This item was submitted to [Loughborough's Research Repository](#) by the author.
Items in Figshare are protected by copyright, with all rights reserved, unless otherwise indicated.

Adaptive stochastic continuation with a modified lifting procedure applied to complex systems

PLEASE CITE THE PUBLISHED VERSION

<https://doi.org/10.1103/PhysRevE.102.032210>

PUBLISHER

American Physical Society

VERSION

AM (Accepted Manuscript)

PUBLISHER STATEMENT

This paper was accepted for publication in the journal Physical Review E and the definitive published version is available at <https://doi.org/10.1103/PhysRevE.102.032210>.

LICENCE

CC BY-NC-ND 4.0

REPOSITORY RECORD

Willers, Clemens, Uwe Thiele, Andrew Archer, David JB Lloyd, and Oliver Kamps. 2020. "Adaptive Stochastic Continuation with a Modified Lifting Procedure Applied to Complex Systems". Loughborough University. <https://hdl.handle.net/2134/12901106.v1>.

Adaptive stochastic continuation with a modified lifting procedure applied to complex systems

Clemens Willers,^{1,2,*} Uwe Thiele,^{1,2,3,†} Andrew J. Archer,^{4,5,‡} David J. B. Lloyd,^{6,§} and Oliver Kamps^{2,¶}

¹*Institut für Theoretische Physik, Westfälische Wilhelms-Universität Münster,
Wilhelm Klemm Str. 9, 48149 Münster, Germany*

²*Center for Nonlinear Science (CeNoS), Westfälische Wilhelms-Universität Münster, Corrensstr. 2, 48149 Münster, Germany*

³*Center for Multiscale Theory and Computation (CMTC),*

Westfälische Wilhelms-Universität, Corrensstr. 40, 48149 Münster, Germany

⁴*Department of Mathematical Sciences, Loughborough University, Loughborough LE11 3TU, United Kingdom*

⁵*Interdisciplinary Centre for Mathematical Modelling,*

Loughborough University, Loughborough LE11 3TU, United Kingdom

⁶*Department of Mathematics, University of Surrey, Guildford, GU2 7XH, United Kingdom*

(Dated: February 6, 2020)

Many complex systems occurring in the natural or social sciences or economics are frequently described on a microscopic level, e.g., by lattice- or agent-based models. To analyse the solution and bifurcation structure of such systems on the level of macroscopic observables one has to rely on equation-free methods like stochastic continuation. Here, we investigate how to improve stochastic continuation techniques by adaptively choosing the model parameters. This allows one to obtain bifurcation diagrams quite accurately, especially near bifurcation points. We introduce lifting techniques which generate microscopic states with a naturally grown structure, which can be crucial for a reliable evaluation of macroscopic quantities. We show how to calculate fixed points of fluctuating functions by employing suitable linear fits. This procedure offers a simple measure of the statistical error. We demonstrate these improvements by applying the approach to give an analysis of (i) the Ising model in two dimensions, (ii) an active Ising model and (iii) a stochastic Swift-Hohenberg equation. We conclude by discussing the abilities and remaining problems of the technique.

I. INTRODUCTION

In many fields of science – such as physics, chemistry, biology, economics and sociology – the behaviour of various systems is captured by mathematical models that describe the time evolution of the system. The models are formulated in terms of state variables and depend on specific parameters. The system can exhibit various stable and unstable solutions, e.g., steady states, time-periodic states or chaotic behaviour. Such general classifications hold for model types ranging from discrete agent-based models to continuum models. They directly apply to deterministic models while for stochastic models they apply when considering certain statistical quantities [1–5]. In this context, we regard molecular dynamics (MD) simulations and similar systems as stochastic systems. Although formally they are deterministic, in practice the very many degrees of freedom cause deterministic chaos that renders them stochastic.

If the models contain nonlinearities, a rich variety of situations may arise, e.g., solutions may change their stability, several states exist at identical parameter values (multistability, hysteresis), the qualitative behaviour of

the system may abruptly change. All this information can be summarized in *bifurcation diagrams* that present how the solutions of the model vary when changing a control parameter. For this, a real quantity is defined as the solution measure or the order parameter which characterizes and distinguishes the different states. This quantity – which can, e.g., be a physical observable, a statistical quantity or a norm – is plotted against the control parameters [6, 7]. Here we restrict ourselves to analyzing steady states of the system, i.e., the solution measure does not need to be averaged over time.

For deterministic models consisting, e.g., of systems of ordinary or partial differential equations, a well-developed and efficient tool for calculating bifurcation diagrams is *numerical path continuation* [8–10]. With these methods one can directly follow branches of solutions and detects changes in their stability and, in consequence, bifurcations where additional solution branches emerge that one can switch onto. In contrast to procedures based on numerical time simulations, continuation can detect and follow branches of both stable and unstable states. It is used in many examples from classical dynamical systems [6, 7] to spatially extended systems such as those occurring in fluid dynamics or reaction-diffusion systems [9, 11–14].

To apply numerical continuation, one normally needs an evolution equation at the particular scale of interest in a closed form. That is, for the order parameter X and a parameter λ (we only regard one-dimensional parameter spaces), one considers a differential equation

$$\dot{X} = f_{\lambda}(X) \quad (1)$$

* clemens.willers@uni-muenster.de

† u.thiele@uni-muenster.de; <http://www.uwethiele.de>; ORCID ID: 0000-0001-7989-9271

‡ ORCID ID: 0000-0002-4706-2204

§ ORCID ID: 0000-0002-1902-0007

¶ ORCID ID: 0000-0003-0986-0878

or an iterative mapping

$$X_{t+\tau} = F_\lambda(X_t) \quad (2)$$

where τ is the time step of the mapping. Here, we are interested in steady states $X_s(\lambda)$ at different values of λ , i.e., in solutions of the equation $f_\lambda(X_s(\lambda)) = 0$ or $F_\lambda(X_s(\lambda)) = X_s(\lambda)$, respectively. Note that a differential equation can be transformed into an iterative map using time-discretisation and integration. So, generally speaking, we need to obtain the roots of a function

$$G_\lambda(X_s(\lambda)) := F_\lambda(X_s(\lambda)) - X_s(\lambda) = 0. \quad (3)$$

Hence, continuation means finding zeros of the function G_λ for subsequent values of the parameter λ . A complete bifurcation diagram is calculated by following all solution branches step by step. In every step, zeros obtained in the previous step or several previous steps can be used to predict a next value. This value is then used as initial guess for the correcting root finding procedure (see, e.g., chapter 10 of Ref. [6]). Advancing step by step in λ in this prediction-correction modus represents the simplest realisation of numerical path continuation.

However, there exist many examples where such equations are only implicitly given. In this case, one can still apply path continuation in an *equation-free* way. That is, one uses a *time stepper* instead of explicit evolution equations [15–20]. This technique is called *stochastic continuation* and makes it possible to apply continuation techniques for a much wider range of systems than those captured by Eqs. (1) or (2), including real-world experiments [21–23].

Typical examples where the equation-free approach is applied are lattice-based or agent-based models [17, 19], MD simulations [24] and kinetic Monte Carlo methods [25]. In all these examples, the model is defined on a microscopic level whereas one is interested in the dynamics of a macroscopic quantity for which an analytical description is not available. The role of the macroscopic quantity is often taken by a statistical quantity, i.e., one analyses ensembles of states. This is the origin of the term *stochastic continuation*, i.e., the method is not limited to stochastic contexts. Variants are also applied to stochastic partial differential equations [26].

In summary, stochastic continuation refers to continuation methodologies that one can apply without having direct access to the dynamics of the quantity of interest. An indirect access is made possible by three steps: *lifting*, *evolving* and *restricting*. First, in the lifting step one creates a suitable microscopic state (or an ensemble of such states) belonging to a specific value of the macroscopic quantity. Then, the state is evolved via a *time stepper*, i.e., the microscopic time evolution is advanced. Finally, in the restricting step the macroscopic quantity belonging to the newly evolved microscopic state is calculated. Taken together, these three steps realize the required time evolution on the macroscopic level and replace the function F_λ .

The present work improves several aspects of the stochastic continuation method. To begin with, in most cases, the output of the time stepper in the evolving step contains statistical fluctuations. We propose an alternative way of handling this issue. Namely, instead of directly applying a root finding method to the fluctuating function G_λ , we propose to evaluate a few (between 10 and 50) function values in a suitable neighbourhood of the initial guess and perform a linear fit; see [27]. The root of the linear fit is found to be a good approximation for the desired steady state. This procedure moderates the influence of fluctuations and works in a simple and stable manner.

Furthermore, the stochastic continuation algorithm involves quite a large number of numerical parameters, e.g., the length of the microscopic time evolution in the evolving step, the number of microscopic realisations employed and the step size of the path continuation step. The corresponding parameter settings have a crucial influence on the quality and precision of the continuation result. In Ref. [19], one possible way to systematically determine these parameters is described. It is based on an analysis of the probability distribution of the order parameter and how it changes during the microscopic time evolution step. This analysis is done during a configuration step prior to the continuation run. As a further improvement, we suggest here to adaptively adjust certain specific parameters on the fly. Especially near bifurcation points, this leads to the results being more accurate.

We also consider the lifting procedure, which is a critical part of stochastic continuation. Being a one-to-many mapping, it is not uniquely defined. In most cases described in the literature, a simple constrained random lifting procedure is used, i.e., a microscopic state with the correct macroscopic observable is chosen randomly without any control of its internal structure [15, 17, 19, 25]. Other lifting procedures found in the literature concern specific examples. In Ref. [28], a lifting technique for MD simulations of dense fluids is described where the positions and velocities of atoms are chosen according to appropriate probability distributions based on the values of the macroscopic quantities. In addition, a detailed study of lifting errors for this particular situation is given. Another example involving MD simulations (water in carbon nanotubes) is discussed in Ref. [29]. There, the lifting is done using a constrained MD simulation with an artificial bias potential. In Ref. [30] a lifting for a grand canonical Monte Carlo simulation of the Larson model for micelle formation is proposed. The lifting is based on a database of microscopic structures which are obtained by performing lengthy time simulations.

Under the assumption that there is a clear separation of the slow macroscopic and the fast microscopic time scales, possible lifting errors can be ‘healed’ during the microscopic time evolution performed during the evolving step. Considering an ensemble of microscopic states, a similar assumption is made: The ensemble is generated according to the low order moment(s) of the appropriate

distribution that are known from the macroscopic observable(s). The unknown higher order moments are expected to establish their values much more rapidly than the low order ones. Through this slaving, errors made in forming the ensemble are healed [19, 25, 31]. The healing process is supposed to constitute a minor part of the algorithm's computation time.

However, in many important cases, the microscopic states show an internal spatial structure that is important for the macroscopic dynamics (see Sec. III). The healing process does develop these internal structures, but this can sometimes only occur after a very long microscopic time evolution, i.e., costing a huge numerical effort. In such cases, stochastic continuation has no advantage for stable states over calculations based on just direct time simulations.

In the literature, various ways to tackle this problem are presented. Ref. [32] uses an algorithm which creates initial states for the evolving step that are very close to a *slow manifold* of the system. However, this only works for specific coupled differential equations. In Ref. [33] a single suitable fixed reference state is used and adjusted according to the values of the macroscopic quantities. However, this does not allow the character of the microscopic states to change over the course of the continuation process and results in similar problems as a constrained random lifting. The lifting techniques used for the MD simulation of water in carbon nanotubes [29] and of micelle formation [30] mentioned above, are able to produce naturally grown microscopic structures. Yet, they require a vast numerical effort because they are based on long time simulations at every single lifting step and on the creation of a database of clustered structures, respectively.

Here, we propose a solution to this problem that consists of a lifting technique which uses a flexibly varying reference state. This facilitates the creation of suitably structured subsequent microscopic states for different values of the macroscopic quantity as one follows the continuation path. In particular, the continuation of stable branches gives much more accurate results with this procedure, which we refer to as *structure lifting*.

After introducing our proposed improvements to the stochastic continuation procedure, we illustrate the approach by applying it to three distinct systems and discuss the results. This allows to focus on the strengths and weaknesses of our method, but we do not give an exhaustive analysis of the individual example systems.

The first system we apply the approach to is the Ising model in two dimensions (2d). This is a simple model for ferromagnetism [34]. It consists of a 2d lattice containing N sites that are each occupied by a spin which can be in one of two possible states, $s_i \in \{-1, 1\}$, with $i = 1, \dots, N$. Every spin interacts with its four nearest neighbours with *interaction strength* J . The Hamiltonian reads

$$H(s) = - \sum_{nn} J s_i s_j \quad (4)$$

where $s = \{s_1, s_2, \dots, s_N\}$ and $\sum_{nn}$ denotes the sum nearest neighbour pairs. The macroscopic observable is the magnetisation m which is the mean value of any of the spins, i.e., $m := \langle s_i \rangle$ where $\langle s_i \rangle$ denotes the time (or ensemble) average value of s_i . The microscopic dynamics of the model is given by the usual Metropolis Monte Carlo algorithm [35]. One is interested in the stable and unstable states, which are characterized by their magnetisation value, which is a function of the reduced temperature $T' := k_B T / J$, where T is the temperature and k_B is Boltzmann's constant. Onsager's analytical solution shows that above a critical temperature T_c the magnetisation is zero while the system shows a spontaneous magnetisation below T_c [36].

The second example we consider is the active Ising model introduced in Ref. [37]. Here, the spins can be associated with particles that not only can flip their spin state but can also move. The preferred direction of motion of a spin depends on its orientation. Depending on the mean particle density, different phases are observed. At low densities, the mean orientation and the mean velocity are zero. This is called the “gas phase”. At high densities, the spins are ordered, i.e., there emerges an overall magnetisation which is connected to an overall directed motion. This is called the “liquid phase”. At intermediate densities, a phase separation occurs, into different regions containing the gas and the liquid states. The total fraction containing the liquid can be used as the macroscopic order parameter (see Sec. IV for details).

Finally, as third example we perform stochastic continuation of the steady inhomogeneous solutions of a stochastic partial differential equation, namely, a stochastic version of the Swift-Hohenberg equation [38]. It incorporates an additive noise field η multiplied by a parameter D which determines the strength of the noise term. The time evolution equation for the field $\Psi(x, y, t)$, which varies with position (x, y) and time t , is

$$\partial_t \Psi = \varepsilon \Psi - \Psi^3 - (1 + \nabla^2)^2 \Psi + \sqrt{D} \eta. \quad (5)$$

Different implementations of stochasticity (i.e. the noise field η) are employed in the literature: One finds additive Gaussian spatiotemporal noise [39–41], independent Gaussian additive and multiplicative noise [42], and spatially global Gaussian noise [38]. Here, we take it to be normally distributed and uncorrelated in space and time. Depending on the control parameter ε and the magnitude of D , this equation has various stable steady states which correspond, e.g., to stripe or square patterns of the field Ψ . A macroscopic quantity which can be used as a solution measure (order parameter) in a bifurcation diagram is the L^2 -norm of the deviations of Ψ from its mean value.

This paper is organized as follows: In Sec. II, we introduce the details of stochastic continuation and the corresponding notations. Additionally, we explain some difficulties with the method and present our proposed improvements. We illustrate the improved method with three examples: the classical Ising model in 2d in Sec. III,

an active Ising model in Sec. IV, and a stochastic Swift-Hohenberg equation in Sec. V. Finally, in Sec. VI, we summarize the capabilities of the proposed approach and discuss some remaining questions.

II. METHOD

A. Numerical continuation and the pseudo-arclength method

The method of stochastic continuation extends the standard equation-based continuation method for deterministic systems. Therefore, we start by briefly explaining the essence of continuation methods for one-dimensional deterministic systems [8–10]. These have a state variable (or order parameter) X that is a real macroscopic quantity. The dynamics is defined by an ordinary differential equation (ODE) of the form in Eq. (1) or by an iterative map of the form in Eq. (2), which depend on a control parameter λ (we only consider one-dimensional parameter spaces). Note that a dynamics described by a partial differential equation (PDE) or an integro-differential equation can be transformed into a system of ODEs, e.g., by discretising in real space using finite-differences or spectral methods. Further, one may employ time-discretisation and integration to transform a differential equation into an iterative map. Therefore, here we only consider the latter.

Of central interest is how steady states $X_s(\lambda)$ depend on the value of the control parameter λ ; i.e. we neglect time-periodic and chaotic states. In other words, the aim is to search for states which satisfy Eq. (3). Presented in a bifurcation diagram as a function of λ , these zeros form different branches (this is a consequence of the implicit function theorem [8]). Applying the continuation method, each branch can be calculated by following it step by step. In every step, a predictor-corrector scheme is applied, whereby different prediction and correction methods are used in the literature.

In the simplest version of continuation, one explicitly changes the parameter λ in small steps. The corresponding steady states $X_s(\lambda)$ are obtained by a root finding procedure applied to the function G_λ in Eq. (3) for a fixed value of λ . In this approach the solution at each step in λ can be obtained by using the solution from the previous step as the initial guess for the solution at the next step. Thinking of this as a ‘prediction’ followed by a ‘correction’, then the prediction step is trivial, and is represented by a horizontal line in the bifurcation diagram, while the corrections correspond to vertical displacements.

Taking two previous steps into account and using a linear extrapolation, it is possible to use this secant predictor to obtain a better initial state in the predictor step. Again, the corrector is a root finding process at fixed λ . The correction at fixed λ is problematic at saddle-node bifurcations because one cannot follow the bifurcation

FIG. 1. Sketch illustrating the basis of numerical continuation. Shown are schemes consisting of a secant predictor and (i) pseudo-arclength corrector and (ii) orthogonal corrector. The secant predictor uses two previous points on the solution branch to obtain an initial guess for the corrector via a linear extrapolation by an arc length ds . The pseudo-arclength corrector employs Newton steps on a circle of radius ds around the previous point to obtain the next point on the branch. The orthogonal corrector searches on a line perpendicular to the direction of the secant.

curve through the turning point (or fold). Instead, the predictor goes beyond the bifurcation point and the corrector either does not converge to any steady solution or converges to a solution on a branch different to the desired one.

To follow a solution branch through a fold, one may instead use the pseudo-arclength method [43]. In this method, a branch is not tracked using the control parameter λ as the main continuation parameter. Instead, an intrinsic and unique property of the bifurcation curve, the arc length s along the branch, is used. In particular, one employs a secant predictor advancing an arc length ds . Then, the predictor consists of Newton steps at fixed arc length, i.e., the true position of the next point on the branch is searched along a circle with radius ds around the previous point, which is sketched as option (i) in Fig. 1. That is, the next point is not searched for at fixed λ but at fixed s , while λ is adapted in the Newton steps and its exact value becomes part of the solution. With the resulting predictor-corrector scheme, it is possible to follow a branch through a saddle-node bifurcation.

Another alternative is to make the correction along a line which is perpendicular to the direction defined by the secant predictor, as sketched as option (ii) in Fig. 1. We refer to this scheme a secant predictor with an orthogonal corrector and use it below in our continuation algorithm. It is simpler to realise than the pseudo-arclength predictor-corrector method and approximates it with a negligible error.

We parameterize the perpendicular line on which the correction is done with a real number u such that the line is given by the set of points (λ_u, X_u) . We search for a zero of the function G_λ along this line. That is, we search for the zero of the real-valued function $\tilde{G}(u) := G_{\lambda_u}(X_u)$.

If we denote a stable steady state as X^* , then when the

state of the system $X < X^*$, we have $X < F_\lambda(X)$ and so $G_\lambda(X)$ is positive. On the other hand, when $X > X^*$, then $G_\lambda(X)$ is negative. Hence, near the steady state, the function G_λ has a negative slope. Near an unstable steady state, it has a positive slope.

We cannot transfer this idea to the function $\tilde{G}$ directly, because, on moving along a curved solution branch, the perpendicular line can have many different orientations. Only for a horizontal segment of a branch and a vertical perpendicular line can the above consideration be directly applied to the function $\tilde{G}$. The meaning of the sign of the slope of the function $\tilde{G}$ may change along curved branches¹. Nevertheless, via the orientation of the perpendicular line, there is a connection between this sign and the stability of the solutions on the branch being continued. So, if a change in the sign of $\tilde{G}$ is detected during the continuation, one has to be aware that a bifurcation point may have been passed and consider the possibly of other branches existing in the bifurcation diagram. This is discussed further in the outlook in Sec. VI.

The numerical continuation schemes just described form a well-developed method to calculate bifurcation diagrams for deterministic systems formulated as ODEs, PDEs or as iterative maps. Many different software packages exist which implement this method, like AUTO-07P [43–45] or PDE2PATH [46]. Furthermore, these contain many additional helpful and sophisticated tools that enable the detection of many types of bifurcations, branch switching to steady and time-periodic states, two-parameter continuation of loci of bifurcation points, like saddle-node, Hopf and period doubling bifurcations. This makes it possible to efficiently determine complex bifurcation diagrams.

B. Stochastic continuation

The continuation methods described in the previous subsection are well suited to investigating deterministic systems. We now explain how the basic concepts can be applied for stochastic systems. For such systems we must distinguish between the macroscopic and the microscopic descriptions and how these are related. In general, we use the terms “macroscopic quantity” and “macroscopic dynamics” for the quantity and dynamics we are interested in on the level of a macroscopic or ensemble-averaged description. These are the quantities that characterise the steady states in the context of bifurcation diagrams. By “microscopic quantity” and “microscopic state” we

FIG. 2. Scheme indicating how the lifting, evolving and restricting steps are used to evaluate a function call of F_λ in the method of stochastic continuation.

refer to individual representations of the stochastic system. Note that the terms “microscopic” and “macroscopic” may not be referring literally to length scales in the system, although of course they can be. For instance, the macroscopic state can be the mean value of a function of the microscopic state (e.g., the mean magnetisation of the spin lattice in the Ising model) or the mean over an ensemble of microscopic states (e.g., a time-average of a quantity described by a Langevin equation).

In the numerical continuation for deterministic systems described above, an equation defining the macroscopic dynamics is explicitly given. However, for a stochastic system the macroscopic dynamics is given only implicitly and the connection between the microscopic and the macroscopic level is often (but not always) of a statistical nature, since it arises from the stochastic dynamics of the microscopic states. Numerical continuation for stochastic systems, referred to as “stochastic continuation”, continues the steady states of the macroscopic quantities and determines their bifurcation diagrams without explicit knowledge of the time evolution of the macroscopic state. Note that by “steady states” of a stochastic system, we mean that the macroscopic quantity obtained by averaging is approximately steady. While microscopic states are not steady in a stochastic system, the solution measure or order parameter used to distinguish different types of microscopic states is always obtained by some averaging procedure (e.g., the mean orientation of all spins located on the Ising lattice at a particular time). Because in practice systems are always of finite size, the macroscopic quantity still slightly fluctuates. This important issue we will further discuss in section II C.

In principle, numerical continuation can be done in all cases where one has access to the function F_λ . For a stochastic system this function is not explicitly given, but one can instead use the microscopic time-stepper; see [15–19]. To do this we have to determine a way to create for each value of the macroscopic quantity X a corresponding microscopic state x or an ensemble of such states $\{x^i\}$. Using such a mapping, one is able to bypass the problem of not having an explicit form for the function F_λ . To obtain the value of the function F_λ at

¹ A straight line can be parameterized into two different directions. So, to be specific, the relation between the sign of the slope of the function G_λ and the stability of the steady state depends on the direction of the parameterization. For one direction, we have stability for a positive slope, while for the other direction it is the negative slope.

X_t , we apply the following three steps: First, a corresponding microscopic state x_t or a corresponding ensemble $\{x_t^i\}$ is created in the so-called *lifting* step. Second, the time evolution is advanced a time interval τ employing the explicitly known microscopic dynamics to obtain the state $x_{t+\tau}$ or the ensemble $\{x_{t+\tau}^i\}$. This is the *evolving* step. Finally, in the *restricting* step, one returns to the macroscopic level by evaluating $X_{t+\tau}$ which is identical to $F_\lambda(X_t)$. This sequence of steps is shown schematically in Fig. 2 and bypasses the task of evaluating the unknown function $F_\lambda(X_t)$. Thus, every time the numerical continuation algorithm calls the macroscopic function F_λ in the procedure to find the roots of the function G_λ or in any other situation, the call of F_λ is replaced by the three-step bypass illustrated in Fig. 2.

Note that the algorithm just described includes two important additional parameters. The first is the time interval τ of the microscopic time evolution. The second is the size M of the microscopic ensemble. Rules for the choice of these are discussed below in Sec. II C 1 and Sec. II C 2, respectively.

Another important issue relates to the specific definition of the lifting map that is used. This depends on the specific example being considered and its microscopic dynamics. We present some general ideas for how to do this in Sec. II C 3 below.

A further issue for stochastic systems relates to the fact that $\tilde{G}$ is also a fluctuating quantity and so to find a zero of $\tilde{G}$ one must deal with the statistical fluctuations. Here we use the C³R method introduced in Ref. [19] to prevent the occurrence of outliers which can lead to poor continuation results. In this method, every approximation which lies outside of a suitably defined interval around the initial guess is repeated a predefined number of times. This greatly reduces the probability of outliers from statistical fluctuations leading to poor results. This is used in combination with the pseudo-adaptive parameter adjustment introduced in the next section.

C. Challenges of Stochastic Continuation – Towards an effective procedure

The stochastic continuation approach presented in the previous section faces a number of important challenges. Here, we discuss these and give methods to remedy the difficulties.

1. Fluctuations in $\tilde{G}$ – Using linear fits

When dealing with stochastic systems, the system size and the ensemble size that can be simulated are always finite. As a consequence, the function $\tilde{G}$ exhibits statistical fluctuations which complicate the root finding and induces uncertainties.

Here, we solve this problem by employing linear fits. Namely, we evaluate a finite number N_{fit} of function val-

FIG. 3. To determine the next point on the branch of steady solutions, the function $\tilde{G}$ is evaluated at several points within an interval around the initial guess (indicated by small crosses) on the line perpendicular to the line defined by the secant predictor. The length of the interval is controlled by the parameter σ that is chosen in such a way that the function $\tilde{G}$ is sufficiently linear and that the required root is contained in it.

ues of $\tilde{G}$ equidistantly spaced on the line perpendicular to the secant prediction that defines the correction step in the orthogonal corrector – see Sec. II A. We use here $10 \leq N_{\text{fit}} \leq 50$. This is done up to a suitable distance from the initial guess, as illustrated in Fig. 3. This yields the function $\tilde{G}$ evaluated in the form of a cloud of points. The cloud is then approximated by a linear best-fit whose root is then employed as the approximation for the zero of $\tilde{G}$. A typical example of this is displayed in Fig. 4. This procedure is very robust if an appropriate range for the evaluation of $\tilde{G}$ and the fitting is defined. Of course, the function $\tilde{G}$ should be sufficiently linear over this range. Thus, on the one hand, the range considered should be sufficiently small for this to be true, but on the other hand, the root should lie within this range, because an extrapolation beyond the range considered can lead to poorly controlled uncertainty. As a result, we must choose a step size for the prediction step to be sufficiently small so that the initial guess for the correction and the zero of $\tilde{G}$ are sufficiently close together. We define a parameter σ that specifies the range to be considered, as indicated in Fig. 3. A good value turns out to be $\sigma = 0.5ds$, as we show next.

The fitting procedure works in a stable way because it moderates the fluctuations of the function $\tilde{G}$. Additionally, it gives the advantage of allowing us to estimate the size of the statistical error $\pm\Delta$ in the calculated zero. As a measure for Δ we use the standard deviation of the horizontal distances of all the calculated points to the fitted line – see Fig. 4. With N evaluated values of the function $\tilde{G}$ with the abscissae $u_1, \dots, u_N$ and the ordinates $\tilde{g}_1, \dots, \tilde{g}_N$ and the linear fit $\tilde{G} = \theta u + b$, we obtain

$$\Delta := \sqrt{\frac{1}{N} \sum_{i=1}^N \left(\frac{\tilde{g}_i - b}{\theta} - u_i \right)^2}. \quad (6)$$

Without our fitting technique, simple numerical root finding methods generally fail, because of the fluctuations

FIG. 4. The plot shows a typical point cloud obtained through evaluations of the function $\tilde{G}(u) = G_{\lambda_u}(X_u)$ on an interval as illustrated in Fig. 3 with $\sigma = 0.05$. The point cloud is shown together with the resulting linear fit used to determine the zero of the fluctuating function. The standard deviation of the horizontal distances of all points to the fit line is indicated as a horizontal red bar. It serves as a measure for the statistical error in the zero. In this particular example, the data is from stochastic continuation of the branch stable steady states of the Ising model at the reduced temperature $T' = 3.1$. Here, T' takes the role of λ . The size of the spin lattice is 100×100 while the parameters of the continuation algorithm are $\tau = 496$, $M = 39$, $\sigma = 0.05$, $ds = 0.1$ and $N_{\text{fit}} = 20$. The values of τ , M and σ are obtained by our adaptive parameter choice (as explained in Sec. II C 2) with $\Delta_d = 0.005$ and $\theta_d = 1.0$. The slope of the fit is $\theta = 0.93$ and the error in the zero is $\Delta = 0.0037$.

of $\tilde{G}$. Independently of the particular fitting technique, it is possible to determine the zeros by applying more advanced root finding methods which take derivatives of the considered function into account [19].

2. Divergence of time scales – Adaptive choice of numerical parameters

Roughly speaking, one can view the dynamics of a typical system to be considered as being composed of a deterministic part together with a stochastic part. The deterministic part controls the slope θ of the function $\tilde{G}$, while the stochastic part causes it to fluctuate. The larger the influence of the deterministic part, the larger is the absolute value of the slope $|\theta|$. For small $|\theta|$, fluctuations dominate and any information we have about the location of the zero is not reliable (cf. the case in Fig. 5). For large $|\theta|$, the uncertainties are small (cf. the case in Fig. 4).

It is well known that on approaching a bifurcation point (sometimes called a “tipping point” [9]), the typical time scales of the system dynamics diverge – often referred to as “critical slowing down”. This implies that if a fixed time span τ is used for the microscopic time evolution, the slope of $\tilde{G}$ tends to zero when approaching the

FIG. 5. Comparing this figure with Fig. 4 shows the influence of the size of the time span τ used for the microscopic time evolution. Here, the point cloud is obtained from evaluating the function $\tilde{G}(u) = G_{\lambda_u}(X_u)$ with $\tau = 10$, which is too small. All the other parameters are as in Fig. 4. The slope of the straight-line fit is $\theta = 0.09$, the error in the zero is $\Delta = 0.0208$.

bifurcation point. This causes a significant loss of precision for these parameter values, which is unfortunate, because this concerns the (generally) most important details of a bifurcation diagram.

Our solution for remedying this issue consist of dynamically adjusting the value of τ to the time scale typical for the system at any given point in the bifurcation diagram. In other words, when the dynamics of the macroscopic observable slows down the microscopic time evolution is prolonged by adapting τ . For example, Fig. 5 displays results for $\tilde{G}$ with $\tau = 10$, which is too small while Fig. 4 shows the corresponding result obtained by adapting τ to 496.

Based on our accumulated experience, it is possible to define suitable values of τ “by hand”. We call this approach a *pseudo adaptive* continuation. We have done this for the first continuation results of the Ising example (cf. Sec. III). Combined with the C³R method and the lifting procedure presented in the next section, very good results can be obtained. However, the pseudo adaptive method is very labour intensive and depends greatly on the user’s experience. This can be avoided by using the fully adaptive adjustment of the parameter τ approach discussed next.

Furthermore, due to the law of large numbers, the size of the microscopic ensemble M also has an influence on the magnitude of fluctuations in $\tilde{G}$ and the precision of the continuation results. For example, Fig. 6 shows $\tilde{G}$ evaluated for $M = 2$, which is too small. This should be compared with the corresponding results obtained from adapting M to 39 displayed in Fig. 4. This comparison illustrates that the parameters τ and M greatly influence the exactness of the method. They also have a large effect on the numerical cost, since the ensemble microscopic time evolutions usually makes the largest contribution to

FIG. 6. Comparing this figure with Fig. 4 shows the influence of the size of the microscopic ensemble M used for the microscopic time evolution. Here, the point cloud is obtained from evaluating $\tilde{G}(u) = G_{\lambda_u}(X_u)$ with $M = 2$, which is too small. All the other parameters are as in Fig. 4. The slope of the fit is $\theta = 1.08$, the error in the zero is $\Delta = 0.0216$.

the computation time. Thus, these parameters should be chosen to be as small as possible and as large as needed. Note also that the parameter σ discussed in Sec. II C 1 is important in this context.

Our proposed adaptive adjustment takes the three parameters τ , M and σ into account. For every parameter, we specify a condition which should be fulfilled to effect a reliable determination of the root of the function $\tilde{G}$ and we implement a rule for modifying the parameter values when these conditions are not fulfilled. For the parameters τ and M we also define upper limits ($\tau_{\max}$ and $M_{\max}$) that reflect the available computing resources. Since these rules are based on estimates, they are iteratively applied until all of the conditions are met. As initial values for the parameters, we use the adjusted parameter values from the previous continuation step. To make the adjustment of τ more efficient, we initially extrapolate the change between the previous two continuation steps. Namely, if the adaptive adjustments in the previously performed continuation steps produced the series of values $\tau_0, \tau_1, \dots, \tau_{n-1}$, then we use

$$\tau_n^{\text{guess}} := \tau_{n-1} + (\tau_{n-1} - \tau_{n-2}) \quad (7)$$

as initial value for τ .

It is useful to start the continuation procedure with low values of τ and M in a region of the bifurcation diagram where no complications are expected, e.g., on a stable branch far away from bifurcation points, and to let the adaptive control work while approaching bifurcation points. Figure 7 gives a flow chart that illustrates how our algorithm adjusts the various parameters. This is discussed below in the order that the parameters are updated.

Most importantly, the interval σ in which we fit $\tilde{G}(u)$ must be appropriate for a linear approximation. Further-

more, the zero should lie within the interval, because extrapolating to outside the interval leads to uncontrolled errors. Thus, if the absolute value of the calculated zero u_0 is greater than σ we set $\sigma = 1.2|u_0|$. However, σ should not be greater than the continuation step size ds . If this condition cannot be fulfilled, we instead decrease the value of ds . If this is not possible and $|u_0|$ remains greater than σ despite all adjustments of the other parameters, we define the root to be zero. That is, we do not change the direction for the next step along the branch of the bifurcation diagram. This can also occur when close to a bifurcation point, since in this situation a reliable result cannot be achieved with an acceptable numerical effort. In this case, one blind step in the previous direction is the best option.

To adjust the parameter τ , we use the slope θ of the linear fit. This slope should be in the interval $(0.8, 1.2)\theta_d$ where θ_d is the desired slope. Based on experience, a value of $\theta_d = 1$ leads to good results. However, it may be reasonable to choose it slightly lower value (down to $\theta_d = 0.5$) to reduce the numerical cost.

For τ and θ we assume a linear relation as the first estimate, i.e. if θ does not lie within the above interval, we set

$$\tau_{\text{new}} := \tau \frac{\theta_d}{\theta}. \quad (8)$$

To adapt the parameter M , we take the error Δ into account. We determine a maximal error (denoted by Δ_d) and change M if Δ lies outside of the interval $(\Delta_d/2, \Delta_d)$. If one considers N independent and identically distributed random variables, the standard deviation of their mean is the standard deviation of one variable times $N^{-1/2}$. Thus, we suppose a linear relation between Δ and $M^{-1/2}$. Including a safety margin, we use

$$M_{\text{new}} = 1.2M \left(\frac{\Delta}{\Delta_d} \right)^2. \quad (9)$$

If the error Δ is still greater than Δ_d after the adjustment of M , we decrease the parameter σ using

$$\sigma_{\text{new}} = 0.9\sigma \frac{\Delta_d}{\Delta}. \quad (10)$$

For the parameters τ and M , it can be useful to set their minimum values to 10, for example. Otherwise, a small adjustment by a factor close to 1 has no effect, since τ and M are natural numbers.

3. Lifting is one-to-many mapping – Structure lifting

The final issue to discuss concerns the lifting procedure needed to initiate microscopic states from a known macroscopic state (see Fig. 2). Obviously, lifting is not unique but corresponds to a one-to-many mapping. This gives a large procedural freedom and requires a detailed

FIG. 7. Flow chart for the algorithm employed to adaptively adjust the parameters τ , M and σ . In each continuation step, the adjustments are applied iteratively until every condition is fulfilled.

discussion of how the lifting is done. In fact, the structure of the microscopic state can influence the resulting dynamics of the macroscopic observable, as illustrated in Fig. 8. In the best case, errors from an inconvenient lifting procedure are healed over a small initial interval within the time span τ of the evolving step. In the worst case, lifting errors crucially impact the macroscopic dynamics leading to wrong results in the root finding pro-

cess. Eventually, this also leads to uncertainties in the choice of the parameter τ . Further, for reasons of computational efficiency, this parameter has to be chosen as small as possible. Consequently, the quality and precision of the evolving step strongly depend on having an effective lifting procedure.

In fact, the ambiguity of the lifting procedure is potentially the most dangerous problem. Often, microscopic

FIG. 8. A sketch illustrating some potential problems with the lifting procedure due to it being a one-to-many mapping. The different possibilities to define a microscopic state belonging to one value of the macroscopic quantity are indicated by the axis with the label “hidden value”. The lines depict different trajectories in phase space and show that large excursions are possible before randomness is “healed out”.

(fluctuating) steady states exhibit specific spontaneously formed structures which are crucial to the time evolution of the macroscopic observable. However, their formation from unstructured microscopic states may take a time span that is too large to be able to accommodate it within the chosen τ . This problem can be tackled if one bases the stochastic continuation on something more refined than simple time simulation from unstructured initial states.

We explain our approach to solve this problem employing our first example, the Ising model. There, the microscopic state corresponds to a particular state of a lattice of spins $\{s_i\}$ while the corresponding macroscopic observable is magnetisation $m = \langle s_i \rangle$ obtained by averaging over the microscopic state. When the magnetisation is near zero and the temperature is near or below T_C , a typical spin configuration exhibits a clustered structure, which has an influence on the dynamics of the magnetisation (examples can be seen below in Sec. III).

A naive mapping in the lifting procedure is to create a spin lattice with the orientations randomly chosen to give the correct magnetisation. However, doing this leads to the spin orientations having no spatial correlations. We call this *random lifting* and illustrate it in Fig. 9. During the time evolution of the evolving step, this disordered structure starts to form clusters. However, the time span τ is not sufficiently large to complete this process. Furthermore, the correlation information is forgotten after the microscopic time evolution is terminated and in the next application of the lifting, a disordered structure is again produced.

Instead, we introduce *structure lifting*, a lifting technique that takes information from previous microscopic time evolutions into account. In particular, the final states of the evolving process are kept, which include all

FIG. 9. Scheme of the standard random lifting procedure (also called structure lifting of 0th order): In continuation step n the microscopic dynamics is initiated with a random state that corresponds to the predicted value of the macroscopic observable X^* .

FIG. 10. Scheme of the structure lifting of 1st order: In continuation step n the microscopic dynamics is initiated with a structured microscopic state obtained by adapting a final state of the microscopic dynamics from the previous continuation step $n - 1$. The adaptation is done randomly and brings the macroscopic observable X^* to the value requested by the prediction step.

the information concerning their internal structure. Then these states are employed in the lifting procedure at the next value of the control parameter. Because successive applications of the lifting step are always done at a value of the macroscopic observable that does not strongly differ from the previous value, only minor adjustments of the microscopic states have to be done. In the simplest case, these can be realized in a spatially random way, as illustrated in Fig. 10. This particular procedure we call *first-order structure lifting*.

To further improve the efficiency of the scheme, it is possible to go one step further and perform a *second-order structure lifting*. Here, information is kept concerning microscopic structures that have resulted from evolving steps in the two previous continuation steps. This allows one to extract information about the typical changes of the structure along the already calculated part of the solution branch. This information is then used to

FIG. 11. Scheme of the structure lifting of 2nd order: In continuation step n the microscopic dynamics is initiated with a structured microscopic state obtained using information from the final states of the microscopic dynamics from the two previous continuation steps $n-2$ and $n-1$. This allows one to adapt the microscopic state mainly in regions where most changes occurred between steps $n-2$ and $n-1$.

form the initial microscopic state in the next lifting procedure in a more precise way, as illustrated in Fig. 11. Namely, the necessary random changes are concentrated in regions where the most changes occurred between the two previous steps. In the example of the Ising model, this means that the minor adjustments are not uniformly distributed in space, as in the first-order structure lifting. Now, with the highest probability they are located at the interfaces between the clusters. We achieve this in practice by employing a Gaussian kernel to increase the probability density of changes in forming the state around the locations where the structure changed previously.

In principle, this concept for the lifting procedure can be extended to n -th order structure lifting. This method produces naturally evolved microscopic states without the need for long relaxation times in every evolving step. Of course, to employ structure lifting at the beginning one needs to provide the algorithm with one or more microscopic states. Depending on the specific situation, these can be generated using random lifting or a long relaxation.

Next, we illustrate the application of the proposed amendments to stochastic continuation via three examples. The first is the Ising model in Sec. III, which we have touched on already, an active Ising model in Sec. IV, and a stochastic Swift-Hohenberg equation in Sec. V.

III. ISING MODEL

The Ising model Hamiltonian is given above in Eq. (4). As mentioned, the macroscopic observable is the magnetisation $m := \langle s_i \rangle$. Depending on the dimensionless temperature $T' := k_B T / J$, different stable and unstable states exist. For the 2d Ising model, an exact solution in terms of the free energy is given by Onsager [36], while Yang first derived the formula for the spontaneous mag-

FIG. 12. The exact (Onsager) bifurcation diagram of the two-dimensional Ising model. The solid and dashed lines indicate the branches of stable and unstable states, respectively. The inset gives a typical spin configuration near the critical temperature (at $T' = 2.285$) that shows spin clusters of different sizes.

netization [47] (also announced by Onsager at a conference in 1949 [34]). As the exact result is available, this example provides a valuable test for the precision of our approach.

The exact solution is displayed in Fig. 12. It shows a pitchfork-like bifurcation at the *critical temperature* T_c , whose numerical value is given by $T'_c = k_B T_c / J \approx 2.269$. At T_c two branches of states of finite magnetisation emerge from the trivial state of zero magnetisation. For temperatures above T_c , the zero magnetisation state is the only state, i.e., it is globally stable. For $T < T_c$, this state is unstable and the states on the two branches emerging at T_c are stable. They are related by symmetry, i.e., they have the same energy. This implies that below T_c the system undergoes a spontaneous symmetry breaking. The exact solution for the magnetisation m is given by

$$m = \pm \left[1 - \left(\sinh \left(\frac{2}{T'} \right) \right)^{-4} \right]^{1/8}. \quad (11)$$

For temperatures $T' < T'_c$, and for $T' \geq T'_c$ we have $m = 0$. Close to T'_c , one writes Eq. (11) as

$$m \approx m_0 |1 - T'/T'_c|^\beta, \quad (12)$$

with the *critical exponent* β . According to the exact solution, $\beta = 1/8 = 0.125$.

Mean-field approximations, such as the Bragg-Williams approximation [48, 49], generally overestimate the critical temperature (e.g. giving $T'_c = 4$) and give the critical exponent to be $\beta = 1/2$, much larger than the exact value [34]. The mean-field exponent $\beta = 1/2$ corresponds to the value expected for a standard pitchfork bifurcation [50].

The microscopic dynamics of the spin lattice in the evolving step is given by the Metropolis algorithm [35].

That is, we randomly choose a lattice site i and then calculate the difference in the energy

$$\Delta E = 2J \sum_{\text{nn of } i} s_i s_j, \quad (13)$$

resulting from flipping the spin at the chosen site. This spin flip is then accepted with a probability of

$$p = \min(e^{-\Delta E/k_B T}, 1). \quad (14)$$

Otherwise, the flip is rejected and the spin configuration remains the same. Note that we use periodic boundary conditions. We then randomly pick another lattice site and repeat the above process. We define one time step τ as being the time to attempt N spin flips, i.e. one attempt per spin on the lattice.

We now present results obtained using the stochastic continuation algorithms described above applied to the Ising model. Those obtained using the pseudo adaptive and the adaptive parameter choice (see Sec. II C 2) are in Figs. 13 and 14, and Figs. 15 and 17, respectively. Note that we treat a change in the size of the microscopic ensemble as being equivalent to changing the lattice size. When we apply the pseudo adaptive parameter choice, the lattice size remains fixed at 1000×1000 . When employing the adaptive parameter choice, we use the smaller lattice sizes of 100×100 and 500×500 because this allows for more flexible adjustment. For all of these lattice sizes, there are negligible finite size effects (cf. Ref. [51]).

At first, we employ random lifting, i.e., to form the initial microscopic state with magnetisation m , we choose the orientation of each spin according to

$$s_i = \begin{cases} 1, & \text{with probability } (1+m)/2 \\ -1, & \text{with probability } (1-m)/2 \end{cases} \quad (15)$$

This procedure produces a spatially disordered spin configuration whose magnetisation is – due to the law of large numbers – very close to the value m . Stochastic continuation results in the bifurcation diagram shown in Fig. 13. Overall, there is a good agreement, except for the stable $m \neq 0$ branches near the critical temperature. Extrapolating these to $m = 0$, we find that this continuation method gives the critical temperature to be $T_c^{\text{cont}} = 2.447$. This differs from the exact value of 2.269 by about 8%, but it is much better than the mean field approximation result.

Employing first-order structure lifting, one can further improve this result (see Fig. 14), getting an extrapolated value of the critical temperature $T_c^{\text{cont}} = 2.296$, which differs only by 1.2% from the exact value. This highlights the importance of the specific way the lifting procedure is done and how important it is to use information regarding the spatial structure of the microscopic state. It also shows the influence this has on the dynamics of the microscopic state. In the Ising example, the clustering of the spin orientations is particularly important for the dynamics near the bifurcation point. The inset of Fig. 12

FIG. 13. Bifurcation diagram of the steady states of the 2d Ising model obtained using the presented stochastic continuation algorithm in the case with random lifting. The blue circles with red error bars are the continuation results while the green dashed lines represent the exact solution. The C³R method is used and a pseudo adaptive adjustment of τ . Namely, τ is linearly increased from 1 to τ_{max} while approaching the bifurcation point (here, first a rough idea of the location of the bifurcation point can be obtained by a continuation run with a fixed value of τ). For the nontrivial branches with $m \neq 0$ we chose $\tau_{\text{max}} = 70$ while for the trivial $m = 0$ branch $\tau_{\text{max}} = 50$ [$\tau_{\text{max}} = 100$] when approaching the critical point from the left [right]. The remaining parameters are $ds = 0.1$, $N_{\text{fit}} = 100$, $M = 1$ and $\sigma = 0.05$. The size of the spin lattice is 1000×1000 .

FIG. 14. Bifurcation diagram showing the $m \neq 0$ steady states of the 2d Ising model obtained using random lifting (blue circles) and first-order structure lifting (green diamonds). The remaining line styles, settings of the continuation algorithm and parameters are as in Fig. 13.

gives a typical spin configuration with cluster formation at a temperature near T_c' .

However, continuation with this first-order structure lifting encounters problems along the trivial branch of zero magnetisation. Namely, the algorithm is not able to follow the unstable solution branch (see the blue cir-

FIG. 15. Bifurcation diagram of the trivial steady states of the 2d Ising model as obtained using first-order structure lifting (blue circles) and second-order structure lifting (green diamonds). The dashed lines belong to the exact solution. The fully adaptive choice of τ , M and σ is used. The remaining parameters are $ds = 0.4$, $N_{\text{fit}} = 10$, $\Delta_d = 0.04$ and $\theta_d = 1.0$. The size of the spin lattice is 500×500 .

cles for $T' < T'_c$ in Fig. 15). Alternatively, employing the second-order structure lifting, which has a higher probability that adjustments are located near the borders of the clusters, result in microscopic states which are more natural and exhibits a reliable microscopic dynamics. The results obtained are shown in Fig. 15. There we see that the second-order structure lifting gives a better result for the unstable part of the trivial $m = 0$ branch.

This example shows that the first-order structure lifting exhibits a specific problem when applied to microscopic states that are strongly clustered. As explained above, the adjustments made to the spin configuration during the lifting process are spatially random. Thus, it is very likely that the adjustments occur at lattice sites which are in the midst of the spatially ordered regions. Then, with a high probability [see eq. (14)] the microscopic dynamics reverses these adjustments resulting in a state with the wrong average magnetization. Hence, the continuation can be led astray. With second-order structure lifting, this issue is improved.

We return now to the $m \neq 0$ nontrivial branch. Continuation of this using first-order structure lifting and adaptive parameter choice while using a small value of Δ_d lead to a result with small statistical errors. This result is shown in Fig. 16 using orange triangles. The curve exhibits very few fluctuations and has very small error bars. Yet, the agreement with the exact solution is not as good as the result from using the pseudo-adaptive first-order structure lifting discussed before. Though it is still much better than the results of random lifting. Note that the error bars do not reach the exact solution. This highlights two important points. First, not surprisingly, our automatic adaptive parameter choice is not as good as the parameter choice chosen by an operator with some physical insight. Second, our error measure Δ

FIG. 16. The critical region from Fig. 14, showing results of the pseudo-adaptive continuation of nontrivial steady states of the 2d Ising model using random lifting (blue circles) and first-order structure lifting (green diamonds) compared to the exact solution (dashed line). The additional line (orange triangles) represents the result of a continuation using first-order structure lifting and adaptive choice of τ , M and σ with parameters $ds = 0.1$, $N_{\text{fit}} = 20$, $\Delta_d = 0.001$, $\theta_d = 1.0$ and spin lattice size 100×100 .

FIG. 17. The data for the nontrivial $m \neq 0$ branch in Fig. 16 obtained with first-order structure lifting and adaptive parameter choice, presented in a log-log plot (blue circles) with a corresponding linear fit of slope 0.216 given as blue dashed line. Note that the leftmost point is not considered for the fit. The green dotted line represents Onsager's exact solution with slope $1/8$. For comparison, the red dotted line shows a slope of 0.5 corresponding to the mean field solution. The required value of T_c is obtained by an extrapolation of the continuation result.

[cf. Eq. (6)] only represents the statistical errors of the stochastic continuation algorithm. It does not include other errors, e.g., those occurring in the lifting procedure. This implies that there are still lifting errors. Of course, our structure lifting can only be an approximation of a perfect lifting operator.

In Fig. 17 we present the results from this last contin-

uation that has small statistical errors in a log-log plot. A linear fit then gives the critical exponent $\beta = 0.216$. We see the significant deviations from the exact results near to the critical temperature and one does not really see a clear power law. Instead, there are a range of values. However, almost all of these are much smaller than the mean field critical exponent of $1/2$. Hence, the continuation gives a much better result than the mean field approximation, both for T_c and also for β . To obtain a clear result for the critical exponent, further improvements of the continuation technique are needed.

IV. AN ACTIVE ISING MODEL

The second example we consider is the active Ising model introduced by Solon and Tailleur in Ref. [52] to capture essential features of the flocking transition occurring in many discrete and continuous models for the collective behaviour of active particles [53–57]. The active Ising model is a nonequilibrium version of a ferromagnetic model with the spin variable having the meaning of a (preferred) state of motion. Here we consider the 2d version from Ref. [37] (whose notation we follow) where the particles' motion in the positive and negative x -direction, is influenced by the spin value. Our lattice has size $L_x \times L_y = 100 \times 100$ and contains N moving particles that carry a spin of ± 1 . Every site of the lattice i can be occupied by an arbitrary number of particles, denoted by $n_i^\pm$ for particles with spin ± 1 . Consequently, for every lattice site i , we have a local particle density $\rho_i := n_i^+ + n_i^-$ and a local magnetisation $m_i := n_i^+ - n_i^-$. Depending on the spin orientations, every particle can move to a neighbouring lattice site or flip its spin at specific rates. The rules controlling the dynamics are:

- A particle located at the site i and having a spin s flips its spin at a rate

$$W(s \rightarrow -s) = \exp\left(-s\beta \frac{m_i}{\rho_i}\right), \quad (16)$$

where $\beta = 1/T$ is an inverse nondimensional temperature.

- A particle with spin s moves with rate D to the upper or lower lattice site (unbiased motion in y -direction), with rate $D(1 + s\varepsilon)$ to its right and with rate $D(1 - s\varepsilon)$ to its left (spin-biased motion in x -direction, where ε measures the strength of the bias).

With these rates, a continuous-time Markov process is defined. For the realization of the dynamics, we use a random-sequential-update algorithm. For this, we discretize the time in steps of Δt . In every step of the algorithm, a particle is chosen randomly and one of the following actions is done:

- With probability $W(s \rightarrow -s)\Delta t$, its spin is flipped.

- With probability $D\Delta t$, it is moved upwards.
- With probability $D\Delta t$, it is moved downwards.
- With probability $D(1 + s\varepsilon)\Delta t$, it is moved to the right.
- With probability $D(1 - s\varepsilon)\Delta t$, it is moved to the left.
- With probability $1 - [4D + W(s \rightarrow -s)]\Delta t$, nothing is done.

After this, the time is incremented by $\Delta t/N$. In order to keep all probabilities smaller than 1 and to minimize the probability that nothing happens, following [52] Δt is chosen to be $(4D + \exp(\beta))^{-1}$.

Here, we choose $\beta = 2$, $\varepsilon = 0.9$ and $D = 1$ and employ the developed stochastic continuation algorithm to analyse the dynamics of the system and how it depends on the average density $\rho_0 = N/(L_x L_y)$, which is determined by N . At low densities, a trivial homogeneous unpolarised state emerges that has zero mean magnetisation (sometimes called “polarisation”), i.e., $\langle m_i \rangle \approx 0$. That is, the particles show no collective behaviour and the state may be seen as a homogeneous gas state. At high densities, we have $\langle m_i \rangle \neq 0$, left-right parity is broken and the particles show a collective motion either to the right or to the left. We call this the “active liquid state”.

At intermediate densities, a phase separation between gas and liquid state is observed and regions of the two states coexist. In particular, one finds dense bands of moving particles as illustrated in the inset of Fig. 18, which shows a snapshot of the local magnetisations. The main panel presents the spatial profile of the local magnetisation averaged along the band, i.e., averaged in the y -direction. Employing a trapezoidal piecewise linear fit, we determine the mean m_l of the local magnetisation in the band of active liquid phase, i.e., the “height” of the plateau. This is then used to define the macroscopic observable

$$\Phi = \frac{\langle m \rangle}{m_l} = \frac{1}{m_l L_x L_y} \sum_i m_i \quad (17)$$

which quantifies the liquid-gas ratio. Varying the average density ρ_0 , the width of the band of the liquid phase changes while the value m_l remains constant in the entire coexistence region.

To analyse the system behaviour employing the developed stochastic continuation algorithm we define a random lifting procedure in a straightforward manner: Given the total particle number N , we choose random lattice sites to place the particles and assign a random spin orientation to all of them using a probability that guarantees a specific value of $\langle m \rangle$.

For densities belonging to coexistence states, we use a hybrid structure lifting: We modulate a coexistence state which is given by the previous continuation step by adding or removing particles at random lattice sites as

FIG. 18. Example of a microscopic magnetisation state of the 2d active Ising model with mean density $\rho_0 = 3.98$, which is in the parameter region where the gas and the liquid state coexist. The inset shows a typical snapshot of the local magnetisation, clearly showing that the collective behaviour results in a band structure. The central band represents the liquid phase which moves to the left through a gaseous background. The main plot gives the spatial profile of the local magnetisation averaged along the band (black solid line) together with a trapezoidal piecewise linear fit (blue solid line). The thin dashed horizontal lines give the magnetisation values of the background gas phase ($m = 0$), of the band of liquid phase ($m = m_l$), and the mean $\langle m \rangle$. The fit is used to extract the macroscopic observable defined as the liquid-gas fraction $\Phi = \langle m \rangle / m_l$. The remaining parameters are $L_x = L_y = 100$, $\beta = 2$, $\varepsilon = 0.9$ and $D = 1$.

well as by changing spin orientations of random particles. However, to maintain the stable coexistence structure, we only do these random adjustments in the region of the boundary between the liquid and the gas phase. The phase boundary is well defined by the sloped parts of the trapezoidal fits (see Fig. 18). This procedure is a targeted first-order structure lifting which acts like a higher order structure lifting. In Fig. 19 we present continuation results for just the branches of stable states, in particular, for the transition between the banded state and the liquid state.

We initiate the continuation with a banded state obtained from a time simulation at $\rho_0 = 3.0$ and then employ continuation to follow the stable branch toward higher mean densities. The result in Fig. 19 is in very good agreement with the results obtained by direct time simulation in Ref. [37]. Namely, the branches corresponding to the banded state and the liquid state are well approximated by straight lines and meet at a bifurcation point at $\rho_0 \approx 5.4$. In particular, it seems that at comparable system sizes, the continuation approach can capture the bifurcation point much better than the direct simulations of Ref. [37]. Their Fig. 10 (left) shows that for a system size of 200×100 there is still a jump in Φ of nearly 20% in the region where the bifurcation point is expected. They have to go to rather large systems to

FIG. 19. Bifurcation diagram of stable steady states of the 2d active Ising model obtained using stochastic continuation. The continuation is initialized on the left, uses second-order structure lifting and adaptive choices for τ , M and σ . The other parameters in the continuation algorithm are: $ds = 0.1$, $N_{\text{fit}} = 15$, $\Delta_d = 0.01$ and $\theta_d = 0.7$. The (rather small) errors are indicated in red.

capture well the bifurcation point.

Note, however, that our approach works less well when initiating the continuation at a high particle density in the liquid phase and descending in ρ_0 . Then, when passing the bifurcation point into the coexistence region, the stable band structure needs a very long time to form. Thus, even with structure lifting, the banded stable states can not be found with an acceptable numerical effort when passing the bifurcation point. This agrees with the general observation that the continuation method generally works better when it is initialized sufficiently away from potential bifurcation points and then subsequently approaching them. In other words, problems should be expected when crossing a bifurcation point onto a branch having a different symmetry.

V. SWIFT-HOHENBERG EQUATION

The third and final example to illustrate the application of the developed stochastic continuation algorithm concerns a stochastic PDE, namely, the stochastic version of the Swift-Hohenberg equation in Eq. (5) [38]. The corresponding deterministic Swift-Hohenberg equation which is obtained when $D = 0$, represent a generic model capturing the dynamics of pattern formation in the vicinity of a monotonic short-wave instability, i.e., a Turing instability [58]. Depending on the control parameter ε , the field Ψ has a variety of different stable states. For $\varepsilon < 0$, the homogeneous state is stable. For a 1d system, a pitchfork bifurcation occurs at $\varepsilon = 0$ and the homogeneous state is unstable for $\varepsilon > 0$. With the present purely cubic nonlinearity, the bifurcation is supercritical, i.e., a branch of stable periodic states emerges towards

FIG. 20. The bifurcation diagram for the deterministic Swift-Hohenberg equation (Eq. (5) with $D = 0$) obtained by standard continuation. The solid and dashed lines are branches of stable and unstable states, respectively. The branch with $\|\delta\Psi\| = 0$ (blue line) is the trivial, homogeneous state which changes its stability at $\varepsilon = 0$, where two branches emerge supercritically. These are a branch of stable stripe patterns (green line) and a branch of unstable square patterns (orange line). The insets show a stripe and a square state, both at $\varepsilon = 0.5$.

$\varepsilon > 0$. Other nonlinearities may result in a subcritical bifurcation and, in general, more intricate behaviour [59].

Here, we investigate the 2d case for a square domain of size $2\pi \times 2\pi$. Then, at $\varepsilon > 0$ branches of stripe patterns (equivalent to the 1d periodic states) and of square patterns simultaneously emerge supercritically. In the deterministic case one can follow them by standard continuation tools, such as PDE2PATH [14]. Figure 20 gives the corresponding bifurcation diagram employing as solution measure the L^2 -norm of the deviations from the mean, i.e.,

$$\|\delta\Psi\| := \sqrt{\langle (\Psi - \langle \Psi \rangle)^2 \rangle}. \quad (18)$$

Note that for the present square domain, the square patterns are always unstable while the stripes are stable. In a rectangular domain suitable for hexagonal patterns (ratio of side lengths of $\sqrt{3}$) one finds stripes and two types of hexagonal patterns and more intricate stability behaviour [58].

Although, there exist many efficient numerical continuation tools to determine bifurcation diagrams for deterministic equations such as Eq. (5) with $D = 0$, here we show that stochastic continuation can also be used. This allows us to check the quality and precision of our method. In Fig. 20 we compare our results with the results obtained with PDE2PATH. To bring the deterministic system into a form suitable for stochastic continuation we need to define macroscopic observables and microscopic states. We take the norm $\|\delta\Psi\|$ as the macroscopic observable and the spatially discretized version of the field Ψ as the microscopic state. We compute the micro-

FIG. 21. The bifurcation diagram obtained using stochastic continuation applied to the deterministic Swift-Hohenberg equation (Eq. (5) with $D = 0$) in 2d. The circles and diamonds represent stable and unstable branches, respectively. For comparison, the dashed line gives the standard continuation results obtained with PDE2PATH. For the branch of homogeneous [stripe] states, the continuation is started on the left at $\varepsilon = -0.1$ [right at $\varepsilon = 0.45$]. First-order structure lifting and adaptive parameter choice for τ and σ are used. However, $M_{\max}$ is fixed at 1 because no ensemble of microscopic states is needed for the deterministic system. The other parameters of the continuation algorithm are $ds = 0.02$ or $ds = 0.05$, $N_{\text{fit}} = 6$, $\Delta_d = 0.01ds$ and $\theta_d = 1.0$. The size of the spatially discrete representation of the field Ψ is 128×128 .

scopic dynamics via a standard pseudo-spectral method on a grid of 128×128 points with periodic boundaries.

Here, our random lifting consists in assigning a normally distributed random variable to every lattice site. By choosing a suitable variance, we obtain a random field for a specific given value of $\|\delta\Psi\|$. When using first-order structure lifting, we adjust the final state of the evolving procedure from the previous continuation step by multiplying it by an appropriate factor which lies near one.

The continuation with first-order structure lifting of the branch belonging to the homogeneous state works without any problems. In this case we have a horizontal branch and the stability is directly given by the slope of the function $\tilde{G}$. Also the branch of stripe states is obtained in good agreement with the result from standard continuation in Fig. 21. However, this is only possible with the structure lifting. With random lifting, during the short time evolution of the evolving step the Laplace operator results in diffusion on short scales and smoothes out noisy structures. Therefore, at the very beginning of the time evolution, a state generated with random lifting always moves towards a homogeneous state, allowing to follow the unstable trivial state. The slope of $\tilde{G}$ then indicates that on short time scales the state appears to be stable.

The unstable branch belonging to the square pattern cannot automatically be followed using our structured lifting algorithm because for positive values of ε it only

FIG. 22. Results from stochastic continuation applied to the stochastic Swift-Hohenberg equation in 1d, Eq. (19). The continuation is initiated at $\varepsilon = 0.45$, uses first-order structure lifting and fully adaptive choices for τ , M and σ . The other parameters of the continuation algorithm are $ds = 0.02$, $N_{\text{fit}} = 20$, $\Delta_d = 0.004$ and $\theta_d = 1.0$. The field $\Psi(x, t)$ is discretized on a spatial grid of 128 cells. The dotted black line represents the continuation result for the corresponding deterministic system (Eq. (19) with $D = 0$).

“sees” the dynamics going from the unstable homogeneous solution to the stable stripe pattern. As the square pattern does not lie “in between” the homogeneous and stripe state, the algorithm does not detect it. However, one can devise other types of structure lifting that make it easier to find unstable solutions like the square pattern; see Sec. VI for a detailed discussion.

Next we consider the $D > 0$ stochastic version of the Swift-Hohenberg equation (5). For the sake of a lower computational effort, here we restrict our attention to the 1d version of the equation with additive noise

$$\partial_t \Psi = \varepsilon \Psi - \Psi^3 - (1 + \partial_x^2) \Psi + \sqrt{D} \eta(x, t) \quad (19)$$

where the noise of strength D is normally distributed and uncorrelated in space and time, i.e., $\langle \eta \rangle = 0$ and $\langle \eta(x', t') \eta(x, t) \rangle = \delta(x' - x) \delta(t' - t)$. We use the value $D = 1$ and a domain size of 2π . Again, we use a pseudo-spectral method with periodic boundaries for the computation of the microscopic dynamic. Hereby, the field $\Psi(x, t)$ is discretized on a spatial grid of 128 cells. The stochastic continuation is initiated on the stable branch of periodic states at $\varepsilon = 0.45$ and proceeds toward smaller ε . The results are displayed in Fig. 22. We see that far from the pitchfork bifurcation of the deterministic system, the branch of periodic states is nearly not influenced by the already rather strong noise. On the other hand, the branch of stable trivial states slightly deviates from $\|\delta\Psi\| = 0$ near the pitchfork bifurcation and has become an imperfect pitchfork bifurcation, as expected when noise is present [41].

VI. CONCLUSION AND OUTLOOK

In the present work we have developed ways to improve the stochastic continuation of stable and unstable steady states in several aspects, mainly concerning the corrector step of the continuation algorithms. In all our work we have employed a secant predictor, however, this can be amended to any commonly used predictor. Our proposals for the corrector step concern all its lifting and evolving steps. The restricting step always consists of a straight forward averaging procedure.

First, we have amended the root finding procedure for the fluctuating function G_λ that results in the evolving step from the microscopic time stepper. Instead of directly applying a root finding method to the fluctuating function G_λ , we evaluate a set of function values in a suitable neighbourhood of the initial guess provided by the predictor and perform a linear fit that is then used to determine the root. We have found that this procedure moderates the fluctuations and works in a stable way.

Our second improvement concerns the choice of numerical parameters. As the stochastic continuation algorithm involves quite a large number of numerical parameters, e.g., the length of the microscopic time evolution in the evolving step, the number of employed microscopic realisations and the step size of the path continuation step. As their choice has a crucial influence on the quality and precision of the continuation results, we suggest to adaptively adjust specific parameters during the continuation run. In particular, in the vicinity of bifurcation points, this has led to more accurate results.

In the third amendment we have considered the lifting procedure that forms a very critical part of stochastic continuation. In most of the cases in the literature, a simple constrained random lifting is used, i.e., a microscopic state with the correct macroscopic observable imposed is chosen randomly. In contrast, we argue that one should exercise some control over its internal spatial structure that is important for the macroscopic dynamics. In particular, we have developed first and second-order structure lifting procedures. In the case of the first-order procedure, final states from the microscopic evolution from the previous continuation step are randomly adapted to the new value of the macroscopic observable. In the case of the second-order procedure, final states from the microscopic evolution from the two previous continuation steps are used to obtain a state at the value of the macroscopic observable through random adjustments only in the region most changed over the previous continuation steps. This procedure can be further refined into an n -th order procedure.

Furthermore, in specific examples, it can be useful to incorporate information about particular solution behaviour into the lifting and/or restricting procedure as it is done by our hybrid lifting for the active Ising model (Sec. IV) and by the weighted lifting and restriction for the agent-based model discussed in Ref. [17].

The abilities and shortcomings of the original method

and of our various proposed amendments have then be discussed using three examples: (i) the classical Ising model in 2d, (ii) an active Ising model, i.e., a kinetic Monte-Carlo model for the collective behaviour of active self-propelled particles, and (iii) the Swift-Hohenberg equation, a generic PDE model for short-scale pattern formation close to a Turing bifurcation. The final example has on the one hand been used to show that our stochastic continuation algorithm can be employed to continue solutions of a deterministic equation (the deterministic Swift-Hohenberg equation in 2d) and, on the other hand, to continue nontrivial steady states of a stochastic PDE (the Swift-Hohenberg equation with additive noise in 1d). These examples have shown that our improvements to the stochastic continuation method are rather important for the quality and precision of the resulting bifurcation diagrams.

We highlight that in example (i) stochastic continuation now provides a bifurcation diagram quite closely resembling Onsager’s analytical solution. Our root finding procedure has allowed for a much closer approach to the bifurcation point, compare Fig. 15 of [19] with the present Fig. 14. Also the fluctuations in the branch of trivial states are much reduced (see Fig. 15). In particular, we have obtained an estimate for the critical exponent that is not very far away from the true value. However, as no clear unique power law emerges there is still space for improvements to the method, in particular, very close to bifurcation points.

This includes in the lifting procedure, as the structure lifting still fails if the microscopic dynamics of the system is too slow. This we have seen in the example (ii), where it was not possible to follow the stable branch through the bifurcation point when initiating the continuation at large values of the control parameter.

Secondly, there can be problems when using the structure lifting to follow unstable steady states. In principle, it should be possible to follow such states, since unstable steady states of spatially extended systems normally correspond to saddle points in a very high-dimensional space of solutions. Within this space there usually exist only one or a few slow unstable modes among many fast stable modes. Therefore, on short time scales one can expect to have a healing process in the evolving step where unstable solutions are approached. This forms the basis of stochastic continuation of unstable states.

We have seen related issues in the deterministic Swift-Hohenberg system [example (iii)], where we have not been able to find the branch of the unstable square pattern, which lies between the unstable branch belonging to the homogeneous solution and the stable branch of the square pattern. The reason for this is that, during continuation with structure lifting, we let the microscopic pattern evolve on its own. As a result, the stable stripe pattern is favoured over the unstable square pattern. To remedy this, one could feed in the missing square pat-

tern once as a basis for the continuation of the missing branch. The idea is to use it only for the first lifting process which is followed by structure liftings as they are done before.

In the example of the deterministic Swift-Hohenberg equation, it is easy to obtain the missing unstable pattern from a PDE2PATH solution, for example. In other examples where only stochastic continuation can be applied, the issue of how one obtains information about potentially existing unstable steady states, such as the square patterns in the Swift-Hohenberg example, is very important. An approach to this problem could be the following. As a system evolves on the way to a stable solution, the state of the system often gets close to unstable solutions, because they behave like saddle points. So, observing the time evolution, starting from an arbitrary state can give clues about possible unstable states. For example, the stripe pattern of the Swift-Hohenberg system evolves out of a homogeneous state for positive values of the parameter ε . But, before the stripes emerge, typically, one sees many localised peaks. These peaks are connected into stripes later on, but they are very similar to the peaks which form the square pattern.

Furthermore, our adaptive adjustments of the numerical parameters makes it possible to obtain more precise results for the zeros of the function $\tilde{G}$. Yet, this can be improved further. Choosing the parameter values by hand in a pseudo adaptive way (as we have done for the results in Fig. 14, for example) is a laborious process. However, human experience (physical insight) accumulated over time still allows one to effectively gauge the obtained plots and linear fits and to adjust parameters by hand and to obtain better results than all of the automated algorithms presented. In the future one can envisage employing machine learning, e.g., using neural networks, for this task. Currently, such methods are quite successful in the solution of visual problems.

Finally, emphasize again the connection between the slope of the function $\tilde{G}$ and the stability of the corresponding solution branch. Moreover, we have seen that the occurrence of a bifurcation should coincide with a slope of zero. However, these connections depend on the direction of the orthogonal corrector which varies while following a curved solution branch. In the future this issue should be elaborated in detail such that one is able to perform a systematic stability analysis of all states on the branches being followed.

ACKNOWLEDGMENTS

We thank Tobias Frohoff-Hülsmann for helpful discussions and for providing the PDE2PATH continuation runs for the deterministic Swift-Hohenberg equation.

- [1] A. Deutsch, P. Maini, and S. Dormann, *Cellular Automaton Modeling of Biological Pattern Formation: Characterization, Applications, and Analysis*, Modeling and Simulation in Science, Engineering and Technology (Birkhäuser, Boston, 2007).
- [2] B. Chopard and M. Droz, *Cellular automata modeling of physical systems*, Aléa-Saclay (Cambridge University Press, 2005).
- [3] P. Tu, *Dynamical systems: An introduction with applications in economics and biology* (Springer-Verlag, Berlin New York, 1994).
- [4] H. Haken, *Information and self-organization: a macroscopic approach to complex systems*, Springer Series in Synergetics (Springer, Berlin Heidelberg, 2006).
- [5] L. M. Pismen, *Patterns and interfaces in dissipative dynamics*, Springer Series in Synergetics (Springer-Verlag, Berlin Heidelberg, 2006).
- [6] Y. A. Kuznetsov, *Elements of applied bifurcation theory*, 3rd ed. (Springer, New York, 2010).
- [7] J. Argyris, G. Faust, M. Haase, and R. Friedrich, *An exploration of dynamical systems and chaos: completely revised and enlarged second edition* (Springer, Berlin Heidelberg, 2015).
- [8] B. Krauskopf, H. M. Osinga, and J. Galan-Vioque, eds., *Numerical continuation methods for dynamical systems* (Springer, Dordrecht, 2007).
- [9] H. A. Dijkstra, F. W. Wubs, A. K. Cliffe, E. Doedel, I. F. Dragomirescu, B. Eckhardt, A. Y. Gelfgat, A. Hazel, V. Lucarini, A. G. Salinger, E. T. Phipps, J. Sanchez-Umbria, H. Schuttelaars, L. S. Tuckerman, and U. Thiele, *Commun. Comput. Phys.* **15**, 1 (2014).
- [10] E. L. Allgower and K. Georg, *Introduction to numerical continuation methods (classics in applied mathematics)* (Society for Industrial Mathematics, 1987).
- [11] Z. Mei, *Numerical bifurcation analysis for reaction-diffusion equations* (Springer, Berlin Heidelberg, 2000).
- [12] J. Sanchez, F. Marques, and J. M. Lopez, *J. Comput. Phys.* **180**, 78 (2002).
- [13] J. Sanchez Umbria and M. Net, *Eur. Phys. J.-Spec. Top.* **225**, 2465 (2016).
- [14] S. Engelnkemper, S. V. Gurevich, H. Uecker, D. Wetzel, and U. Thiele, in *Computational Modeling of Bifurcations and Instabilities in Fluid Mechanics*, Computational Methods in Applied Sciences, vol 50, edited by A. Gelfgat (Springer, 2019) pp. 459–501.
- [15] C. Theodoropoulos, Y. H. Qian, and I. G. Kevrekidis, *Proc. Natl. Acad. Sci. U. S. A.* **97**, 9840 (2000).
- [16] I. G. Kevrekidis and G. Samaey, *Ann. Rev. Phys. Chem.* **60**, 321 (2009).
- [17] D. Avitabile, R. B. Hoyle, and G. Samaey, *SIAM J. Appl. Dyn. Syst.* **13**, 1583 (2014).
- [18] C. Kuehn, *SIAM J. Sci. Comput.* **34**, A1635 (2012).
- [19] S. A. Thomas, D. J. B. Lloyd, and A. C. Skeldon, *Physica A* **464**, 27 (2016).
- [20] D. Barkley, I. G. Kevrekidis, and A. M. Stuart, *SIAM J. Appl. Dyn. Syst.* **5**, 403 (2006).
- [21] J. Sieber and B. Krauskopf, *Int. J. Bifurcation Chaos* **17**, 2579 (2007).
- [22] J. Sieber and B. Krauskopf, *Nonlinear Dyn.* **51**, 365 (2008).
- [23] D. A. Barton, B. P. Mann, and S. G. Burrow, *J. Vib. Control* **18**, 509 (2012).
- [24] Y. Frederix, G. Samaey, C. Vandekerckhove, T. Li, E. Nies, and D. Roose, *Discret. Contin. Dyn. Syst.-Ser. B* **11**, 855 (2009).
- [25] A. G. Makeev, D. Maroudas, and I. G. Kevrekidis, *J. Chem. Phys.* **116**, 10083 (2002).
- [26] C. Kuehn, *SIAM-ASA J. Uncertain. Quantif.* **3**, 762 (2015).
- [27] R. Pasupathy and S. Kim, *ACM Trans. Model. Comput. Simul.* **21** (2011).
- [28] Y. Frederix, G. Samaey, C. Vandekerckhove, T. Li, E. Nies, and D. Roose, *Discret. Contin. Dyn. Syst.-Ser. B* **11**, 855 (2009).
- [29] S. Sriraman, I. G. Kevrekidis, and G. Hummer, *Phys. Rev. Lett.* **95**, 130603 (2005).
- [30] D. I. Kopelevich, A. Z. Panagiotopoulos, and I. G. Kevrekidis, *J. Chem. Phys.* **122**, 044907 (2005).
- [31] J. Sieber, C. Marschler, and J. Starke, *SIAM J. Appl. Dyn. Syst.* **17**, 2574 (2018).
- [32] C. Gear, T. Kaper, I. G. Kevrekidis, and A. Zagaris, *SIAM J. Appl. Dyn. Syst.* **4**, 711 (2004).
- [33] C. Siettos, M. D. Graham, and I. G. Kevrekidis, *J. Chem. Phys.* **118**, 10149 (2003).
- [34] R. J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Elsevier Science, 2016).
- [35] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
- [36] L. Onsager, *Phys. Rev.* **65**, 117 (1944).
- [37] A. P. Solon and J. Tailleur, *Phys. Rev. E* **92**, 042119 (2015).
- [38] A. Hutt, *Europhys. Lett.* **84**, 34003 (2008).
- [39] J. Viñals, E. Hernández-García, M. San Miguel, and R. Toral, *Phys. Rev. A* **44**, 1123 (1991).
- [40] K. R. Elder, J. Viñals, and M. Grant, *Phys. Rev. Lett.* **68**, 3024 (1992).
- [41] G. Agez, M. G. Clerc, E. Louvergneaux, and R. G. Rojas, *Phys. Rev. E* **87**, 042919 (2013).
- [42] T. Taneike, T. Nihei, and Y. Shiwa, *Phys. Lett. A* **303**, 212 (2002).
- [43] E. Doedel, H. B. Keller, and J. P. Kernevez, *Int. J. Bifurcation Chaos* **1**, 493 (1991).
- [44] E. J. Doedel, B. E. Oldeman, A. R. Champneys, F. Dercole, T. F. Fairgrieve, Y. Kuznetsov, R. C. Paffenroth, B. Sandstede, X. J. Wang, and C. H. Zhang, *Auto-07p: continuation and bifurcation software for ordinary differential equations* (Concordia University, 2012) <http://indy.cs.concordia.ca/auto/>.
- [45] E. Doedel, H. B. Keller, and J. P. Kernevez, *Int. J. Bifurcation Chaos* **1**, 745 (1991).
- [46] H. Uecker, D. Wetzel, and J. D. M. Rademacher, *Numer. Math. Theor. Meth. Appl.* **7**, 58 (2014).
- [47] C. N. Yang, *Phys. Rev.* **85**, 808 (1952).
- [48] L. E. Reichl, *A modern course in statistical physics* (Wiley-VCH, Weinheim, 2016).
- [49] A. P. Hughes, U. Thiele, and A. J. Archer, *Am. J. Phys.* **82**, 1119 (2014).
- [50] S. H. Strogatz, *Nonlinear Dynamics and Chaos* (Westview Press, New York, 2014).
- [51] A. E. Ferdinand and M. E. Fisher, *Phys. Rev.* **185**, 832 (1969).

- [52] A. P. Solon and J. Tailleur, Phys. Rev. Lett. **111**, 078101 (2013).
- [53] J. Toner, Y. H. Tu, and S. Ramaswamy, Ann. Phys. **318**, 170 (2005).
- [54] F. Peruani, A. Deutsch, and M. Bär, Phys. Rev. E **74**, 030904 (2006).
- [55] H. Chaté, F. Ginelli, G. Grégoire, F. Peruani, and F. Raynaud, Eur. Phys. J. B **64**, 451 (2008).
- [56] T. Ihle, Phys. Rev. E **83**, 030901 (2011).
- [57] A. P. Solon, H. Chaté, and J. Tailleur, Phys. Rev. Lett. **114**, 068101 (2015).
- [58] M. C. Cross and P. C. Hohenberg, Rev. Mod. Phys. **65**, 851 (1993).
- [59] J. Burke and E. Knobloch, Phys. Rev. E **73**, 056211 (2006).