact values. The bottom row shows 1024 samples generated using the target distribution and the trained diffusion model. From Ref. [20].

to go into are plenty. Gauge theories can be included combining insights from stochastic quantisation and gauge-equivariant networks [6, 33, 34]. The first application to a U(1) gauge theory can be found in Ref. [22]. Fermions can be included implicitly, with their presence imprinted on bosonic field configurations generated in theories with fermions. An interesting direction is to apply diffusion models to theories with a sign or complex action problem, learning the (real and semi-positive) distribution from configurations generated by complex Langevin dynamics, which is not known a priori [35–38]. This is further discussed in Ref. [24]. Finally, in all cases it is important to make the algorithm exact, by including an efficient accept-reject step, and demonstrating an improvement over existing algorithms, e.g. by evading critical slowing down. Work in all these directions is currently in progress.

Acknowledgments

We thank ECT* and the ExtreMe Matter Institute EMMI at GSI, Darmstadt, for support during the ECT*/EMMI workshop *Machine learning for lattice field theory and beyond* in June 2023. GA is supported by STFC Consolidated Grant ST/X000648/1. LW thanks the DEEP-IN working group at RIKEN-iTHEMS for support. KZ is supported by the CUHK-Shenzhen University development fund under grant No. UDF01003041 and UDF03003041, and Shenzhen Peacock fund under No. 2023TC0179.

Research Data and Code Access – Details of the code and data presented can be found in Ref. [20].

References

[1] D. Boyda et al., *Applications of Machine Learning to Lattice Quantum Field Theory*, in *Snowmass 2021*, 2202.05838.

[2] 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].

[3] G. Kanwar, *Flow-based sampling for lattice field theories*, *PoS* **LATTICE2023** (2024) 114 [2401.01297].

[4] S. Duane, A.D. Kennedy, B.J. Pendleton and D. Roweth, *Hybrid Monte Carlo*, *Phys. Lett. B* **195** (1987) 216.

[5] 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 [1904.12072].

[6] 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 [2003.06413].

[7] D.J. Rezende and S. Mohamed, *Variational Inference with Normalizing Flows*, in *International conference on machine learning*, pp. 1530–1538, PMLR, 2015 [1505.05770].

[8] 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].

[9] 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].

[10] 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 [2007.07115].

[11] 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 [2105.12481].

[12] 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 [2302.14082].

[13] 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].

[14] P. de Haan, C. Rainone, M.C.N. Cheng and R. Bondesan, *Scaling Up Machine Learning For Quantum Field Theory with Equivariant Continuous Flows*, 2110.02673.

[15] 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 [2207.00283].

[16] M. Caselle, E. Cellini and A. Nada, *Sampling the lattice Nambu-Goto string using Continuous Normalizing Flows*, *JHEP* **02** (2024) 048 [2307.01107].

[17] H. Wu, J. Köhler and F. Noé, *Stochastic normalizing flows*, *Advances in Neural Information Processing Systems* **33** (2020) 5933 [2002.06707].

[18] M. Caselle, E. Cellini, A. Nada and M. Panero, *Stochastic normalizing flows as non-equilibrium transformations*, *JHEP* **07** (2022) 015 [2201.08862].

[19] K. Zhou, G. Endrődi, L.-G. Pang and H. Stöcker, *Regressive and generative neural networks for scalar field theory*, *Phys. Rev. D* **100** (2019) 011501.

[20] L. Wang, G. Aarts and K. Zhou, *Diffusion models as stochastic quantization in lattice field theory*, *JHEP* **05** (2024) 060 [2309.17082].

[21] L. Wang, G. Aarts and K. Zhou, *Generative Diffusion Models for Lattice Field Theory*, in *37th Conference on Neural Information Processing Systems*, 2023 [2311.03578].

[22] 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*, 2024 [2410.19602].

[23] 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*, 2024 [2410.21212].

[24] D.E. Habibi, G. Aarts, L. Wang and K. Zhou, *Diffusion models learn distributions generated by complex Langevin dynamics*, *PoS* **LATTICE2024** (2025) 039 [2412.01919].

[25] Y. Hirono, A. Tanaka and K. Fukushima, *Understanding Diffusion Models by Feynman's Path Integral*, 2403.11262.

[26] K. Fukushima and S. Kamata, *Stochastic quantization and diffusion models*, 2411.11297.

[27] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu and M. Chen, *Hierarchical Text-Conditional Image Generation with CLIP Latents*, 2204.06125.

[28] R. Rombach, A. Blattmann, D. Lorenz, P. Esser and B. Ommer, *High-Resolution Image Synthesis with Latent Diffusion Models*, *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022) 10684 [2112.10752].

[29] J. Sohl-Dickstein, E.A. Weiss, N. Maheswaranathan and S. Ganguli, *Deep unsupervised learning using nonequilibrium thermodynamics*, in *Proc. 32nd Int. Conf. Int. Conf. Mach. Learn. - Vol. 37*, pp. 2256–2265, 2015 [1503.03585].

[30] G. Parisi and Y.S. Wu, *Perturbation theory without gauge fixing*, *Sci. China, A* **24** (1980) 483.

[31] P.H. Damgaard and H. Hüffel, *Stochastic quantization*, *Phys. Rept.* **152** (1987) 227.

[32] D. Alvestad, R. Larsen and A. Rothkopf, *Towards learning optimized kernels for complex Langevin*, *JHEP* **04** (2023) 057 [2211.15625].

[33] 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].

[34] M. Favoni, A. Ipp, D.I. Müller and D. Schuh, *Lattice Gauge Equivariant Convolutional Neural Networks*, *Phys. Rev. Lett.* **128** (2022) 032003 [2012.12901].

[35] G. Parisi, *On complex probabilities*, *Phys. Lett. B* **131** (1983) 393.

[36] G. Aarts and I.-O. Stamatescu, *Stochastic quantization at finite chemical potential*, *JHEP* **09** (2008) 018 [0807.1597].

[37] E. Seiler, D. Sexty and I.-O. Stamatescu, *Gauge cooling in complex Langevin for QCD with heavy quarks*, *Phys. Lett. B* **723** (2013) 213 [1211.3709].

[38] G. Aarts, *Introductory lectures on lattice QCD at nonzero baryon number*, *J. Phys. Conf. Ser.* **706** (2016) 022004 [1512.05145].