

Chaotic macroscopic phases in one-dimensional oscillators

Antonio Politi^{1,a}, Arkady Pikovsky^{2,b}, and Ekkehard Ullner^{1,c}

¹ Institute for Complex Systems and Mathematical Biology and SUPA, University of Aberdeen, Aberdeen AB24 3UE, UK

² Department of Physics and Astronomy, Potsdam University, Karl-Liebknecht-Str 24/25, 14476 Potsdam, Germany

Received 14 February 2017 / Received in final form 26 March 2017
Published online 21 June 2017

Abstract. The connection between the macroscopic description of collective chaos and the underlying microscopic dynamics is thoroughly analysed in mean-field models of one-dimensional oscillators. We investigate to what extent infinitesimal perturbations of the microscopic configurations can provide information also on the stability of the corresponding macroscopic phase. In ensembles of identical one-dimensional dynamical units, it is possible to represent the microscopic configurations so as to make transparent their connection with the macroscopic world. As a result, we find evidence of an intermediate, mesoscopic, range of distances, over which the instability is neither controlled by the microscopic equations nor by the macroscopic ones. We examine a whole series of indicators, ranging from the usual microscopic Lyapunov exponents, to the collective ones, including finite-amplitude exponents. A system of pulse-coupled oscillators is also briefly reviewed as an example of non-identical phase oscillators where collective chaos spontaneously emerges.

1 Introduction

Collective chaos (CC), i.e., irregular behavior on macroscopic scales, is known to emerge in a large class of dynamical systems. It was first observed in an ensemble of mean-field coupled logistic maps [1] and then confirmed in Stuart-Landau oscillators [2, 3]. This latter setup is particularly relevant, as the dynamics of the single units is described by the normal form of the Hopf bifurcation and thus naturally applies to a wide variety of physical contexts. This regime is, however, not limited to mean-field models: it has been observed in sparsely connected networks, where as few as 60 connections per oscillator suffice to sustain global fluctuations [4]. CC has been observed also in spatially extended systems with long range coupling whenever the decay rate of the coupling strength is sufficiently slow [5]. In this context the large-scale dynamics

^a e-mail: a.politi@abdn.ac.uk

^b e-mail: pikovsky@uni-potsdam.de

^c e-mail: e.ullner@abdn.ac.uk

manifests itself as hydrodynamic (long-wavelength) modes whose amplitudes do not scale to zero upon increasing the system size.

Altogether, CC has raised the attention of many researchers because of a puzzling feature: the emergence of large-scale correlations in an environment of chaotic (and thus supposedly uncorrelated) units. Formally speaking, CC can be understood as the outcome of the application of a nonlinear operator (Perron-Frobenius or Liouville, depending whether the time variable is either discrete or continuous) to a probability density in phase space. Normally, such operators are linear, but the mean-field feedback makes them nonlinear and this paves the way to a possibly chaotic dynamics over macroscopic scales. Practically, it is however still unclear under which conditions CC arises and, more importantly, how microscopic and macroscopic dynamics are connected to one another. In mean-field models, the identical oscillators respond to the same self-consistent signal. In the case of logistic (non-invertible) maps, the single units behave chaotically and separate exponentially from one another while keeping some correlations due to the common value of the forcing signal. In the case of Stuart-Landau oscillators, the single dynamical systems are periodic, but can and actually behave chaotically under the action of the mean field, so that the resulting scenario is qualitatively similar.

By moving further along the sequence of dynamical systems, one could ask whether CC can emerge in mean-field coupled identical phase-oscillators. On the one hand, it looks plausible to expect a positive answer, since the corresponding Liouville operator is nonlinear very much as for the previous setups. On other hand, one-dimensional phase oscillators never behave chaotically under the action of an “external” modulation. Actually, they are expected to synchronize with the input so that no irregularity can be maintained. In fact, we are not aware of any CC being generated by such models. The only seemingly contradictory evidence is that of pulse coupled theta neurons, where a chaotic collective dynamics indeed emerges. However, this occurs in the presence of delay, which makes the neuron dynamics no-longer one-dimensional [6].

If disorder is introduced in the form of a distribution of the connection strengths, the single units are no longer subject to the same mean-field and the above argument does not apply. In fact collective chaos can arise. Most of the studies of this type of models focused on balanced states, where the sum of positive and negative feedbacks average to zero (a prominent example is Ref. [7]). In this case, the model structure is significantly different, since the sum of all feedbacks must be scaled by a factor \sqrt{N} rather than by the usual N (N is the number of oscillators) to keep the feedback strength finite in the thermodynamic limit. Given the substantially different action of the feedback, we do not discuss this type of models any further in the paper. Disorder can be also due to the diversity among the single units still under the action of a common mean field. The above argument about the impossibility to sustain an irregular dynamics in one-dimensional phase-oscillators has to be revisited. In fact, even though all oscillators tend to “synchronize” because of a negative conditional Lyapunov exponent (conditioned to a given feedback signal), each of them synchronizes in a different way, leading, for instance to different effective frequencies. Therefore, it is not obvious a priori whether an irregular dynamics can be maintained. The Kuramoto model [8] is the best known example of such a type of setup. However, no chaotic collective dynamics emerges therein. CC has been instead found in a model of leaky-integrate-and-fire neurons [9]. One might argue that this is due to the presence of delay as in reference [6]. However, the phenomenology is remarkably different in the neural model, since the microscopic Lyapunov exponents are all negative (with the exception of the first exponent, that is exactly equal to zero). This type of behavior goes under the name of stable chaos [10]. In finite systems it is confined to exponentially long but finite times. Finally, CC has been observed also in a model of

pulse coupled phase oscillators, with no delay and without any discontinuity in the evolution equation [11].

In this paper, we thoroughly analyse the emergence of CC in mean-field models of one-dimensional oscillators. In the context of identical units we introduce a new way of representing the microscopic configurations which allows for a natural comparison between the microscopic and macroscopic world. In particular, one can test to what extent macroscopically small perturbations can be assimilated to microscopically small ones and thereby understand whether the macroscopic linear stability analysis is contained in the microscopic analysis. In Section 2, we first introduce the new approach, based on a reordering of the microscopic variables, and then investigate a simple mean-field model, where all typical indicators of instability can be and actually are implemented. In Section 3 we briefly revisit a model of pulse coupled phase-oscillators to investigate the robustness of CC in a context where the microscopic dynamics cannot be chaotic in itself: in particular we test the robustness against variations in the amount of disorder and in the shape of the phase response curve. Finally in Section 4 we discuss the consistency among the various approaches, summarize the current understanding, and propose new lines of research for further progress.

2 Identical oscillators

In this section we revisit the problem of how microscopic and macroscopic dynamics are related to one another. The direct simulation of large but finite ensembles suffices to characterize the behavior of a given system in the thermodynamic limit. Therefore, it is legitimate to ask whether the standard Lyapunov exponents which allow, for instance, determining the fractal dimension of the microscopic attractor, are of any usefulness for the characterization of the collective behavior, as well.

In the literature one finds different claims. On the one hand, it has been found that in some model, the most relevant exponents which control the stability of the collective motion can be singled out within the microscopic Lyapunov spectra [12]. In other cases, it has been found that the macroscopic stability is controlled by the evolution of finite perturbations [13].

The microscopic configuration of a system composed of N single dynamical units is fully identified by the N -dimensional vector $\mathbf{U}(t) = \{U_1(t), U_2(t), \dots, U_i(t), \dots, U_N(t)\}$, where U_j is a scalar variable which takes values in a given interval. The corresponding evolution rule is (if time is discrete)

$$U_j(t+1) = F(U_j(t), Z), \quad (1)$$

where we assume that all elements are identical and that the coupling is contained in the mean field Z , a function of all the single variables (typically an average).

On the other hand, on the macroscopic level the object of study is the probability density $\rho(U, t)$ to find a single variable in the interval $[U, U + dU)$ at time t . Its evolution equation is the nonlinear Perron-Frobenius operator,

$$\rho(U, t+1) = \sum_i \frac{\rho(F_i^{-1}(U, Z, t), t)}{|F_i'|} \quad (2)$$

where i labels the possibly multiple preimages of U . The nonlinearity of the rule is contained in the (yet unspecified) dependence of Z on ρ .

The two descriptions are manifestly different from one another, starting from the phase-space dimension, which is finite (N) dimensional in the former case, while it is




Fig. 1. Microscopic representation of the cumulative distribution function for an ensemble of 10,000 uncoupled logistic maps ($a = 1.86$).

infinite dimensional in the latter one. In order to carry on a quantitative comparison of the two representations it is at least necessary to identify an observable that is sufficiently meaningful and informative in both setups. We find that the probability density is an appropriate variable, since it can be suitably defined also in the microscopic representation. For this to be seen, it is sufficient to reorder the variables from the minimum to the maximum one, to obtain \mathbf{V} , where $V_j(t) = U_{n(j)}(t)$ and $n(j)$ denotes the label of the j th smallest variable at time t . Given that the single maps are indistinguishable, this representation is microscopically equivalent to the standard one. Moreover, j/N , if plotted versus V_j can be interpreted as a cumulative distribution function,

$$\mathcal{P}(U, t) = \int_{V_{min}}^U dV \rho(V, t) \quad (3)$$

of the macroscopic probability $\rho(V, t)$.

This is shown in Figure 1, where j/N is plotted versus $V_j(t)$ for an ensemble of uncoupled logistic maps, $U_j(t+1) = 1 - aU_j^2(t)$. Over the scales that can be resolved by the human eye, such a curve is indistinguishable from a typical snapshot obtained by integrating the Perron-Frobenius operator. In order to uncover the microscopic character of j/N , it would be necessary to increase the resolution. Anyway, the reordering of the local variable allows performing a microscopic analysis of an observable (the density) that is intrinsically macroscopic¹. In fact, since the vector \mathbf{V} can be used as a kernel density estimator, the distance between microscopic configurations can be turned into a distance between cumulative distribution functions.

It is instructive to understand how the map operates, once the ordered representation is adopted. In Figure 2a a configuration composed of 7 variables is reported as a sequence of properly ordered dots along the horizontal axis. The iterate of each variable is determined by the height of the function (see the corresponding vertical dashed lines). Because of the shape of the map, the ordering of the points on the left of the maximum (1–4) is preserved, while the order of the following ones (5–7) is inverted. The complexity of the dynamics comes from the fact that the two sequences intertwine, so that, for instance, the image of the 7th point has to be relabelled as the new 2nd one. In so far as one is interested in performing a linear stability analysis,

¹ Here the one-dimensional character of the dynamical units is crucial. The same idea cannot be straightforwardly implemented in two-dimensional oscillators.




Fig. 2. Interpretation of the logistic map evolution in terms of ordered variables. The initial ordering (1–7) is transformed into (1,7,2,6,3,5,4) in one iterate (a). Panel (b) contains the image of the 6th variable while it moves in between the 5th and 7th one while leaving all the other variables unchanged. The two kinks correspond to the points where the images of the 3rd and the 2nd point are crossed.

i.e., if only infinitesimal perturbations are introduced, the label used to identify each variable is totally irrelevant: the same Lyapunov exponents are obtained as they are dynamical invariants. The story is different if, however, finite perturbations are considered, since the iterates of the two configurations can be characterized by a different ordering. Imagine, for instance that the configuration in Figure 2a is perturbed by shifting V_6 . By definition, V_6 is bound to take values between V_5 and V_7 . Within this range, its iterate can take either values larger than the image of V_3 or smaller than the image of V_2 : whenever this happens a relabelling has to be performed. As a result, the actual “true” iterate of V_6 follows the curve reported in panel (b).

The order exchanges induce two kinks, which can be interpreted as a sort of nonlinear effect, reminiscent of what happens in several models displaying stable chaos [10]: the crucial difference is that here an order exchange does not amplify an otherwise small perturbation, but it rather contributes to squeeze it.

In order to compare generic microscopic configurations, it is convenient to introduce the Euclidean distance

$$\delta = \sqrt{\frac{1}{N} \sum_j (V_j - W_j)^2}, \quad (4)$$

where both V_j and W_j are properly ordered.

From now on, for a while, we assume that the maps evolve independently (no coupling is present). We proceed by evolving a set of N randomly chosen initial conditions until they are distributed according to the invariant measure. Let \mathbf{V} denote a typical configuration once the transient has died out, and perturb it by a tiny amount (very small, but large enough to be accounted for in double-accuracy calculations), to thereby generate the configuration \mathbf{W} . We then evolve the two configurations, reordering them (separately) wherever necessary and measure the Euclidean distance, which is thereby averaged over a set of many different initial conditions to get rid of statistical fluctuations. The results for a set of logistic maps are plotted in Figure 3. In the upper panel we report the evolution of the logarithm of the average distance δ for four different ensembles of orbits ($N = 10^3, 10^4, 10^5$, and 10^6). There, we see that an initial exponential increase, consistent with the chaotic properties of the map, is followed by a saturation due to the fact that the maximal distance is finite. We can also see that the growth rate starts to slow down at progressively shorter times, when N is increased. A clearer picture is obtained by looking at the “instantaneous”




Fig. 3. (a) Time evolution of the average distance between densities as determined from equation (4), for $N = 10^3$ (black), 10^4 (red), 10^5 (green), and 10^6 (blue). Averages are performed over 10000 realizations. In panel (b), the evolution of the rate S (see Eq. (5)) is plotted. Finally in the remaining panels (c–e), S is plotted versus δ for different scaling of the distance.

growth rate of δ , i.e., at the logarithmic derivative

$$S = \ln \delta(t+1) - \ln \delta(t). \quad (5)$$

The time dependence of S can be appreciated in Figure 3b. All curves start from the same plateau (leaving aside the first few iterates which are needed for the perturbation orientation to settle). This plateau is nothing but the Lyapunov exponent of the single map. Interestingly and unexpectedly we see also a second plateau at later times, whose length progressively increases with N , and whose height is about half of the standard Lyapunov exponent. Such a plateau is eventually followed by a convergence to zero that occurs in the same way in all cases, i.e. independently of N .

A more appropriate representation is obtained by plotting the growth rate S versus the current value of δ . This is done in panels (c–e) after rescaling the distance in different ways. In panel c, no rescaling at all is implemented: there we appreciate that the drops in the growth rate occur at progressively smaller distances, when N is increased. Its scaling behavior can be understood in the two following panels. In panel (d) we see that the first drop occurs when the distance is about $1/N$. It is natural to expect the drop to occur when a large enough number of order exchanges have occurred, which damp the standard exponential growth through the saturation outlined in Figure 2. It is also reasonable to assume that order exchanges occur when the separation between the two trajectories is of the same order as the separation between consecutive variables, i.e., $1/N$. A simple calculation shows that assuming $(V_i - W_i)$ is on the order $\mathcal{O}(1/N)$ implies that δ itself is of the same order. In panel (e) we finally see that the second drop, which leads to the final saturation, occurs when $\delta \approx 1/\sqrt{N}$. The collapse of the final part of the curve means that the (asymptotic) maximal distance is of order $1/\sqrt{N}$ i.e., it is just due to statistical fluctuations. Remarkably, the range of distances between the two drops is characterized by a constant rate




Fig. 4. Time evolution of the distance between two microscopic configurations: (i) chosen at equilibrium and differing by 10^{-13} in each component (lower set of curves, same as in Fig. 3a); (ii) one at equilibrium, while the other one confined to the initial interval $[-0.3, -0.1]$. The different curves correspond to different ensembles: $10^3, 10^4, 10^5, 10^6$, from top to bottom. The two exponential behaviours correspond to the largest microscopic Lyapunov exponent ($\lambda = 0.564$, dashed lines), and the largest macroscopic one ($\lambda = -0.188$, solid curve).

and its extension grows upon increasing N : this suggests that the microscopic and macroscopic worlds are separated by an intermediate *mesoscopic* regime.

In this setup, we started with two configurations that are microscopically different and monitored their distance until it becomes maximal. From a macroscopic point of view, the two configurations are, however, indistinguishable from the very beginning as they both correspond to an “equilibrium” state. It is interesting to test this approach while starting from macroscopically different configurations. More precisely, we choose the configuration \mathbf{V} still as an equilibrium one, while W is generated as a random sequence of independent values restricted to the interval $[-0.3, -0.1]$. The results after averaging over 10,000 realizations for different ensemble sizes are plotted in Figure 4 (see the circles) and compared to the previous setup (dashed curves).

There we see that for any ensemble-size the curves obtained with the two setups eventually converge to the same value as they should since the final state corresponds to two randomly picked microscopic realizations of the same macrostate. The convergence is obviously from above in this latter setup. More interestingly, the initial decrease is common to all sizes, suggesting the method is capturing a true collective property. We indeed claim that the initial slope, around -0.188 , corresponds to the largest (least negative) eigenvalue of the Perron-Frobenius operator.²

2.1 Nonlinearly coupled Bernoulli maps

So far, we have considered an ensemble of uncoupled maps, where no CC is expected to arise. Here, we introduce and discuss a simple model of nonlinearly coupled Bernoulli maps (NCB)

$$\theta_j(t+1) = F(\theta_j, Z) = 2[\theta_j(t) + \varepsilon(1 - 2a^2 Z^2) \sin \theta_j], \quad (6)$$

² In the case of uncoupled maps, only macroscopic relaxational phenomena are to be expected.




Fig. 5. Four snapshots of the probability density of the NCB model. The black, red, green, and blue curves correspond to consecutive time units.




Fig. 6. Rate S for the NCB model, versus the distance for three different scaling of δ .

where

$$Z = \frac{1}{N} \sum \cos \theta_j(t) \quad (7)$$

and the variable θ is an angle. As a result, the typical discontinuity of the Bernoulli maps is smoothed out, since θ is equivalent to $\theta + 2\pi$. The NCB model is a bit ad hoc, since the coupling is nonlinear and one might argue that the origin of the collective dynamics is somehow artificial. However, the question is not to explain the origin of collective behaviour but to characterize this regime. The NCB model has been numerically studied for $a = 2.2$ and $\varepsilon = -1.15$. The presence of a collective dynamics is testified by the fluctuations of the density that one can appreciate in Figure 5, where the probability density is reconstructed by iterating an ensemble of $N = 10^6$ maps, ordering all the variables and finally computing the “derivative” after partitioning the θ axis in disjoint intervals containing 10,000 points each. The four curves refer to four different consecutive times, sampled after a suitable transient has elapsed. Macroscopic differences are transparent.

By following the same protocol adopted for the logistic maps, we have again determined the instantaneous growth rate of the average distance δ between initially close distributions. The results are plotted in Figure 6 for three different scaling assumptions. The unscaled data are shown in panel (a). The mutual comparison reveals that $N = 10^3$ is still affected by strong finite size corrections. Other than that, a two-plateau structure is visible also in this case, although the lower one is not perfectly flat and tends to develop a bump upon increasing N . In panel (b), we see that the first drop scales again as $1/N$, while the second drop does not decrease as $1/\sqrt{N}$. This is a consequence of the presence of a macroscopic dynamics. In fact, according to




Fig. 7. (a) Temporal evolution of the average d_c (dashed line) and d_s (solid). Data are obtained averaging at fixed times over 10,000 realizations. Black, red, green and blue lines correspond to $N = 10^3$, 10^4 , 10^5 and 10^6 , respectively. (b) The logarithmic derivative of the curves reported in panel (a). The horizontal segment marks the collective Lyapunov exponent.

the rescaling adopted in panel (c) a collective motion manifests itself as δ -values that grow when multiplied by \sqrt{N} . Actually, the range of values approximately above zero quantifies the range of macroscopic deviations that is accessible in the microscopic simulations. Part of the lower plateau should be associated with the mesoscopic region, so that it is not clear how strong is the macroscopic instability. The study of larger ensembles is necessary for a more quantitative assessment.

The logarithmic derivative S quantifies the exponential growth of finite distances between nearby distributions. A similar quantity, the finite-amplitude Lyapunov (FAE) exponent [14], has been introduced with reference to the difference between scalar observables such as the order parameter which quantifies the strength of the collective dynamics. In both cases, one deals with finite perturbations. The difference is that the former approach deals with high-dimensional variables.

For the sake of completeness, we now discuss the outcome of the latter approach, as well. This is done by considering two observables: $Z = \langle \cos \theta \rangle$ and $\Xi = \langle \sin \theta \rangle$. Z is the order parameter entering the definition of the NCB model and its deviations are denoted with d_c . The deviation of Ξ is instead denoted with d_s . The average perturbation amplitude at different times and for different values of N is plotted in Figure 7a. There, we see that both distances initially grow exponentially and eventually saturate. The saturation value of d_c is independent of N , while that of Ξ reaches increasingly small values when N is increased. This tells us that only Z is able to capture the collective dynamics. In fact, for symmetry reasons, $\Xi = 0$ (see, e.g., Fig. 5) so that the fluctuations displayed by the finite averages are purely statistical.

More detailed information can be extracted from the time-dependence of the FAE, i.e., the logarithmic derivative

$$\Gamma = \ln d(t+1) - \ln d(t) , \quad (8)$$

reported in Figure 7b. There we clearly see a different behavior for the stability of Z and Ξ : d_s grows from the very beginning at the maximum rate, while d_c slows down before recovering and growing with the same rate as d_s . We justify the discrepancy in the following terms: d_c and d_s represent two different ways of measuring the distance between two configurations. The initially different growth rate exhibited by d_c is a purely linear effect due to the orientation of the overall vector: if the true norm of a given vector grows exponentially, it may well be that some component grows




Fig. 8. Lyapunov spectrum of the NCB model for $N = 160$ (black line), 320 (red), 640 (green), and 1280 (blue). The dashed line corresponds to the conditional Lyapunov exponent of a single map subject to an externally given $Z(t)$. In the inset the first (upper curve) and second (lower curve) Lyapunov exponents are plotted versus the network size.

slower (or even faster) if they are (accidentally) negligible. Anyway, after an initial transient, both variables exhibit the same (first) plateau, which coincides also with the similar plateau observed in Figure 6, where it was determined by analysing the full distribution. Furthermore, we can anticipate that the height of the plateau is also consistent with the maximum microscopic Lyapunov exponent (see Fig. 8 and the final section for an additional discussion). At later times, the growth rate of d_s drops to zero, showing that nonlinear effects are only responsible for a saturation, while the rate of d_c exhibits a second short but increasing plateau with approximately the same height as the second plateau displayed in Figure 6. Finally, the height of this second plateau is approximately close but a bit larger than the expected value of the macroscopic exponent (see the next subsection).

2.2 Lyapunov exponents

The microscopic Lyapunov exponents can be computed by linearizing the equations (6), (7). The results are plotted in Figure 8 under the usual scaling Ansatz, i.e., by ordering the Lyapunov exponents from the largest to the smallest one and plotting them versus the intensive variable $\nu = (i - 1/2)/N$ ³.

The shape of the spectrum varies with the system size. In the absence of collective dynamics it was found that the asymptotic Lyapunov spectrum is perfectly flat with two “singularities” at its extrema [16]. The value of the plateau coincides with the conditional single-map Lyapunov exponent, obtained by treating the mean field $Z(t)$ as an external modulation. On the other hand, the deviations observed for ν close to 0 and 1 are the consequence of an extreme sensitivity to the coupling [15–17]. Furthermore, finite-size deviations can be of order $1/(\ln N)$. The scenario reported in Figure 8 does not look very different. The dashed line corresponds to the conditional single-map Lyapunov exponent: the Lyapunov spectrum seems to slowly converge towards this value (the deviations are not incompatible with a $1/(\ln N)$ convergence). Nevertheless, there is a difference: a finite gap between the first and the second Lyapunov exponent. This is clearly visible in the inset of Figure 8, where the first two exponents are plotted versus the system size: while the first one increases with

³ The 1/2 term is included to minimize finite-size effects [15].




Fig. 9. Collective Lyapunov spectrum for the NCB model.

N , the second exponent decreases, indicating that the finite gap persists in the thermodynamic limit. It is not clear whether the same occurs also close to the minimum exponent (i.e., for $\nu \approx 1$). If we now go back to the first plateau found while studying finite perturbations (see Figs. 6 and 7), one can notice that their height, approximately equal to 0.76, is substantially smaller than the maximum Lyapunov exponent (≈ 0.93). In the final section we give possible explanations for this difference. Finally, notice that the whole spectrum is composed of positive exponents, this meaning that the dimension of the microscopic attractor coincides with the phase-space dimension (this pathology is due to the non invertible nature of the single maps).

Besides microscopic Lyapunov exponents, one can compute the macroscopic ones from the linearization of the Perron-Frobenius equation. The advantage of the NCB model is the smoothness of the invariant measure (see the snapshot reported in Fig. 5), accompanied by the angular character of the local variables. Therefore, it is convenient to expand the evolution equation in Fourier modes, since a few of them suffice to provide an accurate description. Let us start formally expanding the probability density $\rho(\theta, t)$,

$$\rho(\theta, t) = \sum_k \psi_k(t) e^{ik\theta}.$$

The corresponding Perron-Frobenius equation can be written as

$$\psi_k(t+1) = \sum_m R(k, m) \psi_m(t), \quad (9)$$

where

$$R(k, m) \equiv \frac{1}{2\pi} \int_0^{2\pi} d\theta e^{i(mU - kF(\theta, Z))}, \quad (10)$$

while

$$Z = \int \cos \theta \rho(\theta, t) d\theta. \quad (11)$$

More specifically,

$$R(k, m) = J_{m-2k}[2k\varepsilon(1 - 2a^2Z^2)], \quad (12)$$

where J_α denotes the Bessel function of order α . In the present case, $Z = \psi_1$ and all ψ_k are real. As a result, the Perron-Frobenius operator reads

$$\psi_k(t+1) = 2k\varepsilon \sin(\psi_1(t)) \sum_m J_{m-2k} \psi_m(t). \quad (13)$$




Fig. 10. Return map of the order parameter Z as obtained by simulating an ensemble of 10^6 maps.

Upon iterating the above equations together with their linearizations, one can easily compute the Lyapunov spectrum: the first 10 exponents are plotted in Figure 9. From the figure we infer that the collective motion is low-dimensional: only the first mode is positive and the expansion is compensated by the second exponent ($\lambda_1 + \lambda_2 < 0$). In fact the dimension, as determined from the Kaplan-Yorke formula, is $1 + \lambda_1/|\lambda_2| = 1.63$. Evidence of the low-dimensional dynamics can be seen in Figure 10, where we plot the return map of the order parameter Z . Statistical fluctuations are hardly visible.

The Lyapunov spectrum is discrete; upon increasing the number of Fourier modes, additional negative exponents add up, which correspond to increasingly stable, high-frequency, modes. In practice, the scenario is reminiscent of the continuum-limit in spatially extended chaos. This is not, however, always the case. Coupled logistic maps appear to be characterized by an infinite dimensional collective dynamics: this has been conjectured by adding noise to the microscopic dynamics (in order to smooth out the singularities of the density) and thereby decreasing the noise amplitude [12, 18]. As a result, it has been found that the dimension increases logarithmically upon decreasing the noise amplitude.

3 Oscillator diversity

CC can emerge also in mean-field models of one-dimensional invertible dynamical systems, under the condition that the oscillators are different from one another [9, 11]. The simplest such setup consists of an ensemble of pulse-coupled phase oscillators,

$$\dot{\phi}_i = \omega_i - \frac{g}{N} R(\phi_i) \sum_j \delta(t - t_j), \quad (14)$$

where $\phi_i \in [0, 1]$ is a phase variable and g the coupling strength. In the absence of coupling, the phase of each oscillator advances according to a (quenched) random bare frequency ω_i . When the coupling is switched on, a δ -like pulse is sent whenever an oscillator reaches the threshold $\phi_i = 1$ and received by all other oscillators (a global




Fig. 11. Different choices of the piecewise linear phase-response curve (PRC). The PRC is symmetric between positive and negative response. The length of the second piece determines the smoothness with δ between 0.0002 and 0.16.

all-to-all coupling is assumed). Finally, the effect of the pulse is gauged by the phase response curve (PRC) $R(\phi)$.

A complete definition of the model requires introducing the distribution of bare frequencies and the PRC. By following reference [11], we assume a uniform distribution within a finite interval $\Delta = \omega_{\max} - \omega_{\min}$ centred in $\bar{\omega} = 1.4$ and a piecewise linear PRC. Rather than giving a not-so-transparent mathematical definition, we illustrate the shape of the PRC in Figure 11. The width δ of the central segment will be used as control parameter to test the role of the PRC shape.

Due to the negative sign in front of the coupling term, wherever $R(\phi)$ is larger (smaller) than zero the coupling has an inhibitory (excitatory) effect, decreasing (increasing) the phase of the receiving neuron. Since preliminary simulations have shown that a change of the average spiking rate does not affect the quality of the overall phenomenology, we have chosen to work with zero-average PRCs. However, this is not a necessary requisite.

The presence of CC can be detected by using various indicators. Given the close relationship with neural models, the *activity* is an appropriate observable to start with. It is defined as

$$\dot{Y} = -\gamma Y + \frac{1}{N} \sum_j \delta(t - t_j),$$

where $\gamma = 5$ defines the time scale of the coarse-graining. For small coupling, the diversity among the oscillators prevails and they fire independently: as a result, the activity is constant (except for statistical fluctuations). Above some threshold, the activity exhibits macroscopic fluctuations. In Figure 12 we plot the standard deviation σ_y of Y versus the coupling strength for various choices of δ (panel a) and of the width Δ of the probability distribution of the bare frequencies (panel b). All data refer to the same system size ($N = 4000$); for some points, we have verified that no substantial changes are found when N is increased, so that the overall scenario can be taken as representative of the thermodynamic limit.

In Figure 12a we see that the parameter region where CC is observed tends to decrease when the steepness of the intermediate branch is decreased. However, the phenomenon is not monotonous (see the three smallest δ values). Interestingly, in panel (b) we see that upon decreasing the disorder, CC extends to progressively smaller coupling strengths. These results show that collective dynamics is a general




Fig. 12. Standard deviation σ_Y of the activity versus the coupling strength g for fixed frequency disorder $\Delta = 1.2$ and different choices of the PRC (a) and fixed PRC ($\delta = 0.1$) but different width of the disorder (b). All simulations refer to a network of $N = 4000$ oscillators and other parameter are as in [11].




Fig. 13. Power spectra of network activity Y for $g = 1.0$ obtained from two different network sizes ($N = 4000$ black, and $N = 64000$ green). The spectra are obtained by transforming time series of 819.175 time units, sampled every 0.025 units and averaged over 50 different realisations.

robust regime and suggest the possibility to develop a perturbative theory in the limit of weak-disorder and small coupling.

A finite standard deviation reveals the presence of a macroscopic dynamics, but it does not tell us how irregular it is. In Figure 13 we report the power spectrum of the activity for different network sizes. We see that it is broad band and substantially independent of the system size (the width of the major peaks does shrink when N is increased). A fractal dimension analysis carried out in [11] has shown that the macroscopic dynamics is high- and possibly infinite-dimensional.

A crucial difference between this regime and that one discussed in the previous section is the microscopic stability of the oscillator dynamics under the action of the mean field Y . This stability can be appreciated by looking at the structure of a snapshot of the distribution of the phases in the CC regime.

In Figure 14a we see that in suitable ranges of the bare frequencies the phases are aligned along seemingly one-dimensional curves. This is a clear indication of a mutual




Fig. 14. (a) Snapshot of the oscillator phases for $N = 64000$, $g = 1.0$, $\delta = 0.1$ and $\Delta = 1.2$; (b) projection onto the ϕ axis to determine the effective probability density.

synchronization that is confirmed by the structured distribution of phases displayed in panel (b) of the same figure. Furthermore, synchronization suggests the stability of oscillator dynamics. A careful analysis of the conditional Lyapunov exponent shows that it is always strictly negative including those frequency regions where the phases are scattered over the $[0, 1]$ interval [11]. At the same time, the microscopic Lyapunov exponents reveal a more structured scenario: the maximum exponent λ_1 is negative for a strictly discontinuous PRC ($\delta = 0$), but it grows with δ , becoming, at some point, positive (for $\delta = 0.1$ it is already so⁴). Furthermore we have found that λ_1 is affected by strong finite size corrections. Therefore, it is not easy to extract useful information from a stability analysis of the type discussed in the previous section. An additional difficulty is given by the presence of sample-to-sample fluctuations that are still nonnegligible in a network of $N = 4000$ oscillators. Here, we have limited ourselves to explore the behavior of small finite perturbations. Given any two close configurations $\{\phi_i\}$ and $\{\psi_i\}$, we have monitored

$$D = \left\langle \sum_i |\cos \phi_i - \cos \psi_i| \right\rangle,$$

where the external average is performed over different initial conditions. The results for $N = 4000$ and $N = 16000$ are reported in Figure 15, where the first thing to notice is the slower growth for $N = 16000$, which is indeed an evidence of the strong finite-size effects, which make any further analysis (such as the identification of two possible slopes) essentially undoable. The spikes, induced by the steep branch, represent an

⁴ The claim reported in reference [11] of a negative exponent for such parameter value is to be considered wrong. It is the consequence of a numerical error. The submission of an Erratum is on the way.




Fig. 15. Time evaluation of the distance between the original and perturbed phases. Each curve has been averaged over 100 different random perturbations for the same set of frequencies $\{\omega_i\}_i$. Each phase Φ_i has been perturbed at time $t = 0$ by adding a uniformly distributed number $U(-10^{-12}, 10^{-12})$. The two lower curved (blue and red) belong to two different sets of $\{\omega_i\}_i$ of size $N = 16000$ whereas the three upper (black, orange, cyan) belong to different samples of size $N = 4000$.

additional obstacle for a more quantitative analysis: they suggest the need to average over much more than 100 initial conditions.

4 Discussion and open problems

In this paper we have revisited the properties of the collective dynamics which emerges in mean-field models of one-dimensional oscillators. So long as the oscillators are identical, the scalar variables, which fully characterize the instantaneous configuration of the single oscillators, can be freely permuted without inducing any change in the evolution. One such permutation is particularly suited to represent a configuration of a set of N oscillators: it corresponds to ordering all the variables from the smallest to the largest one at each time step. Its appeal comes from the fact that any finite perturbation of the single variables can be interpreted as a perturbation of the corresponding probability density and a distance between any two configurations interpreted as a variation of the density itself. This interpretation would be impossible if the label of each oscillator were left unchanged. While the change of labels does not modify the tangent space dynamics (the Lyapunov exponents are dynamical invariants) it does affect the evolution of finite perturbations: the reason is precisely the relabelling of the variables that may be different for the two configurations that are to be compared. As illustrated in Figure 2, this change of labels manifests itself as discontinuities in the first derivative of the actual evolution rule, which turns out to be piecewise smooth. The typical length of a smooth interval scales as $1/N$: this suggests (and it is confirmed by our simulations) that as soon as a perturbation amplitude becomes larger than $1/N$, it does no longer evolve according to the rules of tangent space dynamics. The same analysis shows that the presence of collective dynamics can be inferred only by monitoring distances larger than $1/\sqrt{N}$. A posteriori, this is obvious, since distances must be larger than the (unavoidable) statistical fluctuations which affect microscopic configurations. Additionally, we have detected a “mesoscopic” range (distances which lie between $1/N$ and $1/\sqrt{N}$), where neither the microscopic nor the macroscopic rule applies. In the case of uncoupled logistic




Fig. 16. Schematic view of a perturbation of a configuration (see the arrows) which can be seen as a perturbation of a (coarse-grained) density.

maps, the mesoscopic phase is not a simple crossover, but a region characterized by its own physical observables. The evidence for this phase in the NCB model is less compelling: further studies are necessary to better understand how far down to small scales can the macroscopic rules extend.

When comparing the various approaches (stability of the distribution, FAE, microscopic Lyapunov exponents, collective Lyapunov exponents), the observables that are expected to match each other do so, but some deviations have been found which need to be better clarified. Here we suggest two possible motivations. First of all, the early-time growth rate of small but finite perturbations may not coincide with the maximum Lyapunov exponent because a possibly too-short transient does not allow for a perfect alignment of the perturbations themselves. Moreover, the very same growth rate has been determined by implementing a simple but not optimal method. For the sake of simplicity, we have averaged the perturbation sizes of different realizations, all estimated at the same times. It would be more appropriate to first detect the time when a given perturbation amplitude reaches a prescribed threshold and thereby average the growth rate, as prescribed by the definition of the FAE [14,15]. Preliminary tests made by applying this latter procedure to monitor the distance between probability densities in the uncoupled case confirm the overall picture, i.e., the existence of three regimes (micro-, meso-, and macro-scopic) and the height of the plateaus.

A further direction that might be worth exploring is the possibility to go beyond the computation of “just” the maximum finite-amplitude Lyapunov exponent. This is not typically doable in a meaningful way because an ensemble of perturbed trajectories does not fill a linear subspace. Therefore linear techniques (such as scalar products) can only extract rough, qualitative, information. However, in the case of CC in mean-field models, perturbations larger than $1/\sqrt{N}$ (but still small) fall in the linear regime of the macroscopic evolution. Therefore, it might be possible to introduce a (macroscopic) scalar product that is robust to the presence of statistical fluctuations. The fact that with 10^6 maps, we are able to produce a quantitatively accurate picture of the macroscopic attractor, whose fractal dimension depends on the second (negative) collective Lyapunov exponent, encourages the exploration of this route.

The dynamics of finite perturbations as it emerges from Figure 2b seems to rule out the possibility that the linear stability of a collective variable is contained in the linear stability of the microscopic equations: the macroscopic stability emerges on scales of order $1/\sqrt{N}$, well above the limit where the microscopic linearization stops being valid, which is of order $1/N$. However, one cannot dismiss the results contained in reference [12], based on the structure of the microscopic covariant Lyapunov vectors [19–21].

A covariant Lyapunov vector is the extension of the concept of an eigenvector of a standard linear operator to the product of time-varying operators. A Lyapunov vector is a collection of (infinitesimal) perturbations of the single variables and, as such, it can be seen as perturbation of the corresponding density. This interpretation is qualitatively illustrated in Figure 16, where a microscopic configuration is represented by the position of 8 variables (see the top squares), while a Lyapunov vector is depicted as a




Fig. 17. Snapshot of two covariant Lyapunov vector corresponding to the largest (a) and smallest (b) microscopic Lyapunov exponent for a the NCB model. The upper curve is obtained for $N = 20480$, while the second corresponds to $N = 10000$. In both cases the vector amplitude is scaled to an Euclidean norm N .

sequence of arrows which identify the shift of the corresponding variables. The resulting perturbed configuration is finally represented by the bottom sequence of squares.

The perturbation shown in Figure 16 can be interpreted as an enhancement of the density in the central part of the interval. A necessary condition for a perturbation of the microscopic variables to correspond to a meaningful perturbation of the density is that the microscopic variations cooperate in regions of size larger than $1/\sqrt{N}$ to build macroscopically observable variations. As shown in reference [20], high-dimensional chaotic systems are typically characterized by localized Lyapunov vectors, i.e., by perturbations which are typically localized in a non-extensive set of points. Accordingly, they cannot be interpreted as macroscopic perturbations and, in particular, they cannot correspond to macroscopic Lyapunov vectors. Extensivity is, therefore, a requisite for a covariant vector to have a macroscopic interpretation. In fact, extensive vectors have been found in some mean-field models [12], especially in an ensemble of noisy logistic maps, where one could make a comparison with the Frobenius-Perron operator. So it is legitimate to ask how it is possible that in some cases the saturation induced by the reordering of the microscopic variables (see Fig. 2) does not affect the evolution of small perturbations at least up to scales of order $1/\sqrt{N}$. If nearby variables are coherently perturbed, their mutual distance does not vary significantly: in principle, such perturbations may lead to a few changes of order, so that the microscopic linearized dynamics can extend to scales where it coincides with the behavior of macroscopically perturbed probability densities. This is precisely the reason why macroscopic and microscopic stability analyses give consistent results for the stability of splay states. A splay state is an asynchronous regime where globally coupled oscillators are evenly distributed over the set of possible phases and maintain their initial ordering, so that no oscillator is ever “overtaken”. Differences are obviously found in the corresponding spectra of eigenvalues, but they can be attributed to the finiteness of microscopic ensembles: other than that, as shown in references [22,23], there is an overall consistency. Whether the absence of variable reordering is a necessary and sufficient condition for a microscopic exponent to percolate to macroscopic scales, is a conjecture that need to be further explored.

We now return to the NCB model for some further tests. From Figures 8 and 9 it follows that the largest collective Lyapunov exponent $\Lambda_{max} \approx 0.335$ is located within the spectrum of microscopic exponents, which ranges from $\lambda_{max} \approx 0.93$ down to $\lambda_{min} = 0.335$. Therefore, we do not expect the first covariant vector to exhibit any

macroscopic structure. A typical snapshot of this vector is illustrated in Figure 17a, plotting its local amplitude u versus the current value of the corresponding angular variable (u is scaled so that the Euclidean norm of the vector is N). There we see that many components are of order $\mathcal{O}(1)$: this is a clear signature of the extensivity. Extensivity alone, is not however a sufficient condition for a vector to be macroscopically meaningful. If the sign of the various components are “randomly” distributed, as in Figure 17a, no macroscopic variation of the corresponding density can emerge from the application of this perturbation. It is nevertheless worth stressing that extensivity is a rather anomalous property of covariant vectors in high dimensional systems; they are typically localized [20]. This is, for instance, the case of the covariant vector that corresponds to λ_{min} , plotted in Figure 17b: there we see that most of the components are exponentially small. This localized structure itself is, however, unexpected. In fact Λ_{max} is rather close to λ_{min} and, on the basis of the observations reported in reference [12], one might expect to find evidence of the macroscopic instability in the structure of the usual microscopic Lyapunov exponents. This is an additional evidence that the relationship between macroscopic and microscopic instabilities is subtler than believed so far.

Finally, a comment on the onset of CC in ensembles of nonidentical oscillators. The model explored in Section 3 is affected by too strong finite-size corrections to be able to draw any quantitative conclusion. It would be desirable to identify a simple setup, characterized by a faster convergence, where one can quantitatively monitor the strength of the instability across a wide range of perturbation amplitudes.

APo and EU wish to acknowledge the Advanced Study Group activity at the Max Planck Institute for the Physics of Complex Systems in Dresden “From Microscopic to Collective Dynamics in Neural Circuits” for the opportunity to develop part of the project.

References

1. K. Kaneko, Phys. Rev. Lett. **65**, 1391 (1990)
2. V. Hakim, W.J. Rappel, Phys. Rev. A **46**, R7347 (1992)
3. N. Nakagawa, Y. Kuramoto, Prog. Theor. Phys. **89**, 313 (1993)
4. S. Luccioli, S. Olmi, A. Politi, A. Torcini, Phys. Rev. Lett. **109**, 138103 (2012)
5. G. Robb, A. Politi, Phys. Rev. E **95**, 040201(R) (2017)
6. D. Pazó, E. Montbrió, Phys. Rev. Lett. **116**, 238101 (2016).
7. H. Sompolinsky, A. Crisanti, H.J. Sommers, Phys. Rev. Lett. **61**, 259 (1988).
8. Y. Kuramoto, *Chemical Oscillations, Waves and Turbulence* (Springer, Berlin, 1984)
9. S. Luccioli, A. Politi, Phys. Rev. Lett. **105**, 158104 (2010)
10. A. Politi, A. Torcini, in *Nonlinear Dynamics and Chaos: Advances and Perspectives*, Underst. Complex Syst. (Springer, Berlin, 2010), pp. 103–129
11. E. Ullner, A. Politi, Phys. Rev. X **6**, 011015 (2016)
12. K.A. Takeuchi, H. Chaté, J. Phys. A **46**, 254007 (2013)
13. M. Cencini, M. Falcioni, D. Vergni, A. Vulpiani, Physica D **130**, 58 (1999)
14. E. Aurell, G. Boffetta, A. Crisanti, G. Paladin, A. Vulpiani, Phys. Rev. Lett. **77**, 1262 (1996)
15. A. Pikovsky, A. Politi, *Lyapunov Exponents*, 1st edition (Cambridge University Press, 2016)
16. K.A. Takeuchi, H. Chaté, F. Ginelli, A. Politi, A. Torcini, Phys. Rev. Lett. **107**, 124101 (2011)
17. H. Daido, Phys. Lett. A **121**, 60 (1987)
18. T. Shibata, T. Chawanya, K. Kaneko, Phys. Rev. Lett. **82**, 4424 (1999)
19. P. Bryant, R. Brown, H.D.I. Abarbanel, Phys. Rev. Lett. **65**, 1523 (1990)

20. F. Ginelli, P. Poggi, A. Turchi, H. Chaté, R. Livi, A. Politi, *Phys. Rev. Lett.* **99**, 130601 (2007)
21. C.L. Wolfe, R.M. Samelson, *Tellus A* **59**, 355 (2007)
22. S. Olmi, A. Politi, A. Torcini, *J. Math. Neurosci.* **2**, 12 (2012)
23. S. Olmi, A. Politi, A. Torcini, *Front. Comput. Neurosci.* **8**, 8 (2014)

Open Access This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.