

A spiral attractor network drives rhythmic locomotion

Angela M. Bruno^{1†} and William N. Frost^{2*} and Mark D. Humphries^{3*}

1. Department of Neuroscience, The Chicago Medical School, Rosalind Franklin University of Medicine and Science, North Chicago, Illinois, USA.
2. Department of Cell Biology and Anatomy, The Chicago Medical School, Rosalind Franklin University of Medicine and Science, North Chicago, Illinois, USA.
3. Faculty of Biology, Medicine, and Health, University of Manchester, Manchester, UK.

† Present address: California Institute of Technology, Division of Biology and Biological Engineering, Pasadena, California.

* Contact: william.frost@rosalindfranklin.edu; mark.humphries@manchester.ac.uk

Running title: A spiral attractor network drives locomotion

Abstract

The joint activity of neural populations is high dimensional and complex. One strategy for reaching a tractable understanding of circuit function is to seek the simplest dynamical system that can account for the population activity. By imaging *Aplysia's* pedal ganglion during fictive locomotion, here we show that its population-wide activity arises from a low-dimensional spiral attractor. Evoking locomotion moved the population into a low-dimensional, periodic, decaying orbit - a spiral - in which it behaved as a true attractor, converging to the same orbit when evoked, and returning to that orbit after transient perturbation. We found the same attractor in every preparation, and could predict motor output directly from its orbit, yet individual neurons' participation changed across consecutive locomotion bouts. From these results, we propose that only the low-dimensional dynamics for movement control, and not the high-dimensional population activity, are consistent within and between nervous systems.

Introduction

The increasing availability of large scale recordings of brain networks at single neuron resolution provides an unprecedented opportunity to discover underlying principles of motor control. However, such long-sought data sets are revealing a new challenge - the joint activity of large neural populations is both complex and high dimensional (Ahrens et al., 2012; Cunningham and Yu, 2014; Yuste, 2015). Population recordings have as many dimensions as neurons, and each neuron's activity can have a complex form. What strategies can we use to expose the hoped-for simplifying principles operating beneath the turbulent surface of real-world brain activity? One route is dimension reduction (Briggman et al., 2006; Cunningham and Yu, 2014; Kobak et al., 2016), which focuses on identifying the

25 components of activity that co-vary across the members of a neural population, shifting
26 the focus from the high dimensional recorded data to a low-dimensional representation of
27 the population.

28 Such low-dimensional signals within joint population activity have been described in
29 neural circuits for sensory encoding (Mazor and Laurent, 2005; Bartho et al., 2009),
30 decision-making (Briggman et al., 2005; Harvey et al., 2012; Mante et al., 2013), navi-
31 gation (Seelig and Jayaraman, 2015; Peyrache et al., 2015), and movement (Levi et al.,
32 2005; Ahrens et al., 2012; Kato et al., 2015). Implicit in such dimension reduction ap-
33 proaches is the hypothesis that the high-dimensional population activity being recorded,
34 while highly heterogenous, is derived from a simpler, consistent low-dimensional system
35 (Brody et al., 2003; Churchland et al., 2010; Kato et al., 2015; Miller, 2016). We sought to
36 directly test this hypothesis by identifying the simplest dynamical system that can account
37 for high dimensional population activity.

38 A useful model to address these questions is the neural control of movement. Movement
39 arises from the mass action of neuron populations (Georgopoulos et al., 1986; Getting,
40 1989; Ahrens et al., 2012; Portugues et al., 2014; Yuste, 2015; Petersen and Berg, 2016).
41 While individual neuron activity can correlate with specific aspects of movement (Chestek
42 et al., 2007; Hatsopoulos et al., 2007; Churchland et al., 2010, 2012), the embedded low
43 dimensional signals in population recordings (Briggman et al., 2005; Levi et al., 2005;
44 Kato et al., 2015) and the intermittent participation of individual neurons across repeated
45 movements in both vertebrates (Carmena et al., 2005; Huber et al., 2012) and invertebrates
46 (Hill et al., 2010, 2015) together suggest that only the collective population activity, and
47 not specifics of single neuron firing, are key to movement control. If so, then finding the
48 underlying dynamical system will be necessary for a parsimonious theory of the neural
49 control of movement (Briggman and Kristan, 2008).

50 In order to identify the simplest dynamical system underlying population activity in
51 movement control, we imaged large populations at single-neuron, single-spike resolution
52 in the pedal ganglion of *Aplysia* during fictive locomotion (Figure 1A). The pedal gan-
53 glion presents an ideal target for testing hypotheses of movement control as it contains the
54 pattern generator (Jahan-Parwar and Fredman, 1979, 1980), motoneurons (Hening et al.,
55 1979; Fredman and Jahan-Parwar, 1980) and modulatory neurons (Hall and Lloyd, 1990;
56 McPherson and Blankenship, 1992) underlying locomotion. Moreover, its fictive locomo-
57 tion is sustained for minutes, ideal for robustly characterising population dynamics. Using
58 this model system, here we find its low-dimensional, underlying dynamical system, test if
59 the low-dimensional signal encodes movement variables, and determine the contribution
60 of single neurons to the low-dimensional dynamics.

61 We show that evoking fictive locomotion caused heterogenous population spiking ac-
62 tivity, but under which always lay a low-dimensional, slowly decaying periodic orbit. This
63 periodic trajectory met the convergence and perturbation criteria for an attractor. Cru-
64 cially, we identify the attractor as a stable, decaying spiral in every preparation. We
65 decoded motoneuron activity directly from the low-dimensional orbit, showing that it di-
66 rectly encodes the relevant variables for movement. Yet we found that individual neurons
67 varied their participation in the attractor between bouts of locomotion. Consequently, only
68 the low-dimensional signal and not the high-dimensional population activity was consis-
69 tent within and between nervous systems. These findings strongly constrain the possible
70 implementations of the pattern generator for crawling in *Aplysia*; and by quantifying the
71 attractor they make possible future testing of how short- and long-term learning change
72 properties of that attractor. Collectively, these results provide experimental support for
73 the long-standing idea that neural population activity is a high-dimensional emergent

74 property of a simpler, low-dimensional dynamical system.

75 Results

76 We sequentially evoked three bouts of fictive locomotion in each of 10 isolated central
 77 nervous system preparations (Figure 1B). Each bout of locomotion was evoked by short
 78 stimulation of the tail nerve P9, mimicking a sensory stimulus to the tail that elicits the
 79 escape locomotion response (Hening et al., 1979); in intact animals, a strong tail stimulus
 80 typically elicits a two-part escape behavior consisting of several cycles of a vigorous arch-
 81 ing gallop, followed by several minutes of a more sedate rhythmic crawl (Jahan-Parwar
 82 and Fredman, 1979; Flinn et al., 2001). We imaged the dorsal pedal ganglion 30 s before
 83 through to 90 s after the evoking stimulus, aiming to capture the population dynamics
 84 initiating and driving the initial gallop before the transition to the crawl. Recorded popu-
 85 lations from the pedal ganglion comprised 120-180 neurons each, representing $\approx 10\%$ of the
 86 network in each recording. The population recordings captured rich, varied single neuron
 87 dynamics within the ganglion's network following the stimulus (Figure 1C). A dominant,
 88 slow (≤ 0.1 Hz) oscillation in neural firing (Figure 1D) is consistent with the periodic
 89 activity necessary to generate rhythmic locomotion. But the variety of single neuron dy-
 90 namics (Bruno et al., 2015) (Figure 1C) and the slowly decaying population firing rate
 91 (Figure 1F) post-stimulus hint at a more complex underlying dynamical system driving
 92 locomotion than a simple, consistent oscillator.

93 Seeking the simplest dynamical system to account for these data, we first show here
 94 that the joint activity of the population meets the necessary conditions for a periodic
 95 attractor (Figure 1F). We identified these as: (1) applying a driving force causes the
 96 system's activity to fall onto a stable, periodic orbit; (2) repeatedly driving the system
 97 causes convergence of its activity to the same orbit; and (3) the system should return to
 98 the periodic orbit after the end of transient perturbation. Figure 1 - figure supplement 1
 99 demonstrates these conditions in a dynamical model of a neural periodic attractor.

100 Joint population activity forms a low-dimensional periodic orbit

101 We first established that under the heterogenous population activity evoked by the tail-
 102 nerve stimulation there was a low dimensional periodic trajectory, consistent with there
 103 being a periodic attractor in the pedal ganglion. Projections of a population's joint activity
 104 into three dimensions typically showed that stimulation caused a strong deviation from
 105 the spontaneous state, which then settled into repeated loops (Figure 2A). Capturing a
 106 significant proportion (80%) of the population variance generally required 4-8 embedding
 107 dimensions (Figure 2B), representing a dimension reduction by more than a factor of 10
 108 compared to the number of neurons. Thus, throughout our analysis, we projected each
 109 evoked program into the number of embedding dimensions needed to capture at least 80%
 110 of the variance in population activity (4-8 dimensions: inset of Figure 2B). However, we
 111 cannot directly visualise this space; therefore we could not tell by visual inspection if the
 112 low-dimensional trajectory repeatedly returned to the same position, and so was truly
 113 periodic.

114 To determine whether population activity in higher dimensions reached a stable peri-
 115 odic orbit, we made use of the idea of recurrence (Lathrop and Kostelich, 1989; Marwan
 116 et al., 2007). For each time-point in the low-dimensional trajectory of the population's
 117 activity, we check if the trajectory passes close to the same point in the future (Figure
 118 2C). If so, then the current time-point *recurs*, indicating that the joint activity of the

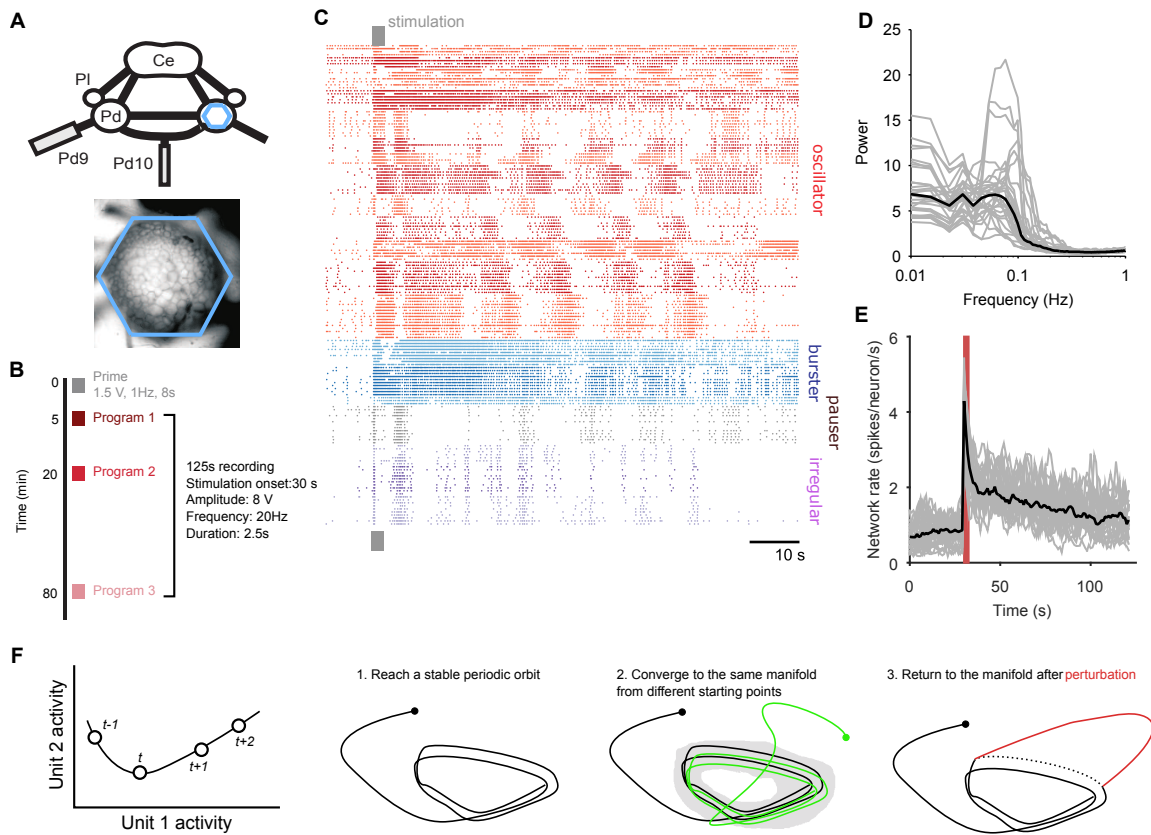


Figure 1: Population dynamics during fictive locomotion. **A** Voltage-sensitive dye recording of the pedal ganglion (Pd) network in an isolated central nervous system preparation (top) using a photodiode array (blue hexagon). The array covered the dorsal surface of the ganglion (bottom). Ce: cerebral ganglion; Pl: pleural ganglion; Pd9/10: pedal nerve 9/10. **B** Stimulus protocol. Three escape locomotion bouts were evoked in each preparation by stimulation of tail nerve Pd9. Parameters are given for the stimulus pulse train. **C** Example population recording. Raster plot of 160 neurons before and after Pd9 nerve stimulation. Neurons are grouped into ensembles of similarly-patterned firing, and ordered by ensemble type (colors) - see Methods. **D** Power spectra of each population's spike-trains, post-stimulation (grey: mean spectrum of each bout; black: mean over all bouts). **E** Network firing rate over time (grey: every bout; black: mean; red bar: stimulation duration. Bins: 1s). **F** Terminology and schematic illustration of the necessary conditions for identifying a periodic attractor (or "cyclical" attractor). Left: to characterise the dynamics of a N -dimensional system, we use the joint activity of its N units at each time-point t - illustrated here for $N = 2$ units. The set of joint activity points in time order defines the system's trajectory (black line). Right: the three conditions for identifying a periodic attractor. In each panel, the line indicates the trajectory of the joint activity of all units in the dynamical system, starting from the solid dot. The manifold of a dynamical system is the space containing all possible trajectories of the unperturbed system - for periodic systems, we consider the manifold to contain all periodic parts of the trajectories (grey shading). In (3), the dashed line indicates where the normal trajectory of the system would have been if not for the perturbation (red line). See Figure 1 - figure supplement 1 for a dynamical model illustrating these conditions.

119 population revisits the same state at least once. The time between the current time-point
120 and when it recurs gives us the period of recurrence. A strongly periodic system would
121 thus be revealed by its population's trajectory having many recurrent points with similar
122 recurrence periods; random or chaotic dynamics, by contrast, would not show a single
123 clustered recurrence period.

124 Plotting recurrent time-points showed that the evoked low-dimensional population
125 activity typically recurred with a regular period (example in Figure 2D). We found strongly
126 periodic recurrence on the scale of 10-15 s in many but not all of the 30 evoked population
127 responses (Figure 2E,F). This reflected the range of stimulation responses from strongly
128 periodic activity across the population to noisy, stuttering, irregular activity (Figure 2 -
129 figure supplement 1). Nonetheless, despite this heterogeneity across stimulus responses,
130 the activity of almost all populations was dominated by a single periodic orbit (Figure
131 2E), robust to the choice of threshold for defining recurrence (Figure 2 - figure supplement
132 2).

133 **Joint population activity meets the conditions for a periodic attractor**

134 The trajectory of a periodic dynamical system remains within a circumscribed region of
135 space – the manifold – that is defined by all the possible states of that system. (We
136 schematically illustrate a manifold by the grey shading in Figure 1F (condition 2), and
137 demonstrate the manifold of our model periodic attractor network in panel C of Figure
138 1 - figure supplement 1). If the population responses of the pedal ganglion are from an
139 underlying periodic attractor, then the population's joint activity should rapidly reach and
140 stay on its manifold when evoked; reach the same manifold every time it is evoked; and
141 return to the manifold when perturbed (these three conditions are schematically illustrated
142 in Figure 1F; see Figure 1 - figure supplement 1 for the corresponding examples from the
143 dynamical model).

144 We found that almost all evoked population responses quickly reached a state of high
145 recurrence, within one oscillation period (Figure 3A), and were thereafter dominated by
146 recurrence, indicating they quickly reached and stayed on the manifold.

147 But does each population response from the same preparation reach the same man-
148 ifold? The key problem in analysing any putative attractor from experimental data is
149 identifying when the experimentally-measured dynamics are or are not on the attractor's
150 manifold, whether due to perturbations of the system or noise in the measurements. More-
151 over, we cannot directly compare time-series between evoked responses because, as just
152 demonstrated, each response may reach the manifold at different times (see also panel C
153 in Figure 1 - figure supplement 1). Thus the set of recurrent time-points allowed us to
154 identify when the joint population activity was most likely on the attractor's manifold,
155 and then to make comparisons between population responses.

156 To determine if sequentially-evoked responses from the same preparation reached the
157 same manifold, we projected all 3 population responses into the same set of embedding
158 dimensions, using only the recurrent points (Figure 3B; Figure 3 - figure supplement 1
159 shows these results are robust to other projections). Falling on the same manifold would
160 mean that every recurrent point in one population response's trajectory would also appear
161 in both the others' trajectories, if noiseless. Consequently, the maximum distance between
162 any randomly chosen recurrent point in population response A and the closest recurrent
163 point in population response B should be small. We defined small here as being shorter
164 than the expected distance between a recurrent point in A and the closest point on a
165 random projection of the activity in the same embedding dimensions. Despite the inherent
166 noise and limited duration of the recordings, this is exactly what we found: pairs of evoked

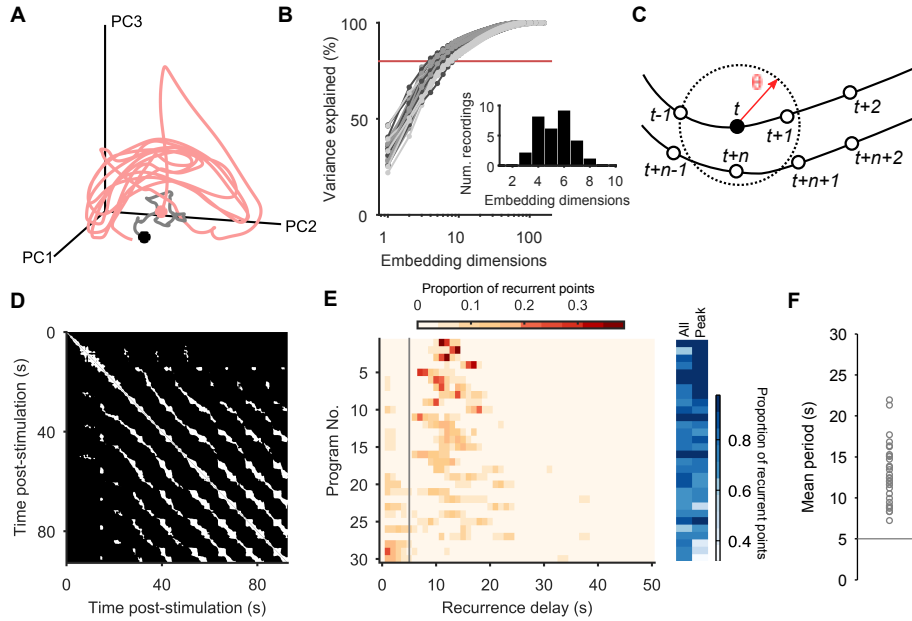


Figure 2: Population dynamics form a low-dimensional periodic orbit. **A** Projection of one evoked population response into 3 embedding dimensions, given by its first 3 principal components (PCs). Dots: start of recording (black) and stimulation (pink); spontaneous activity is shown in grey. Smoothed with 2 s boxcar window. **B** Proportion of population variance explained by each additional embedding dimension, for every evoked population response ($n = 30$; light-to-dark grey scale indicates stimulations 1 to 3 of a preparation). We chose a threshold of 80% variance (red line) to approximately capture the main dimensions: beyond this, small gains in explained variance required exponentially-increasing numbers of dimensions. Inset: Histogram of the number of PCs needed to explain 80% variance in every recorded population response. **C** Quantifying population dynamics using recurrence. Population activity at some time t is a point in N -dimensional space (black circle), following some trajectory (line and open circles); that point *recurs* if activity at a later time $t + n$ passes within some small threshold distance θ . The time n is the recurrence time of point t . **D** Recurrence plot of the population response in panel A. White squares are recurrence times, where the low-dimensional dynamics at two different times passed within distance θ . We defined θ as a percentile of all distances between two points; here we use 10%. Stimulation caused the population’s activity to recur with a regular period. Projection used 4 PCs. **E** Histograms of all recurrence times in each population response (threshold: 10%), ordered top-to-bottom by height of normalised peak value. Vertical line indicates the minimum time we used for defining the largest peak as the dominant period for that population response. Right: density of time-points that were recurrent, and density of recurrence points with times in the dominant period. **F** Periodic orbit of each evoked population response, estimated as the mean recurrence time from the dominant period.

167 population responses from the same preparation fell close to each other throughout (Figure
168 3C), well in excess of the expected agreement between random projections of the data onto
169 the same embedding dimensions.

170 We also checked that this convergence to the same manifold came from different initial
171 conditions. The initiating stimulation is a rough kick to the system – indeed a fictive
172 locomotion bout can be initiated with a variety of stimulation parameters (Bruno et al.,
173 2015) – applied to ongoing spontaneous activity. Together, the stimulation and the state
174 of spontaneous activity when it is applied should give different initial conditions from
175 which the attractor manifold is reached. We found that the stimulation caused population
176 responses within the same preparation to diverge far more than in either the spontaneous
177 activity or after coalescing to the manifold (Figure 3D). Thus, a wide range of initial
178 driven dynamics in the pedal ganglion population converged onto the same manifold.

179 Previous studies have used the consistency of pairwise correlations between neurons
180 across conditions as indirect evidence for the convergence of population activity to an
181 underlying attractor (Yoon et al., 2013; Peyrache et al., 2015). The intuition here is
182 that neurons whose activity contributes to the same portion of the manifold will have
183 simultaneous spiking, and so their activity will correlate across repeated visits of the
184 population’s activity to the same part of the manifold. To check this, we computed the
185 pairwise similarity between all neurons within an evoked population response (Figure 3E),
186 then correlated these similarity matrices between responses from the same preparation.
187 We found that pair-wise similarity is indeed well-preserved across population responses in
188 the same preparation (Figure 3F). This also shows that the apparent convergence to the
189 same manifold is not an artefact of our choice of low-dimensional projection.

190 In many population responses, we noticed spontaneous perturbations of the low-
191 dimensional dynamics away from the trajectory (examples in Figure 3 - figure supplement
192 2), indicated by sudden falls in the density of recurrent points (Figure 3G). That is, pertur-
193 bations could be detected by runs of contiguous points on the population trajectory that
194 were not recurrent. As each spontaneous perturbation was a cessation of recurrence in a
195 trajectory accounting for 80% of the co-variation between neurons, each was a population-
196 wide alteration of neuron activity (see example rasters in Figure 3 - figure supplement 2).
197 In most cases (90%), the population dynamics returned to a recurrent state after the
198 spontaneous perturbation (Figure 3H; Figure 3 - figure supplement 2, panel B), consistent
199 with the perturbation being caused by a transient effect on the population. The two pertur-
200 bations that did not return to a recurrent state were consistent with the end of the evoked
201 fictive locomotion and a return to spontaneous activity (Figure 3 - figure supplement 2,
202 panel A). Of those that returned, all but three clearly returned to the same manifold
203 (Figure 3I); for those three, the spontaneous perturbation appeared sufficient to move the
204 population dynamics into a different periodic attractor (Figure 3 - figure supplement 2,
205 panel C). Potentially, these are the known transitions from the escape gallop to normal
206 crawling (Flinn et al., 2001). The low dimensional dynamics of the pedal ganglion thus
207 meet the stability, manifold convergence, and perturbation criteria of a periodic attractor
208 network.

209 **Heterogenous population activity arises from a common attractor**

210 While these results show the existence of a periodic orbit on an attractor in the evoked
211 population responses, they cannot address whether these arise from the same putative
212 attractor within and, crucially, between animals. To determine if there is a common
213 underlying attractor despite the heterogeneity in spiking patterns across the population
214 responses (Figure 2 - figure supplement 1), we introduced a statistical approach to quan-

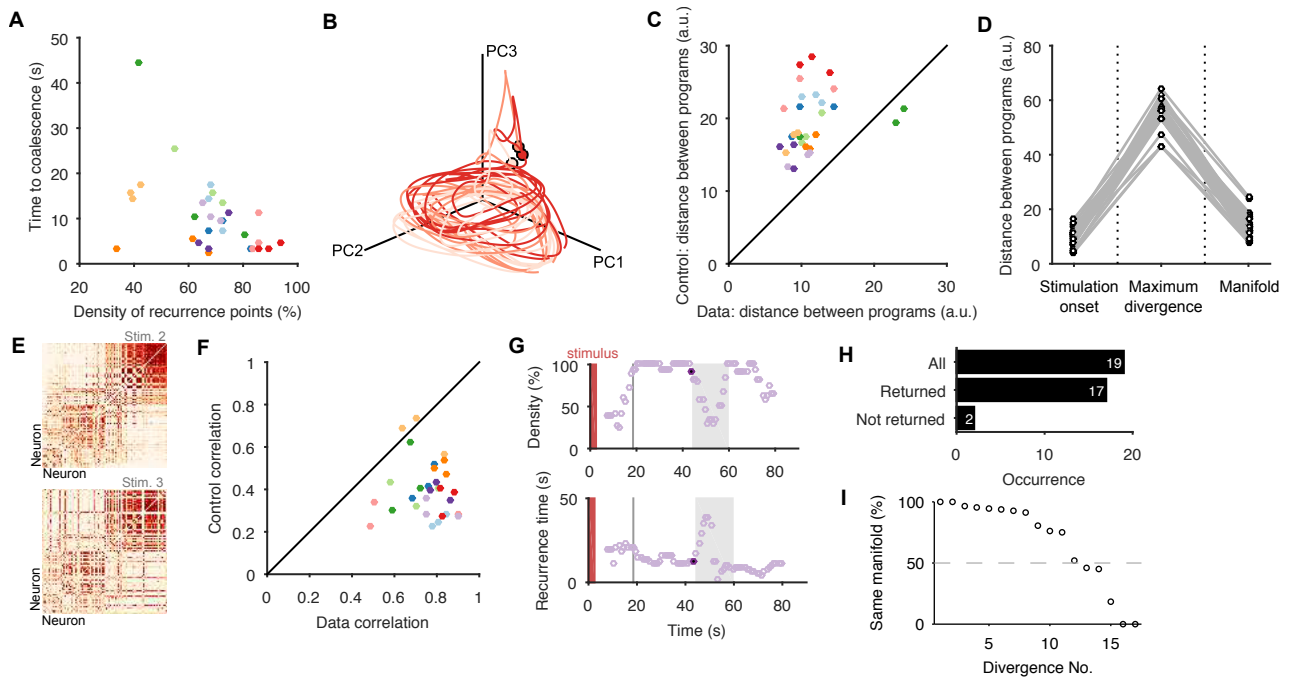


Figure 3: Low dimensional population dynamics meet the conditions for a periodic attractor. **A** Distribution of the time the population dynamics took to coalesce onto the attractor from the stimulation onset, and the subsequent stability of the attractor (measured by the proportion of recurrent points). Colours indicate evoked responses from the same preparation. The coalescence time is the mid-point of the first 5 second sliding window in which at least 90% of the points on the population trajectory recurred in the future. **B** Projection of three sequential population responses from one preparation onto the same embedding dimensions. Dots are time of stimulus offset. **C** Sequential population responses fall onto the same manifold. Dots indicate distances between pairs of population responses in the same preparation; color indicates preparation. Control distances are from random projections of each population response onto the same embedding dimensions - using the same time-series, but shuffling the assignment of time series to neurons. This shows how much of the manifold agreement is due to the choice of embedding dimensions alone. The two pairs below the diagonal are for response pairs (1,2) and (1,3) in preparation 4; this correctly identified the unique presence of apparent chaos in response 1 (see Figure 3 - figure supplement 1). **D** Distances between pairs of population responses from the same preparation in three states: the end of spontaneous activity (at stimulus onset); between stimulation onset and coalescence (the maximum distance between the pair); and after both had coalesced (both reaching the putative attractor manifold; data from panel C). **E** Example neuron activity similarity matrices for consecutively evoked population responses. Neurons are ordered according to their total similarity in stimulation 2. **F** Correlation between pairs of neuron similarity matrices (Data) compared to the expected correlation between pairs of matrices with the same total similarity per neuron (Control). Values below the diagonal indicate conserved pairwise correlations between pairs of population responses within the same preparation. The two pairs on the diagonal are response pairs (1,3) and (2,3) in preparation 7; this correctly identified the unique presence of a random walk in response 3 (see Figure 3 - figure supplement 1). **G** Spontaneous divergence from the trajectory. For one population response, here we plot the density of recurrence points (top) and the mean recurrence delay in 5s sliding windows. Coalescence time: grey line. The sustained “divergent” period of low recurrence (grey shading) shows the population spontaneously diverged from its ongoing trajectory, before returning. Black dot: pre-divergence window (panel I). **H** Breakdown of spontaneous perturbations across all population responses. Returned: population activity became stably recurrent after the perturbation. **I** Returning to the same manifold. For each of the 17 “Returned” perturbations in panel H, the proportion of the recurrent points in the pre-divergence window that recurred after the divergent period, indicating a return to the same manifold or to a different manifold.

215 tifying the low-dimensional trajectory. We first fitted a linear model of the local dynamics
 216 around each time point in the low-dimensional projection (see Methods). For each N -
 217 dimensional point $P(t)$ in this projection, we fitted the N -dimensional model $\dot{P}^* = \mathbf{A}P^*$
 218 to the trajectory forwards and backwards in time from point $P(t)$. In this model, the
 219 change in the trajectory over time \dot{P}^* in the neighbourhood of point $P(t)$ is determined
 220 by the values of the $N \times N$ matrix \mathbf{A} . The maximum eigenvalue of A thus tells us
 221 whether the trajectory around point $P(t)$ is predominantly expanding or contracting in
 222 the N -dimensional projection, and whether or not it is rotating (Strogatz, 1994).

223 By fitting the linear model to each point on the trajectory we obtained time-series of
 224 the maximum eigenvalues, describing the local dynamics at each point along the trajec-
 225 tory. The time-series of eigenvalues typically showed long periods of similar magnitude
 226 eigenvalues, corresponding to the recurrent points (Figure 4A). Consequently, by then
 227 averaging over the eigenvalues obtained only for recurrent points, we could potentially
 228 capture the dynamics of the underlying attractor. Doing so, we found that the evoked
 229 population responses had highly clustered maximum eigenvalues (Figure 4B,C), and thus
 230 highly similar underlying dynamics despite the apparent heterogeneity of spike-train pat-
 231 terns between them. The dominance of negative complex eigenvalues implies the pedal
 232 ganglion network implements a contracting periodic orbit - it is a stable spiral attractor
 233 (Figure 4D).

234 In most population responses, the low-dimensional trajectory had negative, complex
 235 eigenvalues in all embedding dimensions, meaning that the spiral attractor completely
 236 characterised the population dynamics (Figure 4 - figure supplement 1). Intriguingly, a
 237 few population responses had a positive real eigenvalue in one low-variance dimension
 238 (Figure 4 - figure supplement 1), implying a simultaneous minor expansion of the popu-
 239 lation trajectory. This corresponded to the appearance of a small sub-set of neurons with
 240 increasing firing rates (Figure 4E).

241 The identification of a stable spiral makes a clear prediction for what should and should
 242 not change over time in the dynamics of the population. The negative complex eigenvalues
 243 mean that the magnitude of the orbit decays over time, corresponding to the decreasing
 244 population spike rate in most evoked responses (Figure 1E). However, a stable spiral
 245 indicates only a decrease in magnitude; it does not mean the orbital period is also slowing.
 246 Consequently, the presence of a stable spiral attractor predicts that the magnitude and
 247 period of the orbit are dissociable properties in the pedal ganglion network.

248 We checked this prediction using the linear model. The linear model estimated a mean
 249 orbital period of around 10 s (Figure 4C), consistent with the directly-derived estimate
 250 from the recurrent points (Figure 2F). This indicated the linear model was correctly cap-
 251 turing the local dynamics of each program. But our linear model also gave us a time-series
 252 of estimates of the local orbital period (Figure 5A), which we could use to check whether
 253 the orbital period was changing during each evoked response. We found that the popula-
 254 tion responses included all possible changes in periodic orbit: slowing, speeding up, and
 255 not changing (Figure 5B). As predicted there was no relationship between the contraction
 256 of the periodic orbit and its change in period (Figure 5C).

257 **The locomotion motor program can be decoded from the low-dimensional** 258 **orbit.**

259 Collectively, these periodic, decaying dynamics are ethologically consistent with locomotion
 260 that comprises a repeated sequence of movements that decays in intensity over time
 261 (Jahan-Parwar and Fredman, 1979; Flinn et al., 2001; Marinesco et al., 2004). If this pu-
 262 tative low-dimensional periodic attractor is the “motor program” for locomotion, then we

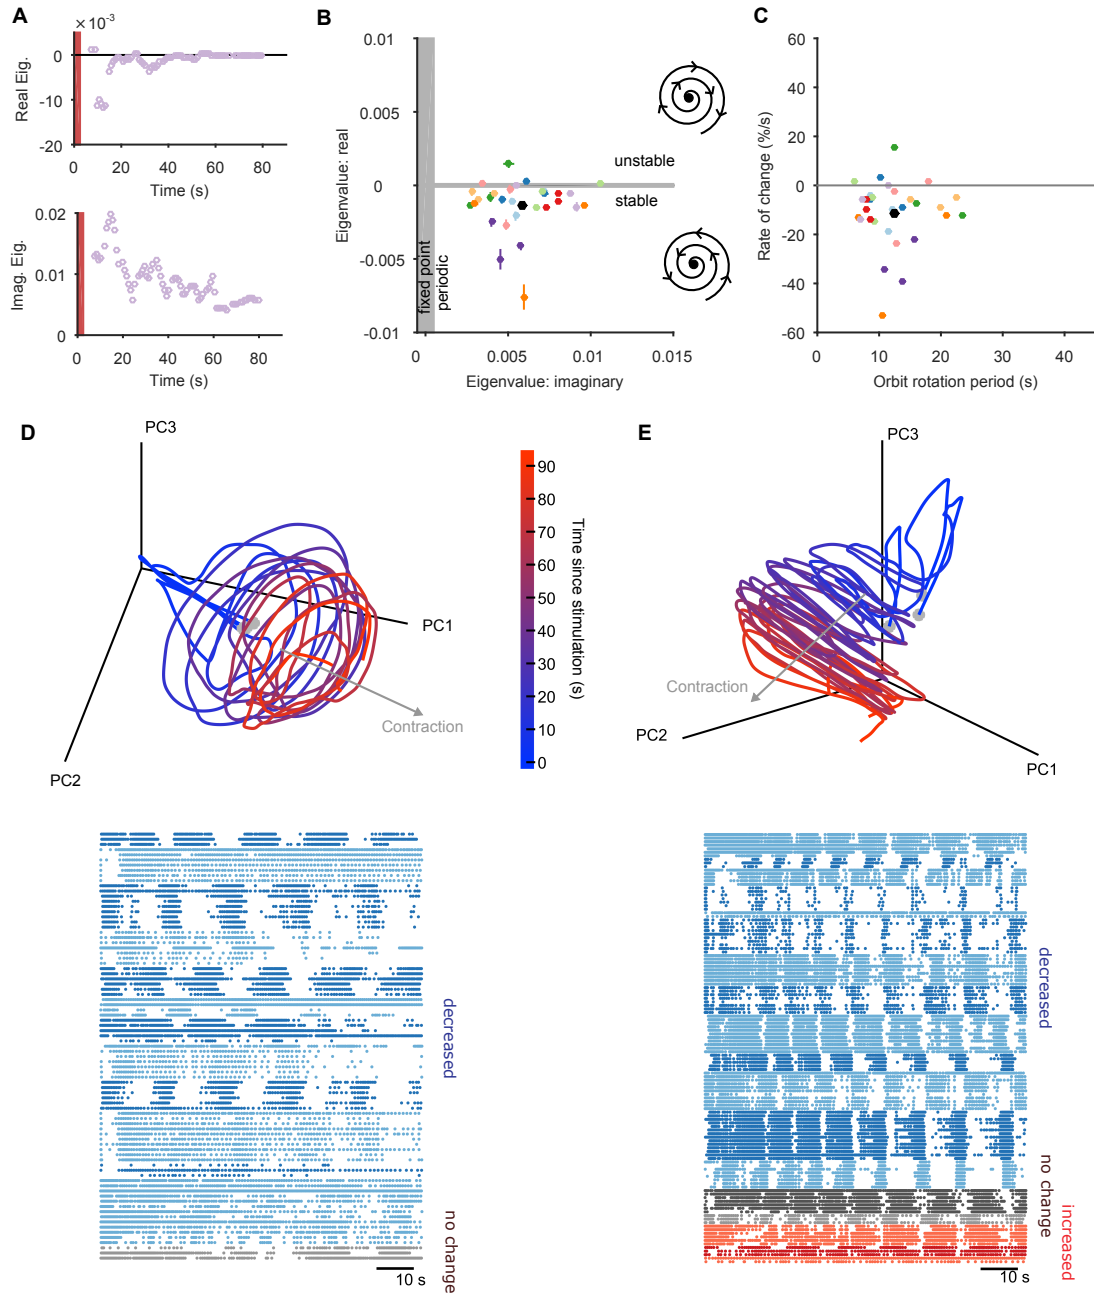


Figure 4: The pedal ganglion contains a spiral attractor.

A Example time-series from one population response of the real (top) and imaginary (bottom) component of the maximum eigenvalue for the local linear model. Points are averages over a 5s sliding window. Red bar indicates stimulus duration. **B** Dominant dynamics for each evoked population response. Dots and lines give means ± 2 s.e.m. of the real and imaginary components of the maximum eigenvalues for the local linear model. Colours indicate responses from the same preparation. Black dot gives the mean over all population responses. Grey shaded regions approximately divide the plane of eigenvalues components into regions of qualitatively different dynamics: fixed point attractor; stable spiral (bottom-right schematic); unstable spiral (top-right schematic). **C** As panel B, converted to estimates of orbital period and rate of contraction. (Note that higher imaginary eigenvalues equates to faster orbital periods, so the ordering of population responses is flipped on the x-axis compared to panel B).

D A preparation with a visible spiral attractor in a three-dimensional projection. Each line is one of the three evoked population responses, colour-coded by time-elapsed since stimulation (grey circle). The periodicity of the evoked response is the number of loops in the elapsed time; loop magnitude corresponds to the magnitude of population activity. The approximate dominant axis of the spiral's contraction is indicated. Bottom: corresponding raster plot of one evoked response. Neurons are clustered into ensembles, and colour-coded by the change in ensemble firing rate to show the dominance of decreasing rates corresponding to the contracting loop in the projection.

E As panel D, but for a preparation with simultaneously visible dominant spiral and minor expansion of the low-dimensional trajectory. The expansion corresponds to the small population of neurons with increasing rates.




Figure 5: The spiral attractor dissociates changes in oscillation period and firing rate. **A** Example of a change in the local estimate of the periodic orbit over a population response; here, slowing over time ($n = 57$ points are each averages over a 5 s sliding window; ρ is weighted Spearman’s rank correlation - see Methods; P from a permutation test). Changes in the periodic orbit were assessed only after coalescence to the manifold (grey line). **B** Histogram of correlations between time elapsed and local estimate of the periodic orbit for each population response (positive: slowing; negative: speeding up). Red bars correspond to population responses with $P < 0.01$ (permutation test). Number of local estimates ranged between 31 and 72 per population response. **C** Relationship between the change in periodic orbit over time and the rate of contraction for each population response (Pearson’s R ; $n = 30$ responses).

263 should be able to decode the locomotion muscle commands from its trajectory. In 3 of the
 264 10 preparations we were able to simultaneously record activity from the P10 nerve that
 265 projects to the neck muscles (Xin et al., 1996) for all three evoked population responses.
 266 The spiking of axons in this nerve should correspond to the specific neck contraction
 267 portion of the cyclical escape locomotion. We thus sought to decode the spiking of P10
 268 directly from the low-dimensional population trajectory (Figure 6A).

269 We first confirmed that each recorded neural population did not appear to contain
 270 any motoneurons with axons in P10, which could make the decoding potentially trivial
 271 (Figure 6 - figure supplement 1). To then decode P10 activity, we used a statistical model
 272 that predicts the firing rate of nerve P10 at each time point, by weighting and summing the
 273 recent history (up to 100 ms) of the trajectory in the low dimensional space, and using a
 274 non-linearity to convert this weighted sum into a firing rate. We controlled for over-fitting
 275 using cross-validation forecasting: we fit the model to a 40 s window of trajectory data,
 276 and predicted the next 10 s of P10 activity (Figure 6B). By sliding the window over the
 277 data, we could assess the quality of the forecast over the entire recording (Figure 6C).

278 The model could accurately fit and forecast P10 activity from the low-dimensional
 279 trajectory in all 9 population responses (Figure 6D). Emphasising the quality of the model,
 280 in Figure 6D we plot example forecasts of the entire P10 recording based on fitting only to
 281 the first 40s window, each example taken from the extremes we obtained for the fit-quality
 282 metrics. Notably, in one recording the population response shutdown half-way through;
 283 yet despite the model being fit only to the 40s window containing strong oscillations, it
 284 correctly forecasts the collapse of P10 activity, and its slight rise in firing rate thereafter.
 285 Thus, the low dimensional trajectory of the periodic attractor appears to directly encode
 286 muscle commands for movement.

287 To confirm this, we asked whether the encoding – as represented by the P10 activity
 288 – was truly low-dimensional. The successful decoding of future P10 activity was achieved
 289 despite needing only 3-5 embedding dimensions to account for 80% variance in the popu-
 290 lation activity for these nine recordings (Figure 6 - figure supplement 2). Increasing the
 291 number of embedding dimensions to account for 90% variance, at least doubling the num-
 292 ber of embedding dimensions, did not improve the forecasts of P10 activity (Figure 6 -
 293 figure supplement 2). These results suggest that the low dimensional population trajectory

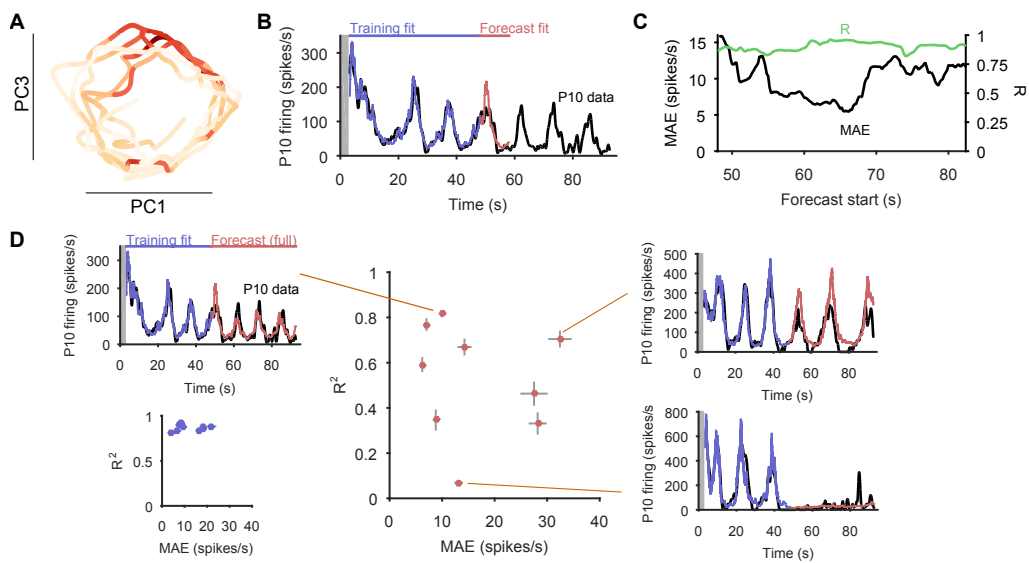


Figure 6: Motor output can be decoded directly from the low-dimensional trajectory of population activity. **A** An example two-dimensional projection of one population’s response trajectory, color-coded by simultaneous P10 firing rate. In this example pair of dimensions, we can see nerve P10 firing is phase-aligned to the periodic trajectory of population activity. **B** Example fit and forecast by the statistical decoding model for P10 firing rate. Grey bar indicates stimulation time. **C** For the same example P10 data, the quality of the forecast in the 10 s after each fitted 40 s sliding window. Match between the model forecast and P10 data was quantified by the fits to both the change (R : correlation coefficient) and the scale (MAE: median absolute error) of activity over the forecast window. **D** Summary of model forecasts for all 9 population responses with P10 activity (main panel). Dots and lines show means \pm 2 s.e.m. over all forecast windows ($N = 173$). Three examples from the extremes of the forecast quality are shown, each using the fitted model to the first 40 s window to forecast the entire remaining P10 time-series. The bottom right example is from a recording in which the population response apparently shutdown half-way through. Inset, lower left: summary of model fits in the training windows; conventions as per main panel.

294 is sufficient to encode the locomotion muscle commands.

295 **Variable neuron participation in stable motor programs**

296 If the low-dimensional trajectory described by the joint activity of the population just
297 is the motor program for locomotion, then how crucial to this program are the firing of
298 individual neurons (Katz et al., 2004; Carmena et al., 2005; Hill et al., 2012; Huber et al.,
299 2012; Carroll and Ramirez, 2013; Hill et al., 2015)? Having quantified the motor program
300 as the low-dimensional activity trajectory, we could uniquely ask how much each neuron
301 participated in each evoked program. We quantified each neuron’s *participation* as the
302 absolute sum of its weights on the principal axes (eigenvectors): large total weights indicate
303 a dominant contribution to the low-dimensional trajectory, and small weights indicate little
304 contribution. So quantified, participation is a contextual measure, giving the contribution
305 to the population trajectory of both a neuron’s firing rate and its synchrony with other
306 neurons, relative to the rate and synchrony of all other neurons in the population (Figure
307 7 - figure supplement 1).

308 Every population response had a long-tailed distribution of participation (Figure 7A),
309 indicating that a minority of neurons dominated the dynamics of any given response.
310 Nonetheless, these neurons were not fixed: many with high participation in one popula-
311 tion response showed low participation in another (Figure 7B,C). To rule out noise effects
312 on the variability of participation (for example, due to the finite duration of recording), we
313 fitted a noise model to the change in participation separately for each preparation (Figure
314 7D,E). Every preparation’s pedal ganglion contained neurons whose change in partici-
315 pation between responses well-exceeded that predicted by the noise model (Figure 7F).
316 Consequently, the contribution of single neurons was consistently and strongly variable
317 between population responses in the same preparation.

318 We also tested for the possibility that hidden within the variation between programs
319 is a small core of neurons that are strongly participating, yet invariant across programs.
320 Such a core of phasically active neurons may, for example, form the basis of a classical
321 central pattern generator. However, in our observed portion of the ganglion we found
322 no evidence for a core of strongly participating, invariant, and phasically active neurons
323 across the preparations (Figure 7 - figure supplement 2).

324 These data show that a neuron’s role within the locomotion motor program is not
325 fixed, but leave open the question of whether single neuron variability causes variation
326 in the program itself. In our analysis, variation between sequentially-evoked population
327 responses is quantified by the distance between their low-dimensional projections (as in
328 Figure 3C). We found that the distance between a pair of population responses did not
329 correlate with either the total change in neuron participation between the two responses
330 (Figure 7G) or the distance between their participation distributions (Figure 7H). The
331 execution of the motor program is thus robust to the participation of individual neurons.

332 **Participation maps identify potential locations of the pattern generator 333 network**

334 To get some insight into the physical substrate of the attractor, we plotted maps of the
335 participation of each neuron in each preparation. We found that neurons with strong
336 participation across the three evoked population responses were robustly located in the
337 caudo-lateral quadrant of the ganglion (Figure 8A,B). Maps of the right ganglion also
338 indicated strong participation in the rostro-medial quadrant; due to the low numbers
339 of maps for each side, it is unclear whether this is a true asymmetry of the ganglia or

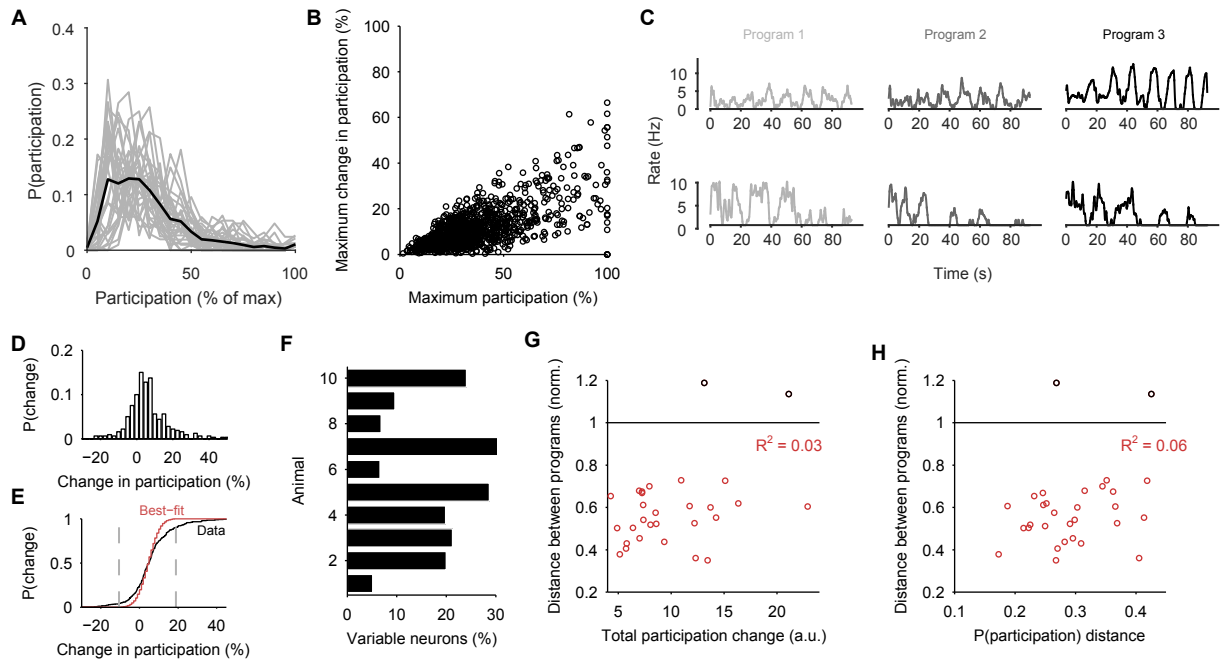


Figure 7: Single neuron participation varies within and between evoked locomotion bouts. **A** Distributions of single neuron participation per evoked population response. We plot the distribution of participation for all neurons in a population (grey line), expressed as a percentage of the maximum participation in that population’s response. Black line gives the mean over all 30 population responses. **B** Change in participation between evoked locomotion bouts. Each dot plots one neuron’s maximum participation over all 3 evoked population responses, against its maximum change in participation between consecutive responses ($n = 1131$ neurons). **C** Two example neurons with variable participation between responses, from two different preparations. **D** Distribution of the change in participation between responses for one preparation. **E** Detecting strongly variable neurons. Gaussian fit (red) to the distribution of change in participation (black) from panel D. Neurons beyond thresholds (grey lines) of mean $\pm 3SD$ of the fitted model were identified as strongly variable. **F** Proportion of identified strongly variable neurons per preparation. **G** Distance between pairs of population responses as a function of the total change in neuron participation between them. Each dot is a pair of responses from one preparation; the distance between them is given as a proportion of the mean distance between each response and a random projection (< 1 : closer than random projections), allowing comparison between preparations (Figure 3C). Black dots are excluded outliers, corresponding to the pairs containing response 1 in preparation 4 with apparent chaotic activity (Figure 3 - figure supplement 1). **H** Distance between pairs of population responses as a function of the distance between the distributions of participation (panel A). Conventions as for panel G.

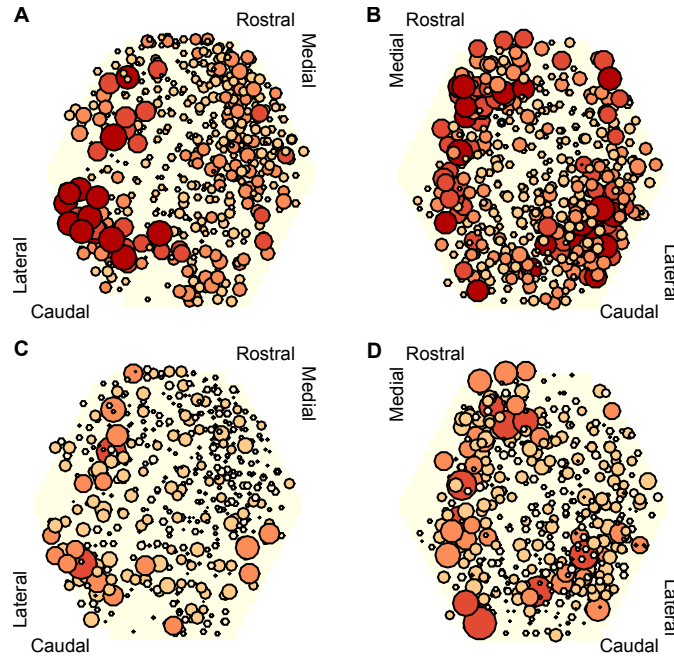


Figure 8: Mapping of participation in the attractor across the ganglion network. Here we plot neuron location with respect to the photodiode array (yellow hexagon). Each plot pools neurons from preparations of the left ($n = 4$ preparations) or right ($n = 4$) ganglia. **A,B** Maps of maximum participation across the three evoked population responses for left (A) and right (B) ganglion recordings. The area of each marker is proportional to the neuron’s maximum participation. Neurons are colour coded (light orange to dark red) by the quintile of their participation across all preparations. **C,D** As for panels (A,B), but plotting the range of participation across the three evoked population responses.

340 simply reflects sampling variation. Neurons with highly variable participation between
 341 population responses (Figure 8C,D) were similarly found in the caudo-lateral quadrants
 342 of both ganglia. Strongly participating neurons were thus confined to specific distributed
 343 regions of the pedal ganglion’s network.

344 These data are consistent with a network-level distribution of the attractor, with a
 345 particularly strong contribution from the caudo-lateral quadrant. Encouragingly, from
 346 a different data-set we previously described this region as containing neural ensembles
 347 that generated a cyclical packet of neural activity, which moved in phase with activity
 348 from the neck-projecting P10 nerve (Bruno et al., 2015). Consequently, both those data
 349 and our new data support our hypothesis that the pattern generator for locomotion is
 350 predominantly located in the caudo-lateral network.

351 Discussion

352 Locomotion networks provide a tractable basis for testing theories of neural dynamics
 353 (Lewis and Kristan, 1998; Briggman et al., 2005; Levi et al., 2005; Briggman and Kris-
 354 tan, 2006; Berg et al., 2007; Bruno et al., 2015; Petersen and Berg, 2016), as they couple
 355 complex dynamics with clearly defined outputs. We took advantage of this to comprehen-
 356 sively test the idea that high-dimensional population activity arises from an underlying
 357 low-dimensional dynamical system: to determine what dynamical system accounts for the
 358 population activity, whether its low-dimensional signal encodes movement, and how sin-
 359 gle neuron activity relates to that signal. We showed here that *Aplysia*’s pedal ganglion

360 contains a spiral attractor, that the low-dimensional signal it generates directly encodes
361 muscle commands, and yet individual neurons vary in their participation in the attractor.

362 **A consistent low-dimensional spiral attractor**

363 Testing the idea that high-dimensional population activity contains a low-dimensional sig-
364 nal has only been possible in the last decade or so, due to the necessary combination
365 of large-scale multi-neuron recording and dimension reduction approaches (Brown et al.,
366 2004; Briggman et al., 2006; Cunningham and Yu, 2014; Kobak et al., 2016). Landmark
367 studies have used this combination to project high-dimensional population activity into
368 a more tractable low-dimensional space. In this space, studies have shown how activity
369 trajectories are different between swimming and crawling (Briggman et al., 2005); dis-
370 tinguish olfactory (Mazor and Laurent, 2005), auditory (Bartho et al., 2009), and visual
371 (Mante et al., 2013) stimuli; and distinguish upcoming binary choices (Harvey et al., 2012).
372 Here we have gone a step further than previous studies by not only observing such low-
373 dimensional signals, but explicitly testing for the first time the type of dynamical system
374 that gives rise to the low-dimensional trajectories and its consistency between animals.

375 Across all 30 evoked population responses examined here, there was a remarkable het-
376 erogeneity of spike-train patterns, from visually evident widespread oscillations to noisy,
377 stuttering oscillations in a minority of neurons (Figure 2 - figure supplement 1). Yet
378 our analysis shows that underpinning this heterogeneity is the same dynamical system:
379 a low-dimensional, decaying, periodic orbit. We found a remarkably consistent periodic-
380 ity and rate of orbital decay across evoked responses within a preparation and between
381 preparations. The stability of these dynamics, and the convergence of population activity
382 to the same manifold, are all consistent with the expected behaviour of a true attractor.
383 Our data thus suggest that only the low-dimensional system and not the high-dimensional
384 population activity are consistent within and between nervous systems.

385 We advance the hypothesis that the properties of the spiral attractor fully determine
386 the parameters of the escape gallop: its frequency, physical distance per cycle, and du-
387 ration. In this hypothesis, the orbital period of the attractor determines the period of
388 the rhythmic gallop – the sequential activity of the neurons in each orbit thus driving
389 the sequential contraction of the muscles driving the escape gallop (Bruno et al., 2015).
390 Further, the amplitude of the orbital period, corresponding to the spike rate of the neural
391 population, could determine the strength of muscle contraction during the escape gallop,
392 allowing control of the physical distance covered by each arching movement. Finally, the
393 contraction rate of the attractor determines the duration of the escape: the faster the
394 contraction rate, the shorter the escape gallop’s duration. The variation of these attractor
395 properties between animals then determines the natural variability in the escape gallop.
396 It follows that changes to parameters of the escape gallop caused by neuromodulation
397 should correlate with changes to the orbital period and/or contraction rate of the at-
398 tractor. For example, the reported increase in gallop duration by systemic injection of
399 serotonin (Marinesco et al., 2004) should correlate with a decreased contraction rate of
400 the attractor. Future work could test this hypothesis by determining the effects of neu-
401 romodulators on the spiral attractor’s properties and correlating those with read-outs of
402 the escape gallop.

403 Treating a neural circuit as a realisation of a dynamical system takes the emphasis
404 away from the details of individual neurons - their neurotransmitters, their ion channel
405 repertoire - and places it instead on their collective action. This allows us to take a
406 Marr-ian perspective (Marr, 1982), which neatly separates the computational, algorithmic,
407 and implementation levels of movement control. The computational problem here is

408 of how to generate rhythmic locomotion for a finite duration; the algorithmic solution is
409 a decaying periodic attractor - a spiral; and the implementation of that attractor is the
410 particular configuration of neurons in the pedal ganglion - one of many possible imple-
411 mentations (Kleinfeld and Sompolinsky, 1988; Pasemann, 1995; Eliasmith, 2005; Rokni
412 and Sompolinsky, 2012). Indeed, a spiral attractor is potentially a general solution to the
413 problem of how to generate a finite rhythmic behaviour.

414 **Insights and challenges of variable neuron participation**

415 We saw the separation of these levels most clearly in the variable participation of the in-
416 dividual neurons between evoked bouts of fictive locomotion. The projection of the pedal
417 ganglion network’s joint activity into a low dimensional space captured the locomotion
418 motor program independently of any single neuron’s activity. Even the most strongly par-
419 ticipating neurons in a given population response could more than halve their participation
420 in other evoked responses. These results suggest that the pedal ganglion’s pattern gener-
421 ator is not driven by neurons that are endogenous oscillators, as they would be expected
422 to participate equally in every response. Rather, this variation supports the hypothesis
423 that the periodic activity is an emergent property of the network.

424 The adaptive function of having variably participating neurons is unknown. One pos-
425 sibility is that, by not relying on any core set of neurons to generate rhythmic activity, the
426 pedal ganglion’s ability to generate locomotion is robust to the loss of neurons. A related
427 possibility is that there is “sloppiness” (Panas et al., 2015) in the pedal ganglion network,
428 such that there are many possible configurations of neurons and their connections able to
429 realise the spiral attractor that drives locomotion (Marder et al., 2015). Such sloppiness
430 allows for a far more compact specification of the developmental program than needing to
431 genetically specify the type and wiring configuration of each specific neuron.

432 The wide variation of single neuron participation between evoked bouts of fictive lo-
433 comotion also raises new challenges for theories of neural network attractors (Marder and
434 Taylor, 2011). While a variety of models present solutions for self-sustaining periodic ac-
435 tivity in a network of neurons (Kleinfeld and Sompolinsky, 1988; Eliasmith, 2005; Rokni
436 and Sompolinsky, 2012), it is unclear if they can account for the variable participation of
437 single neurons. A further challenge is that while the variable participation of individual
438 neurons does not affect the underlying program, clearly it takes a collective change in
439 single neuron activity to transition between behaviours - as, for example, in the transition
440 from galloping to crawling in *Aplysia*. What controls these transitions, and how they
441 are realised by the population dynamics, is yet to be explored either experimentally or
442 theoretically.

443 **Possible implementations of rhythmic locomotion by the pedal ganglion 444 network**

445 Our results nonetheless argue against a number of hypotheses for the implementation of
446 rhythmic locomotion by the pedal ganglion. As noted above, such single neuron variability
447 between sequential locomotion bouts argues against the generation of rhythmic activity by
448 one or more independent neurons that are endogenous oscillators. Our results also argue
449 against the existence of many stable periodic states in this network (Pasemann, 1995).
450 Such meta-stability would manifest as changes in periodicity following perturbation. Our
451 results show that spontaneous divergences from the attractor overwhelmingly returned to
452 the same attractor.

453 How then might the pedal ganglion network implement a spiral attractor? Our data
454 were collected from an isolated central nervous system preparation, in which the modu-
455 latory influence of neurons outside the pedal ganglion cannot be discounted (Jing et al.,
456 2008). Nonetheless, as the pedal ganglion contains the central pattern generator for loco-
457 motion (Jahan-Parwar and Fredman, 1980), we can suggest how that generator is realised.
458 Our results here support the hypothesis that the periodic activity is an emergent property
459 of the ganglion’s network. We know the pedal ganglion contains a mix of interneurons
460 and motorneurons (Fredman and Jahan-Parwar, 1980), and that the motorneurons are
461 not synaptically coupled (Hening et al., 1979), suggesting they read-out (and potentially
462 feedback to) the dynamics of an interneuron network. An hypothesis consistent with our
463 results here is that the ganglion contains a recurrent network of excitatory interneurons,
464 centred on the caudo-lateral quadrant, which feed-forward to groups of motorneurons
465 (Bruno et al., 2015). This recurrent network embodies the attractor, in that stimulation
466 of the network causes a self-sustained packet of activity to sweep around it (Bruno et al.,
467 2015). We see this as the periodic trajectory of joint population activity (cf Figure 2A,
468 Figure 3B).

469 **Multiple periodic attractors and multi-functional circuits**

470 Our data further suggest that the pedal ganglion network supports at least two stable
471 states, the spontaneous activity and the stable-spiral attractor. Reaching the stable-spiral
472 attractor from the spontaneous activity required long-duration, high-voltage pedal nerve
473 stimulation (Figure 1; Bruno et al., 2015). In dynamical systems terms, this suggests that
474 the spontaneous state’s basin of attraction is large: most perturbations return to that
475 state, and it takes a large perturbation to move into a different basin of attraction.

476 Multiple co-existing periodic attractors in a single network is also a challenge for cur-
477 rent theories. While point attractor networks, such as Hopfield networks, can have vast
478 number of stable states defined by different arrangements of the equilibrium activity of
479 their neurons (Miller, 2016), a stable periodic attractor network typically has only two
480 stable states: silence and periodic activity. The co-existence of stable spontaneous and
481 periodic states in the same network suggests that something must reconfigure the network
482 to sustain periodic activity (Calin-Jageman et al., 2007); otherwise, irrespective of the
483 stimulation, the network would always return to the spontaneous state. One possibility
484 in the pedal ganglion is that serotonin alters the effective connections between neurons:
485 escape galloping is both dramatically extended by systemic injection of serotonin alongside
486 tail stimulation (Marinesco et al., 2004), and evoked by stimulating serotonergic command
487 neurons CC9/CC10 in the cerebral ganglion (Jing et al., 2008). Future experimental work
488 should thus test the stability of the spontaneous state, and test how manipulating sero-
489 tonin affects reaching and sustaining the stable-spiral attractor.

490 There are potentially more stable states within the pedal ganglion’s network. The
491 long-lasting crawl that follows the escape gallop is slower and omits the periodic arching
492 of the body (Flinn et al., 2001). We saw three perturbations of the attractor activity
493 that were suggestive of a transition to a different, slower periodic orbit (e.g. panel C in
494 Figure 3 - figure supplement 2), consistent with a transition from galloping to crawling.
495 Such crawling is also the animal’s normal mode of exploration (Leonard and Lukowiak,
496 1986), and so the “crawling” attractor must be reachable from the spontaneous state
497 too. *Aplysia*’s exploratory head-wave, moving its head side-to-side presumably to allow
498 its tentacles and other head sensory organs to sample the environment (Leonard and
499 Lukowiak, 1986), is also controlled by motorneurons in the pedal ganglion (Kuenzi and
500 Carew, 1994). Previous studies of the *Aplysia*’s abdominal ganglion (Wu et al., 1994), the

501 leech segmental ganglion (Briggman and Kristan, 2006), and the crustacean stomatogastric
502 ganglion (reviewed in Marder and Bucher, 2007) have described multi-functional networks
503 in which the same neurons are active in different motor behaviours. Our work here is
504 consistent with the hypothesis that such multi-function is due to the neurons participating
505 in different attractors realised by same network (Briggman and Kristan, 2008). Further
506 work is needed to map the pedal ganglion network’s dynamics to the full range of *Aplysia*
507 motor behaviour.

508 Outlook

509 Finding and quantifying the attractor required new analytical approaches. We introduce
510 here the idea of using recurrence analysis to solve two problems: how to identify periodic
511 activity in a high-dimensional space; and how to identify when the recorded system is and
512 is not on the manifold of the attractor. By extracting the times when the population ac-
513 tivity is on the manifold, we could then quantify and characterise the attractor, including
514 identifying transient perturbations, and estimating changes in orbital period. Crucially,
515 these manifold-times let us further introduce the idea of using linear models as a statistical
516 estimator, to identify the type of attractor, and compare the detected attractor’s param-
517 eters within and between preparations. Our analysis approach thus offers a road-map for
518 further understanding the dynamics of neural populations.

519 There is rich potential for understanding spontaneous, evoked or learning-induced
520 changes in the dynamics of populations for movement control. The dynamics of movement
521 control populations transition between states either spontaneously or driven by external
522 input (Briggman et al., 2005; Levi et al., 2005). Our recurrence approach allows both
523 the detection of transitions away from the current state (Figure 3) and the characteri-
524 sation of the attractor in the new state. For learning, taking an attractor-view allows
525 us to distinguish three distinct ways that short (Stopfer and Carew, 1988; Katz et al.,
526 1994; Hill et al., 2015) or long-term (Hawkins et al., 2006) plasticity could change the
527 underlying attractor: by changing the shape of the manifold; by changing the rate of
528 movement of the low-dimensional signal on the manifold; or by changing the read-out of
529 the manifold by downstream targets. Such insights may contribute to the grand challenge
530 of systems neuroscience, that of finding simplifying principles for neural systems in the
531 face of overwhelming complexity (Koch, 2012; Yuste, 2015).

532 Materials and methods

533 **Data and code availability** Bandpassed optical data, spike-sorted data, and available
534 P10 nerve recordings are hosted on CRCNS.org at: [DOI to follow].

535 All research code is available under a MIT License from (Humphries, 2017): <https://github.com/mdhumphries/AplysiaAttractorAnalysis>
536

537 **Imaging** Full details of the *Aplysia californica* preparation are given in Bruno et al.
538 (2015). Briefly, the cerebral, pleural and pedal ganglia were dissected out, pinned to the
539 bottom of a chamber, and maintained at 15 – 17°C. Imaging of neural activity used the
540 fast voltage sensitive absorbance dye RH-155 (Anaspec), and a 464-element photodiode
541 array (NeuroPDA-III, RedShirtImaging) sampled at 1600 Hz. Optical data from the 464
542 elements were bandpass filtered in Neuroplex (5 Hz high pass and 100 Hz low pass Butter-
543 worth filters), and then spike-sorted with independent component analysis in MATLAB
544 to yield single neuron action potential traces (the independent components), as detailed

545 in (Hill et al., 2010). Rhythmic locomotion motor programs were elicited using 8V 5ms
 546 monophasic pulses delivered at 20Hz for 2.5s via suction electrode to pedal nerve 9. A
 547 separate suction electrode was attached to pedal nerve 10 to continuously monitor the
 548 locomotion rhythm (Xin et al., 1996). Evoked activity could last for many minutes; our
 549 system allowed us to capture a maximum of ≈ 125 s, divided between 30 s of spontaneous
 550 activity and 95 s of evoked activity. The stimulation protocol (Figure 1B) used short (15
 551 mins) and long (60 mins) intervals between stimulations, as the original design also sought
 552 effects of sensitisation.

553 **Spike-train analysis** Power spectra were computed using multi-taper spectra routines
 554 from the Chronux toolbox (Bokil et al., 2010). We computed the power spectrum of each
 555 neuron’s spike-train post-stimulation, and plot means over all spectra within a recorded
 556 population, and the mean over all mean spectra. We computed the spike-density function
 557 $f(t)$ for each neuron by convolving each spike at time t_s with a Gaussian G : $f(t) =$
 558 $\sum_{t_0 < t_s < t_1} G(t_s) / \int_{t_0}^{t_1} G(t^*) dt^*$, evaluated over some finite window between t_0 and t_1 (see
 559 Szucs, 1998). We set the window to be $\pm 5\sigma$, and evaluated the convolution using a time-
 560 step of 10 ms. We defined the standard deviation σ of the Gaussian by the median inter-
 561 spike interval of the population: $\sigma = \{\text{median ISI in population}\} / \sqrt{12}$ (see Humphries,
 562 2011).

563 To visualise the entire population’s spiking activity (Figure 1C), we cluster neurons
 564 by the similarity of their firing patterns using our modular deconstruction toolbox (Bruno
 565 et al., 2015). Different dynamical types of ensembles were identified by the properties of
 566 their autocorrelograms: tonic, oscillator, burster, or pauser - see (Bruno et al., 2015) for
 567 details. We also assigned each neuron in the ensemble the same dynamical label, which
 568 we use in the analysis of Figure 7 - figure supplement 2. To demonstrate the firing rate
 569 change of each ensemble (Figure 4), we first counted the number of spikes emitted by that
 570 ensemble in 20s windows, advanced in 5s steps from the onset of stimulation. We then
 571 correlated (Pearson’s R) the time of each window against its spike count: ensembles were
 572 classified as decreasing rate if $R < -0.2$, and increasing if $R > 0.2$.

573 **Model network** We used a three-neuron network to demonstrate the dynamical prop-
 574 erties of a periodic attractor as realised by neurons (Figure 1 - figure supplement 1).
 575 Each neuron’s membrane dynamics were given by $\tau_a \dot{a}_i = -a_i(t) + c_i(t) + \sum_{j=1}^3 w_{ji} r_j(t) -$
 576 $\gamma y_i(t)$, with adaptation dynamics $\tau_y \dot{y}_i = -y_i(t) + r_i(t)$, and output firing rate $r_i(t) =$
 577 $\max\{0, a_i(t)\}$. Weights $w_{ji} \leq 0$ give the strength of inhibitory connections between
 578 the neurons, each of which receives a driving input c_i . This model, due to Matsuoka
 579 (Matsuoka, 1985, 1987), generates self-sustained oscillation of network firing rates given
 580 constant scalar inputs $c_i(t) = c$, despite each neuron not being an endogenous oscilla-
 581 tor: consequently the oscillations are an emergent property of the network. The time
 582 constants of membrane τ_a and adaptation τ_y dynamics, together with the strength of
 583 adaptation γ , determine the periodicity of the oscillations (Matsuoka, 1985, 1987). Here
 584 we use $\tau_a = 0.025$ s, $\tau_y = 0.2$ s, and $\gamma = 2$; input was $c_i = 3$ throughout except where
 585 noted.

586 **Recurrence analysis** Low dimensional projections of the joint population activity were
 587 obtained for each program using standard principal components analysis, applied to the
 588 covariance matrix of the spike-density functions. The d leading eigenvectors W_i of the co-
 589 variance matrix define the d principal dimensions, and the d corresponding eigenvalues are
 590 proportional to the variance accounted for by each dimension. The projection (the “prin-

591 cipal component”) onto each of the chosen dimensions is given by $p_i(t) = \sum_{k=1}^n W_i^k f^k(t)$,
 592 where the sum is taken over all n neurons in the analyzed population.

593 We used recurrence analysis (Lathrop and Kostelich, 1989; Marwan et al., 2007) to
 594 determine if the low-dimensional projection contained a stable periodic orbit. To do so, we
 595 checked if the low-dimensional projection $P(t) = (p_1(t), p_2(t), \dots, p_d(t))$ at time t recurred
 596 at some time $t + \delta$ in the future. Recurrence was defined as the first point $P(t + \delta) =$
 597 $(p_1(t + \delta), p_2(t + \delta), \dots, p_d(t + \delta))$ that was less than some Euclidean distance θ from
 598 $P(t)$. The recurrence time of point $P(t)$ is thus δs . Contiguous regions of the projection’s
 599 trajectory from $P(t)$ that remained within distance θ were excluded. Threshold θ was
 600 chosen based on the distribution of all distances between time-points, so that it was scaled
 601 to the activity levels in that particular program. Throughout we use the 10% value of that
 602 distribution as θ for robustness to noise; similar periodicity of recurrence was maintained
 603 at all tested thresholds from 2% upwards (Figure 2 - figure supplement 2).

604 We checked every time-point t between 5s after stimulation until 10s before the end of
 605 the recording (around 7770 points per program), determining whether it was or was not
 606 recurrent. We then constructed a histogram of the recurrence times using 1s bins to detect
 607 periodic orbits (Figure 2E): a large peak in the histogram indicates a high frequency of the
 608 same delay between recurrent points, and thus a periodic orbit in the system. All delays
 609 less than 5s were excluded to eliminate quasi-periodic activity due to noise in otherwise
 610 contiguous trajectories. Peaks were then defined as contiguous parts of the histogram
 611 between empty bins, and which contained more than 100 recurrent points. Programs had
 612 between one and four such periodic orbits. The peak containing the greatest number of
 613 recurrent points was considered the dominant periodic orbit of the program; the majority
 614 of programs had more than 50% of their recurrent points in this peak (blue-scale vectors
 615 in Figure 2E). The mean orbit period of the program was then estimated from the mean
 616 value of all recurrence times in that peak.

617 We measured the attractor’s stability as the percentage of all points that were in
 618 periodic orbits. Evolving dynamics of each program were analysed using 5 s sliding win-
 619 dows, advanced in steps of 1 s. We defined the “coalescence” time of the attractor as the
 620 mid-point of the first window in which at least 90% of the points on the trajectory were
 621 recurrent.

622 **Testing convergence to the same manifold** To determine if sequentially-evoked
 623 programs had the same manifold, we determined how closely the trajectories of each pair
 624 of programs overlapped in the low-dimensional space. We first projected all three programs
 625 from one preparation onto the principal axes of first program, to define a common low-
 626 dimensional space. For each pair of programs (A, B) in this projection, we then computed
 627 the Hausdorff distance between their two sets of recurrent points, as this metric is suited to
 628 handling tests of closeness between irregularly shaped sets of points. Given the Euclidean
 629 distances $\{d(A, B)\}$ from all recurrent points in A to those in B , and vice-versa $\{d(B|A)\}$,
 630 this is the maximum minimum distance needed to travel from a point in one program
 631 to a point in the other (namely $\max\{\min\{d(A, B)\}, \min\{d(B, A)\}\}$). To understand if
 632 the resulting distances were close, we shuffled the assignment of time-series to neurons,
 633 then projected onto the same axes giving shuffled programs A^*, B^* . These give the
 634 trajectories in the low-dimensional space determined by just the firing patterns of neurons.
 635 We then computed the shuffled Hausdorff distance $\max\{\min\{d(A, B^*)\}, \min\{d(B, A^*)\}\}$.
 636 The shuffling was repeated 100 times. Mean \pm 2SEM of the shuffled distances are plotted
 637 in (Figure 3C); the error bars are too small to see.

638 To check the robustness of the convergence to the same manifold, we repeated this

639 analysis starting from a common set of principal axes for the three programs, obtained
 640 using principal component analysis of their concatenated spike-density functions. We plot
 641 the results of this analysis in panel A of Figure 3 - figure supplement 1.

642 As a further robustness control, we sought evidence of the manifold convergence in-
 643 dependent of any low-dimensional projection. We made use of the idea that if neurons
 644 are part of sequential programs on a single manifold, then the firing of pairs of neurons
 645 should have a similar time-dependence between programs (Yoon et al., 2013; Peyrache
 646 et al., 2015). For each pair of programs (A, B) from the same preparation, we computed
 647 the similarity matrix $S(A)$ between the spike-density functions of all neuron pairs in A ,
 648 and similarly for B , giving $S(B)$. We then computed the correlation coefficient between
 649 $S(A)$ and $S(B)$: if A and B are on the same manifold, so their pairwise correlations
 650 should themselves be strongly correlated. As a control we computed a null model where
 651 each neuron has same total amount of similarity as in the data, but its pairwise similarity
 652 with each neuron is randomly distributed (Humphries, 2011). The expected value of pair-
 653 wise correlation between neurons i and j under this model is then $E_{ij} = s_i s_j / T$, where
 654 (s_i, s_j) are the total similarities for neurons i and j , and T is the total similarity in the
 655 data matrix. For comparison, we correlated $S(A)$ with E , and plot these as the control
 656 correlations in Figure 3E.

657 **Testing return to the same manifold after perturbation** We detected divergences
 658 of the trajectory away from the putative manifold, indicating spontaneous perturbations
 659 of population dynamics. We first defined potential perturbations after coalescence as a
 660 contiguous set of 5s windows when the density of recurrent points was below 90% and fell
 661 below 50% at least once. The window with the lowest recurrence density in this divergent
 662 period was labelled the divergent point. We removed all such divergent periods whose
 663 divergent point fell within 2 oscillation cycles of the end of the recording, to rule out a fall
 664 in recurrence due solely to the finite time horizon of the recording. For the remaining 19
 665 divergent periods, we then determined if the population activity returned to a recurrent
 666 state after the divergent point; that is, whether the density of recurrence returned above
 667 90% or not. The majority (17/19) did, indicating the perturbation returned to a manifold.

668 For those 17 that did, we then determined if the recurrent state post-divergence was
 669 the same manifold, or a different manifold. For it to be the same manifold after the
 670 spontaneous perturbation, then the trajectory before the perturbation should recur after
 671 the maximum divergence. To check this, we took the final window before the divergent
 672 period, and counted the proportion of its recurrent delays that were beyond the end of
 673 the divergent period, so indicating that the dynamics were in the same trajectory before
 674 and after the divergence. We plot this in Figure 3H.

675 **Statistical estimation of the attractor’s parameters** We introduce here a statistical
 676 approach to analysing the dynamics of low-dimensional projections of neural activity time-
 677 series obtained from experiments. We first fitted a linear model around each point on the
 678 low-dimensional trajectory to capture the local dynamics. For each point $P(t)$, we took the
 679 time-series of points before and after $P(t)$ that were contiguous in time and within $2.5 \times \theta$
 680 as its local neighbourhood; if less than 100 points met these criteria $P(t)$ was discarded.
 681 We then fitted the dynamical model $\dot{P}^* = AP^*$ that described the local evolution of
 682 the low-dimensional projection P^* by using linear regression to find the Jacobian matrix
 683 A ; to do so, we used the selected local neighbourhood time-series as P^* , and their first-
 684 order difference as \dot{P}^* . The maximum eigenvalue $\lambda = a + ib$ of A indicates the dominant
 685 local dynamics (Strogatz, 1994), whether contracting or expanding (sign of the real part

686 a of the eigenvalue), and whether oscillating or not (existence of the complex part of
 687 the eigenvalue i.e. $b \neq 0$). The other eigenvalues, corresponding to the $d - 1$ remaining
 688 dimensions, indicate other less-dominant dynamics; usually these were consistent across all
 689 dimensions (Figure 4 - figure supplement figure 1). We fitted A to every point $P(t)$ after
 690 the stimulation off-set, typically giving ≈ 5000 local estimates of dynamics from retained
 691 $P(t)$. The dominant dynamics for the whole program were estimated by averaging over
 692 the real a and the complex b parts of the maximum eigenvalues of the models fitted
 693 to all recurrent points in the dominant periodic orbit. The linear model's estimate of
 694 the orbit rotation period was estimated from the complex part as $\omega = 2\pi b\Delta t$, with the
 695 sampling time-step $\Delta t = 0.01$ s here. The linear model's estimate of the contraction rate
 696 is $\exp(a/\Delta t)$, which we express as a percentage.

697 **Tracking changes in periodicity over a program** We tracked changes in the oscil-
 698 lation period by first averaging the recurrence time of all recurrent points in a 5s sliding
 699 window. We then correlated the mean time with the time-point of the window to look
 700 for sustained changes in the mean period over time, considering only windows between
 701 coalescence and the final window with 90% recurrent points. We used a weighted version
 702 of Spearman's rank to weight the correlation in favour of time windows in which the tra-
 703 jectory was most clearly on the periodic orbit, namely those with a high proportion of
 704 recurrent points and low variation in recurrence time. The weighted rank correlation is:
 705 given vectors x and y of data rankings, and a vector of weights w , compute the weighted
 706 mean $m = \sum_i w_i x_i / \sum_i w_i$ and standard deviation $\sigma_{xy} = \sum_i w_i (x_i - m_x)(y_i - m_y) / \sum_i w_i$,
 707 and then the correlation $\rho = \sigma_{xy} / \sqrt{\sigma_{xx}\sigma_{yy}}$. We used the weight vector: $w_i = s_i^{-1} Q_i$, where
 708 s_i is the standard deviation of recurrence times in window i , and Q_i is the proportion of
 709 recurrent points in window i . P-values were obtained using a permutation test with 10000
 710 permutations.

711 **Decoding motor output** We decoded P10 activity from the low-dimensional trajec-
 712 tory of population activity using a generalised linear model. We first ruled out that any
 713 simultaneously recorded neuron was a motoneuron with an axon in nerve P10, by check-
 714 ing if any neurons had a high ratio of locking between their emitted spikes and spikes
 715 occurring at short latency in the P10 recording. Figure 6 - figure supplement 1 shows that
 716 no neuron had a consistent, high ratio locking of its spikes with the P10 activity.

717 We convolved the spikes of the P10 recording with a Gaussian of the same width as
 718 the spike-density functions of the simultaneously recorded program, to estimate its con-
 719 tinuous firing rate f_{10} . We fitted the model $f_{10}(t) = \exp\left(\beta_0 + \sum_{i=1}^d \sum_{h=1}^m \beta_{i,h} P_i(t-h)\right)$
 720 to determine the P10 firing rate as a function of the past history of the population activity
 721 trajectory. Using a generalised linear model here allows us to transform the arbitrary co-
 722 ordinates of the d -dimensional projection $P(t)$ into a strictly positive firing rate. Fitting
 723 used glmfit in MATLAB R2014. To cross-validate the model, we found the coefficients β
 724 using a 40s window of data, then forecast the P10 firing rate f_{10}^* using the next 10 seconds
 725 of population recording data as input to the model. Forecast error was measured as both
 726 the median absolute error and the correlation coefficient R between the actual and forecast
 727 P10 activity in the 10s window. The fitting and forecasting were repeated using a 1s step
 728 of the windows, until the final 40s+10s pair of windows available in the recording.

729 We tested activity histories between 50 and 200ms duration, with time-steps of 10ms,
 730 so that the largest model for a given program had $d \times 20$ coefficients. These short windows
 731 were chosen to rule out the contributions of other potential motoneurons in the population
 732 recording that would be phase offset from neck contraction (as 200 ms is 2% of the typical

733 period). All results were robust to the choice of history duration, so we plot results
 734 using history durations that had the smallest median absolute error in forecasting for that
 735 program.

736 **Single neuron participation** We quantified each neuron’s participation in the low-
 737 dimensional projection as the L1-norm: the absolute sum of its weights on the principal
 738 axes (eigenvectors) for program m : $\rho_i^m = \sum_{j=1}^d |\lambda_j^m W_j^m(i)|$, where the sum is over the d
 739 principal axes, $W_j^m(i)$ is the neuron’s weight on the j th axis, and λ_j^m is the axis’ corre-
 740 sponding eigenvalue. Within a program, participation for each neuron was normalised to
 741 the maximum participation in that program. To fit a noise model for the variability in
 742 participation between programs, we first computed the change in participation for each
 743 neuron between all pairs of programs in the same preparation. We then fit a Gaussian
 744 model for the noise, using an iterative maximum likelihood approach to identify the likely
 745 outliers; here the outliers are the participation changes that are inconsistent with stochas-
 746 tic noise. In this approach, we compute the mean and variance of the Gaussian from the
 747 data, eliminate the data-point furthest from the estimate of the mean, re-estimate the
 748 mean and variance, and compute the new log likelihood of the Gaussian model without
 749 that data-point. We iterate elimination, re-estimation, and likelihood computation until
 750 the likelihood decreases. The final model (mean and variance) found before the decrease
 751 is then the best-fit Gaussian model to the bulk of the data. Neurons whose maximum
 752 change in participation exceeded a threshold of the mean $\pm 3\text{SD}$ of that best-fit model
 753 were then considered “strongly variable” neurons.

754 We asked whether the variation in low-dimensional dynamics of sequentially-evoked
 755 programs was a consequence of the degree of variation in single neuron participation.
 756 Between a pair of consecutively evoked programs, we quantified the variation in their
 757 low dimensional dynamics as the Hausdorff distance between them, normalised by the
 758 mean distance between their random projections. This normalisation allowed us to put
 759 all programs on a single scale measuring the closeness relative to random projections,
 760 such that 1 indicates equivalence to a random projection, < 1 indicates closer than ran-
 761 dom projections, and > 1 indicates further apart than random projections. For a given
 762 pair of programs, we quantified the variability of individual neurons’ participation in two
 763 ways: by summing the change in participation of each neuron between the programs; and
 764 by computing the Hellinger distance between the two distributions of participation (one
 765 distribution per program).

766 **Participation maps** Each neuron’s (x,y) location on the plane of the photodiode array
 767 could be estimated from the weight matrix from the independent component analysis of
 768 the original 464 photodiode time-series; see (Bruno et al., 2015) for full details. We were
 769 able to reconstruct locations for all neurons in 8 of the 10 recorded preparations; for
 770 the other two preparations, partial corruption of the original spike-sorting analysis data
 771 prevented reconstructions of some neuron locations in one; for the other, we could not
 772 determine on what side it was recorded. We merged all left or right ganglion recordings
 773 on to a common template of the photodiode array. The marker sizes and colour codes for
 774 each neuron were proportional to the normalised maximum participation of that neuron
 775 (Figure 8A,C) and to the range of normalised maximum participation across the three
 776 programs (Figure 8B,D).

777 **Acknowledgements** We thank R. Petersen and B. Mensh for comments on drafts,
 778 A. Singh for suggesting the Hausdorff distance, and J. Wang for technical assistance.

779 M.D.H was supported by a Medical Research Council Senior non-Clinical Fellowship. W.F.
780 was supported by NIH R01NS060921 and NSF 1257923. A.B. was supported by NIH
781 F31NS079036.

782 References

- 783 Ahrens, M. B., Li, J. M., Orger, M. B., Robson, D. N., Schier, A. F., Engert, F. and Por-
784 tugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish.
785 *Nature* *485*, 471–477.
- 786 Bartho, P., Curto, C., Luczak, A., Marguet, S. L. and Harris, K. D. (2009). Population
787 coding of tone stimuli in auditory cortex: dynamic rate vector analysis. *Eur J Neurosci*
788 *30*, 1767–1778.
- 789 Berg, R. W., Alaburda, A. and Hounsgaard, J. (2007). Balanced inhibition and excitation
790 drive spike activity in spinal half-centers. *Science* *315*, 390–393.
- 791 Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S. and Mitra, P. P. (2010). Chronux: a
792 platform for analyzing neural signals. *J Neurosci Methods* *192*, 146–151.
- 793 Briggman, K. L., Abarbanel, H. D. I. and Kristan, Jr, W. (2005). Optical imaging of
794 neuronal populations during decision-making. *Science* *307*, 896–901.
- 795 Briggman, K. L., Abarbanel, H. D. I. and Kristan, W. B. (2006). From crawling to
796 cognition: analyzing the dynamical interactions among populations of neurons. *Curr*
797 *Opin Neurobiol* *16*, 135–144.
- 798 Briggman, K. L. and Kristan, W. B. (2008). Multifunctional pattern-generating circuits.
799 *Annu Rev Neurosci* *31*, 271–294.
- 800 Briggman, K. L. and Kristan, Jr, W. B. (2006). Imaging dedicated and multifunctional
801 neural circuits generating distinct behaviors. *J Neurosci* *26*, 10925–10933.
- 802 Brody, C. D., Romo, R. and Kepecs, A. (2003). Basic mechanisms for graded persistent
803 activity: discrete attractors, continuous attractors, and dynamic representations. *Curr*
804 *Opin Neurobiol* *13*, 204–211.
- 805 Brown, E. N., Kass, R. E. and Mitra, P. P. (2004). Multiple neural spike train data
806 analysis: state-of-the-art and future challenges. *Nat Neurosci* *7*, 456–461.
- 807 Bruno, A. M., Frost, W. N. and Humphries, M. D. (2015). Modular Deconstruction
808 Reveals the Dynamical and Physical Building Blocks of a Locomotion Motor Program.
809 *Neuron* *86*, 304–318.
- 810 Calin-Jageman, R. J., Tunstall, M. J., Mensh, B. D., Katz, P. S. and Frost, W. N. (2007).
811 Parameter space analysis suggests multi-site plasticity contributes to motor pattern
812 initiation in *Tritonia*. *J Neurophysiol* *98*, 2382–2398.
- 813 Carmena, J. M., Lebedev, M. A., Henriquez, C. S. and Nicolelis, M. A. L. (2005). Sta-
814 ble ensemble performance with single-neuron variability during reaching movements in
815 primates. *J Neurosci* *25*, 10712–10716.
- 816 Carroll, M. S. and Ramirez, J.-M. (2013). Cycle-by-cycle assembly of respiratory network
817 activity is dynamic and stochastic. *Journal of neurophysiology* *109*, 296–305.

-
- 818 Chestek, C. A., Batista, A. P., Santhanam, G., Yu, B. M., Afshar, A., Cunningham,
819 J. P., Gilja, V., Ryu, S. I., Churchland, M. M. and Shenoy, K. V. (2007). Single-
820 neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* *27*,
821 10742–10750.
- 822 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P.,
823 Ryu, S. I. and Shenoy, K. V. (2012). Neural population dynamics during reaching.
824 *Nature* *487*, 51–56.
- 825 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. and Shenoy, K. V.
826 (2010). Cortical preparatory activity: representation of movement or first cog in a
827 dynamical machine? *Neuron* *68*, 387–400.
- 828 Cunningham, J. P. and Yu, B. M. (2014). Dimensionality reduction for large-scale neural
829 recordings. *Nat Neurosci* *17*, 1500–1509.
- 830 Eliasmith, C. (2005). A unified approach to building and controlling spiking attractor
831 networks. *Neural Comput* *17*, 1276–1314.
- 832 Flinn, J. M., Gochman, P., Wanschura, P. and Chandhoke, V. (2001). The effect of
833 dopamine receptor blockade on motor behavior in *Aplysia californica*. *Pharmacol*
834 *Biochem Behav* *69*, 425–430.
- 835 Fredman, S. M. and Jahan-Parwar, B. (1980). Role of pedal ganglia motor neurons in
836 pedal wave generation in *Aplysia*. *Brain Res Bull* *5*, 179–193.
- 837 Georgopoulos, A. P., Schwartz, A. B. and Kettner, R. E. (1986). Neuronal population
838 coding of movement direction. *Science* *233*, 1416–1419.
- 839 Getting, P. A. (1989). Emerging principles governing the operation of neural networks.
840 *Annu Rev Neurosci* *12*, 185–204.
- 841 Hall, J. D. and Lloyd, P. E. (1990). Involvement of pedal peptide in locomotion in *Aplysia*:
842 modulation of foot muscle contractions. *J Neurobiol* *21*, 858–868.
- 843 Harvey, C. D., Coen, P. and Tank, D. W. (2012). Choice-specific sequences in parietal
844 cortex during a virtual-navigation decision task. *Nature* *484*, 62–68.
- 845 Hatsopoulos, N. G., Xu, Q. and Amit, Y. (2007). Encoding of movement fragments in the
846 motor cortex. *J Neurosci* *27*, 5105–5114.
- 847 Hawkins, R. D., Clark, G. A. and Kandel, E. R. (2006). Operant conditioning of gill
848 withdrawal in *Aplysia*. *J Neurosci* *26*, 2443–2448.
- 849 Hening, W. A., Walters, E. T., Carew, T. J. and Kandel, E. R. (1979). Motorneuronal
850 control of locomotion in *Aplysia*. *Brain Res* *179*, 231–253.
- 851 Hill, E. S., Moore-Kochlacs, C., Vasireddi, S. K., Sejnowski, T. J. and Frost, W. N. (2010).
852 Validation of independent component analysis for rapid spike sorting of optical recording
853 data. *J Neurophysiol* *104*, 3721–3731.
- 854 Hill, E. S., Vasireddi, S. K., Bruno, A. M., Wang, J. and Frost, W. N. (2012). Variable
855 Neuronal Participation in Stereotypic Motor Programs. *PLoS One* *7*, e40579.

-
- 856 Hill, E. S., Vasireddi, S. K., Wang, J., Bruno, A. M. and Frost, W. N. (2015). Memory
857 Formation in Tritonia via Recruitment of Variably Committed Neurons. *Curr Biol* *25*,
858 2879–2888.
- 859 Huber, D., Gutnisky, D. A., Peron, S., O’Connor, D. H., Wiegert, J. S., Tian, L., Oertner,
860 T. G., Looger, L. L. and Svoboda, K. (2012). Multiple dynamic representations in the
861 motor cortex during sensorimotor learning. *Nature* *484*, 473–478.
- 862 Humphries, M. D. (2011). Spike-train communities: finding groups of similar spike trains.
863 *J Neurosci* *31*, 2321–2336.
- 864 Humphries, M. D. (2017). `AplysiaAttractorAnalysis`.
865 <https://github.com/mdhumphries/AplysiaAttractorAnalysis>.
866 941b08dc5bc9958eff74c59827cbdac131e225f3.
- 867 Jahan-Parwar, B. and Fredman, S. M. (1979). Neural control of locomotion in Aplysia:
868 role of the central ganglia. *Behav Neural Biol* *27*, 39–58.
- 869 Jahan-Parwar, B. and Fredman, S. M. (1980). Motor program for pedal waves during
870 Aplysia locomotion is generated in the pedal ganglia. *Brain Res Bull* *5*, 169–177.
- 871 Jing, J., Vilim, F. S., Cropper, E. C. and Weiss, K. R. (2008). Neural analog of arousal:
872 persistent conditional activation of a feeding modulator by serotonergic initiators of
873 locomotion. *J Neurosci* *28*, 12349–12361.
- 874 Kato, S., Kaplan, H. S., Schrödel, T., Skora, S., Lindsay, T. H., Yemini, E., Lockery, S.
875 and Zimmer, M. (2015). Global Brain Dynamics Embed the Motor Command Sequence
876 of *Caenorhabditis elegans*. *Cell* *163*, 656–669.
- 877 Katz, P. S., Getting, P. A. and Frost, W. N. (1994). Dynamic neuromodulation of synaptic
878 strength intrinsic to a central pattern generator circuit. *Nature* *367*, 729–731.
- 879 Katz, P. S., Sakurai, A., Clemens, S. and Davis, D. (2004). Cycle period of a network
880 oscillator is independent of membrane potential and spiking activity in individual central
881 pattern generator neurons. *J Neurophysiol* *92*, 1904–1917.
- 882 Kleinfeld, D. and Sompolinsky, H. (1988). Associative neural network model for the
883 generation of temporal patterns. Theory and application to central pattern generators.
884 *Biophys J* *54*, 1039–1051.
- 885 Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C. E., Kepecs, A., Mainen, Z. F.,
886 Qi, X.-L., Romo, R., Uchida, N. and Machens, C. K. (2016). Demixed principal com-
887 ponent analysis of neural population data. *Elife* *5*, e10989.
- 888 Koch, C. (2012). Systems biology. Modular biological complexity. *Science* *337*, 531–532.
- 889 Kuenzi, F. M. and Carew, T. J. (1994). Head waving in *Aplysia californica*. III. Inter-
890 ganglionic pathways underlying the coordination and control of searching movements.
891 *J Exp Biol* *195*, 75–90.
- 892 Lathrop, D. P. and Kostelich, E. J. (1989). Characterization of an experimental strange
893 attractor by periodic orbits. *Phys Rev A* *40*, 4028–4031.
- 894 Leonard, J. L. and Lukowiak, K. (1986). The Behavior of *Aplysia californica* Cooper
895 (Gastropoda; Opisthobranchia): I. Ethogram. *Behaviour* *98*, 320–360.

-
- 896 Levi, R., Varona, P., Arshavsky, Y. I., Rabinovich, M. I. and Selverston, A. I. (2005).
897 The role of sensory network dynamics in generating a motor program. *J Neurosci* *25*,
898 9807–9815.
- 899 Lewis, J. E. and Kristan, Jr, W. (1998). A neuronal network for computing population
900 vectors in the leech. *Nature* *391*, 76–79.
- 901 Mante, V., Sussillo, D., Shenoy, K. V. and Newsome, W. T. (2013). Context-dependent
902 computation by recurrent dynamics in prefrontal cortex. *Nature* *503*, 78–84.
- 903 Marder, E. and Bucher, D. (2007). Understanding circuit dynamics using the stomatogas-
904 tric nervous system of lobsters and crabs. *Ann Rev Physiol* *69*, 291–316.
- 905 Marder, E., Goeritz, M. L. and Otopalik, A. G. (2015). Robust circuit rhythms in small
906 circuits arise from variable circuit components and mechanisms. *Curr Opin Neurobiol*
907 *31C*, 156–163.
- 908 Marder, E. and Taylor, A. L. (2011). Multiple models to capture the variability in biological
909 neurons and networks. *Nat Neurosci* *14*, 133–138.
- 910 Marinesco, S., Wickremasinghe, N., Kolkman, K. E. and Carew, T. J. (2004). Serotonergic
911 modulation in *Aplysia*. II. Cellular and behavioral consequences of increased serotonergic
912 tone. *J Neurophysiol* *92*, 2487–2496.
- 913 Marr, D. (1982). *Vision*. W. H. Freeman, San Francisco.
- 914 Marwan, N., Romano, M. C., Thiel, M. and Kurths, J. (2007). Recurrence plots for the
915 analysis of complex systems. *Physics Reports* *438*, 237–329.
- 916 Matsuoka, K. (1985). Sustained oscillations generated by mutually inhibiting neurons with
917 adaptation. *Biol Cybern* *52*, 367–376.
- 918 Matsuoka, K. (1987). Mechanisms of frequency and pattern control in the neural rhythm
919 generators. *Biol Cybern* *56*, 345–353.
- 920 Mazor, O. and Laurent, G. (2005). Transient dynamics versus fixed points in odor repre-
921 sentations by locust antennal lobe projection neurons. *Neuron* *48*, 661–673.
- 922 McPherson, D. R. and Blankenship, J. E. (1992). Neuronal modulation of foot and body-
923 wall contractions in *Aplysia californica*. *J Neurophysiol* *67*, 23–28.
- 924 Miller, P. (2016). Dynamical systems, attractors, and neural circuits. *F1000Research* *5*,
925 992.
- 926 Panas, D., Amin, H., Maccione, A., Muthmann, O., van Rossum, M., Berdondini, L. and
927 Hennig, M. H. (2015). Sloppiness in spontaneously active neuronal networks. *J Neurosci*
928 *35*, 8480–8492.
- 929 Pasemann, F. (1995). Characterization of periodic attractors in neural ring networks.
930 *Neural Networks* *8*, 421 – 429.
- 931 Petersen, P. C. and Berg, R. W. (2016). Lognormal firing rate distribution reveals promi-
932 nent fluctuation-driven regime in spinal motor networks. *eLife* *5*, e18805.
- 933 Peyrache, A., Lacroix, M. M., Petersen, P. C. and Buzsáki, G. (2015). Internally organized
934 mechanisms of the head direction sense. *Nat Neurosci* *18*, 569–575.

-
- 935 Portugues, R., Feierstein, C. E., Engert, F. and Orger, M. B. (2014). Whole-brain activity
936 maps reveal stereotyped, distributed networks for visuomotor behavior. *Neuron* *81*,
937 1328–1343.
- 938 Rokni, U. and Sompolinsky, H. (2012). How the brain generates movement. *Neural*
939 *Comput* *24*, 289–331.
- 940 Seelig, J. D. and Jayaraman, V. (2015). Neural dynamics for landmark orientation and
941 angular path integration. *Nature* *521*, 186–191.
- 942 Stopfer, M. and Carew, T. J. (1988). Development of sensitization in the escape locomotion
943 system in *Aplysia*. *J Neurosci* *8*, 223–230.
- 944 Strogatz, S. H. (1994). *Nonlinear Dynamics and Chaos*. Westview Press, Cambridge, MA.
- 945 Szucs, A. (1998). Applications of the spike density function in analysis of neuronal firing
946 patterns. *J Neurosci Methods* *81*, 159–167.
- 947 Wu, J. Y., Cohen, L. B. and Falk, C. X. (1994). Neuronal activity during different
948 behaviors in *Aplysia*: a distributed organization? *Science* *263*, 820–823.
- 949 Xin, Y., Weiss, K. R. and Kupfermann, I. (1996). An identified interneuron contributes
950 to aspects of six different behaviors in *Aplysia*. *J Neurosci* *16*, 5266–5279.
- 951 Yoon, K., Buice, M. A., Barry, C., Hayman, R., Burgess, N. and Fiete, I. R. (2013).
952 Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat*
953 *Neurosci* *16*, 1077–1084.
- 954 Yuste, R. (2015). From the neuron doctrine to neural networks. *Nat Rev Neurosci* *16*,
955 487–497.