

A spiral attractor network drives locomotion in *Aplysia*

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.

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

Abstract

The neural control of motor behaviour arises from the joint activity of large neuron populations. Unknown is what underlying dynamical system generates this joint activity. Here we show that the network-wide activity driving locomotion of the sea-slug *Aplysia* is a low-dimensional spiral attractor. We imaged large populations at single-spike resolution from the *Aplysia*'s pedal ganglion during fictive locomotion. Evoking locomotion rapidly moved the population activity from irregular spontaneous activity into a low-dimensional, periodic, decaying, orbit. This orbit was a true attractor: the activity returned to the same orbit after transient perturbation; and repeatedly evoking locomotion caused the activity to converge on the same low-dimensional orbit. To show this attractor was the locomotion program, we accurately decoded simultaneous recordings of neck motorneuron activity directly from the low-dimensional orbit. Our results provide direct evidence that neural circuits are periodic attractors. Consequently, they support the long-held hypothesis that population activity is an emergent property of a simpler underlying system.

Neural control of motor behaviour arises from the mass action of neuron populations¹⁻⁶. While individual neuron activity can correlate with specific aspects of movement^{7,8}, the intermittent participation of individual neurons across repeated movements⁹⁻¹² suggests that only the collective population activity, and not specifics of single neuron firing, are key to motor control. Thus, a parsimonious theory of motor control requires unravelling the underlying system embodied by a population¹³⁻¹⁸.

The idea that high-dimensional population activity is an emergent property of an underlying low-dimensional dynamical system has a long history in neural theory^{19,20}. A key hypothesis is that stable repetitions of population activity arise from an underlying, low-dimensional attractor. Stable repetitions of low-dimensional population activity have been described in arm movements^{15,16}, navigation systems^{21,22}, and sensory encoding^{23,24}, but the underlying dynamical system has not been determined. It is thus unknown whether high-dimensional population activity can emerge from a true low-dimensional attractor.

To directly test for the presence of an underlying attractor in motor control, we imaged population activity at single-neuron, single-spike resolution in the pedal ganglion of *Aplysia* during fictive locomotion (Fig 1a). The pedal ganglion presents an ideal target for this problem as it contains the pattern generator^{25,26}, motoneurons^{27,28} and modulatory neurons^{29,30} underlying locomotion.

Using this model system, we show here that evoking fictive locomotion caused heterogeneous population activity, but under which always lay a low-dimensional, slowly decaying periodic orbit. This periodic trajectory met the convergence and perturbation criteria for an attractor. Crucially, we identify the attractor as a stable, decaying spiral, and show there is this same, single attractor within the pedal ganglion of all tested animals. While individual neurons varied their participation in the attractor between bouts of locomotion, we could nonetheless decode motoneuron activity directly from the low-dimensional orbit. Our results thus provide experimental support for the long-standing theory that neural population activity is a high-dimensional emergent property of a low-dimensional system.

Results

We sequentially evoked 3 bouts of fictive locomotion in each of 10 isolated central nervous system preparations (Fig 1b). Each bout of locomotion was evoked by short stimulation of the tail nerve P9, mimicking a sensory stimulus to the tail that elicits the escape locomotion response²⁷; in intact animals, a strong tail stimulus typically elicits a two-part escape behavior consisting of several cycles of a vigorous arching gallop, followed by several minutes of a more sedate rhythmic crawl^{25,31}. Recorded populations from the pedal ganglion comprised 120-180 neurons each, representing $\approx 10\%$ of the network in each recording. The population recordings captured rich, varied single neuron dynamics within the ganglion's network following the stimulus (Fig 1c). A dominant, slow (≤ 0.1 Hz) oscillation in neural firing (Fig. 1d) is consistent with the periodic activity necessary to generate rhythmic locomotion. But the variety of single neuron dynamics³² (Fig 1c) and the slowly decaying population firing rate (Fig. 1e) post-stimulus hint at a more complex underlying dynamical system driving locomotion than a simple, consistent oscillator.

We show here that the joint activity of the population meets the necessary conditions for a periodic attractor (Fig. 1f). We identified these as: (1) applying a driving force causes the system's dynamics to fall onto a stable, periodic orbit; (2) repeatedly driving the system causes convergence of the dynamics to the same orbit; and (3) the system should return to the periodic orbit after the end of transient perturbation. Supplementary Fig. 1 demonstrates these conditions in a dynamical model.

Joint population activity forms a low-dimensional periodic orbit

We first show that under the heterogeneous population activity evoked by the tail-nerve stimulation there is a low dimensional periodic trajectory, consistent with there being a periodic attractor in the pedal ganglion. Projections of a population's joint activity into three dimensions typically showed that stimulation caused a strong deviation from the spontaneous state, which then settled into repeated loops (Fig. 2a). Capturing a significant proportion of the population variance generally required 4-8 embedding dimensions (Fig. 2b), representing a dimension reduction by more than a factor of 10 compared to the number of neurons. However, we cannot directly visualise this space; therefore we cannot tell by visual inspection if the low-dimensional trajectory repeatedly returns to the same position, and so is truly periodic.

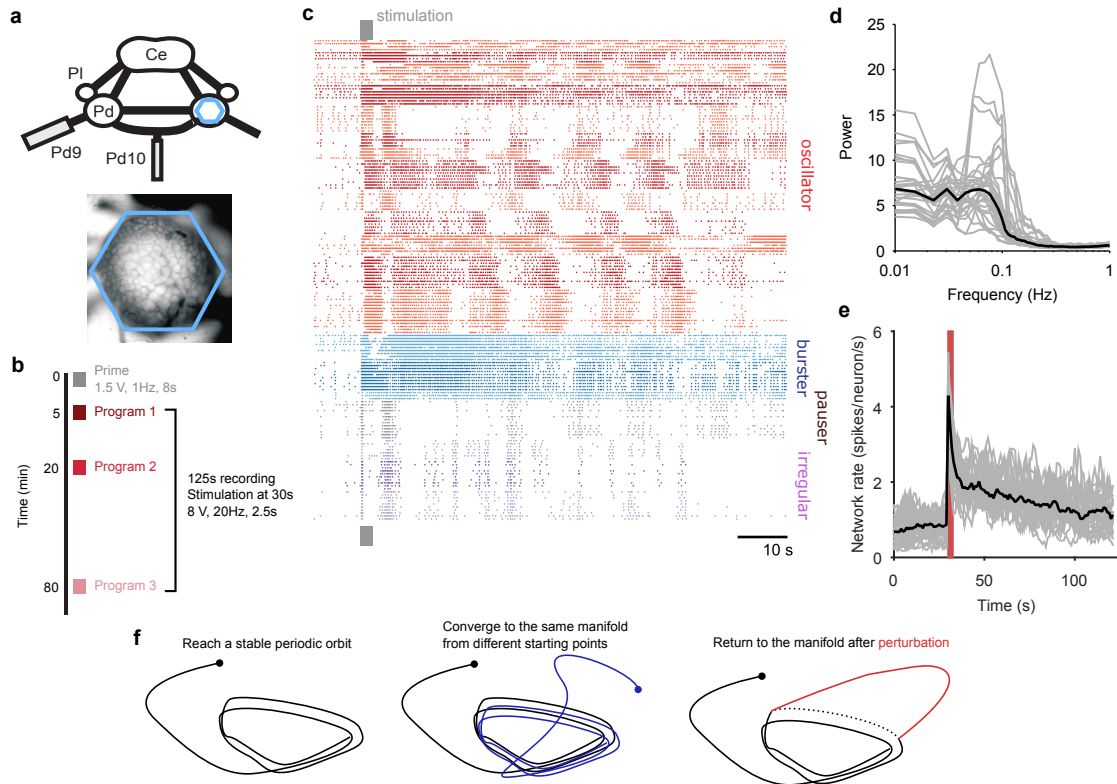


Figure 1: Population dynamics during fictive locomotion. **a**, Voltage-sensitive dye recording of the pedal ganglion (Pd) network in an isolated central nervous system preparation (top) using a photodiode array (blue hexagon). The array covered the dorsal surface of the ganglion (bottom). Ce: cerebral 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. **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 program; black: mean over programs). **e** Network firing rate over time (grey: every program; black: mean; red bar: stimulation duration. Bins: 1s). **f** Schematic illustration of the necessary conditions for identifying a periodic attractor (or "cyclical" attractor). Dot indicates start time; line indicates trajectory of the dynamical system; dashed line indicates deviation from the normal trajectory. See Supplementary Fig. 1 for a dynamical model illustrating these conditions.

To determine whether population activity in higher dimensions reached a stable periodic orbit, we made use of the idea of recurrence^{33,34}. For each time-point in the low-dimensional trajectory of the population's activity, we check if the trajectory passes close to the same point in the future (Fig. 2c). If so, then the current time-point *recurs*, indicating that the joint activity of the population revisits the same state at least once. The time between the current time-point and when it recurs gives us the period of recurrence. A strongly periodic system would thus be revealed by its population's trajectory having many recurrent points with similar recurrence periods; random or chaotic dynamics, by contrast, would not show a single clustered recurrence period.

Plotting recurrent time-points showed that the evoked low-dimensional population activity typically recurred with a regular period (example in Fig. 2d). We found strongly periodic recurrence on the scale of 10-15 s in many but not all of the 30 evoked population responses (Fig. 2e,f). This reflected the range of stimulation responses from strongly periodic activity across the population to noisy, stuttering, irregular activity (Supplementary Fig. 2). Nonetheless, despite this heterogeneity across stimulus responses, the activity of almost all populations was dominated by a single periodic orbit (Fig. 2e), robust to the choice of threshold for defining recurrence (Supplementary Fig. 3).

Joint population activity meets the conditions for a periodic attractor

The trajectory of the population activity describes a circumscribed region of low-dimensional space - the manifold - within which it remains. To truly be an attractor (Fig. 1f; Supplementary Fig. 1), the system's dynamics should rapidly reach and stay on the manifold when evoked; reach the same manifold every time it is evoked; and return to the manifold when perturbed. We found that almost all evoked population responses quickly reached the attractor, within one oscillation period (Fig. 3a), and were then dominated by recurrence, indicating they stayed on the manifold.

The key problem in analysing any putative attractor from experimental data is identifying when the experimentally-measured dynamics are or are not on the attractor, whether due to perturbations of the system or noise in the measurements. The set of recurrent time-points allowed us to identify when the joint population activity was most likely on the attractor.

To determine if sequentially-evoked responses from the same preparation reached the same manifold, we projected all 3 population responses into the same set of embedding dimensions, using only the recurrent points (Fig. 3b; Supplementary Fig. 4 shows these results are robust to other projections). Falling on the same manifold would mean that every recurrent point in one population would also appear in both others, if noiseless. Consequently, the maximum distance between any randomly chosen recurrent point in population response A and the closest recurrent point in population response B should be small. We defined small here as being shorter than the expected distance between a recurrent point in A and the closest point on a random walk of the activity in the same embedding dimensions. Despite the inherent noise and limited duration of the recordings, this is exactly what we found: pairs of evoked population responses from the same preparation fell close to each other throughout (Fig. 3c), well in excess of the expected agreement between random projections of the data onto the same embedding dimensions.

We also checked that this convergence to the same manifold came from different initial conditions. The initiating stimulation is a rough kick to the system - indeed a fictive locomotion bout can be initiated with a variety of stimulation parameters³² - applied to ongoing spontaneous activity. Together, the stimulation and the state of spontaneous activity when it is applied should give different initial conditions from which the locomotion

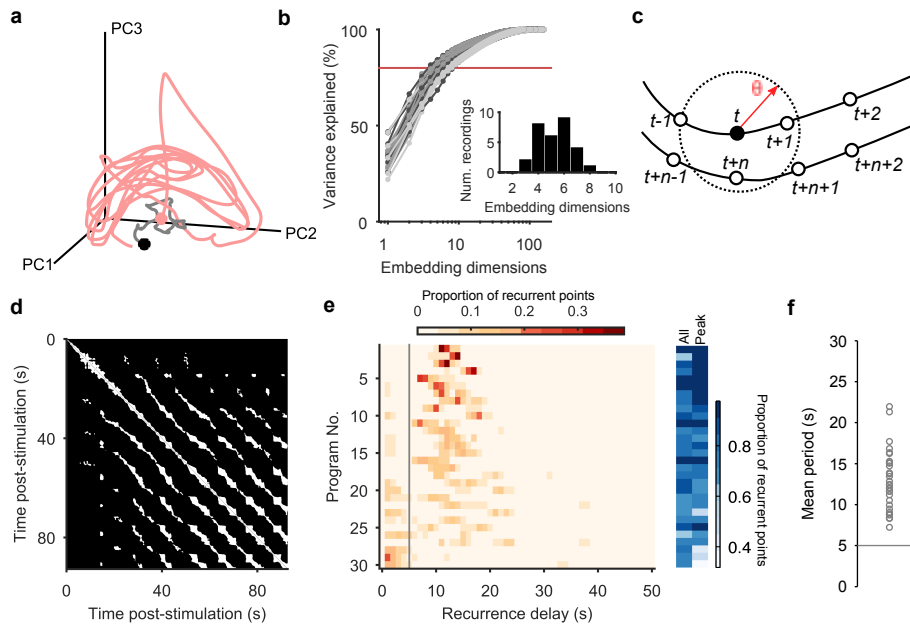


Figure 2: Population dynamics form a low-dimensional periodic orbit. **a**, Projection of one evoked population response into 3 embedding dimensions, given by its first 3 principal components (PCs). Dots: start of recording (black) and stimulation (pink); spontaneous activity is shown in grey. Smoothed 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 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.

manifold is reached. We found that the stimulation caused population responses within the same preparation to diverge far more than in either the spontaneous activity or after coalescing to the manifold (Fig. 3d). Thus, a wide range of initial driven dynamics in the pedal ganglion population converged onto the same manifold.

Previous studies have used the consistency of pairwise correlations between neurons across conditions as indirect evidence for the convergence of population activity to an underlying attractor^{22,35}. The intuition here is that neurons whose activity contributes to the same portion of the manifold will have simultaneous spiking, and so their activity will correlate across repeated visits to the manifold. To check this, we computed the pairwise similarity between all neurons within an evoked population response (Fig. 3e), then correlated these similarity matrices between responses from the same preparation. We found that pair-wise similarity is indeed well-preserved across population responses in the same preparation (Fig. 3f). This also shows that the apparent convergence to the same manifold is not an artifact of our choice of low-dimensional projection.

In many population responses, we noticed spontaneous perturbations of the low-dimensional dynamics away from the trajectory (examples in Supplementary Fig. 5), indicated by sudden falls in the density of recurrent points (Fig. 3g). That is, perturbations could be detected by runs of contiguous points on the population trajectory that were not recurrent. In most cases (90%), the population dynamics returned to a recurrent state after the spontaneous perturbation (Fig. 3h); the two perturbations that did not were consistent with the end of the evoked fictive locomotion and a return to spontaneous activity (Supplementary Fig. 5). Of those that returned, all but three clearly returned to the same manifold (Fig. 3i); for those three, the spontaneous perturbation appeared sufficient to move the population dynamics into a different periodic attractor (Supplementary Fig. 5). Potentially, these are the known transitions from the escape gallop to normal crawling³¹. The low dimensional dynamics of the pedal ganglion thus meet the stability, manifold convergence, and perturbation criteria of a periodic attractor network.

Heterogenous population activity arises from a common attractor

While these results show the existence of a periodic orbit on an attractor in the evoked population responses, they cannot address whether these arise from the same putative attractor within and, crucially, between animals. To determine if there is a common underlying attractor despite the heterogeneity in spiking patterns across the population responses (Supplementary Fig. 2), we introduced a statistical approach to quantifying the low-dimensional trajectory. We fitted a linear model of the local dynamics around each time point (see Methods), which allowed us to then estimate the local dynamics based on the eigenvalues of that model: whether the trajectory is expanding away or contracting towards this point, and whether or not it is rotating.

The time-series of eigenvalues typically showed long periods of similar magnitude eigenvalues, corresponding to the recurrent points (Fig 4a). Consequently, by averaging over the local dynamics only for recurrent points, we could potentially capture the dynamics of the underlying attractor. Doing so, we found that the evoked population responses had highly clustered eigenvalues (Fig. 4b,c), and thus highly similar underlying dynamics despite the apparent heterogeneity of spike-train patterns between them. The dominance of negative complex eigenvalues implies the pedal ganglion network implements a low-dimensional system containing a stable spiral - a decaying periodic orbit.

The identification of a stable spiral makes a clear prediction for what should and should not change over time in the dynamics of the population. The negative complex eigenvalues mean that the magnitude of the orbit decays over time, corresponding to the

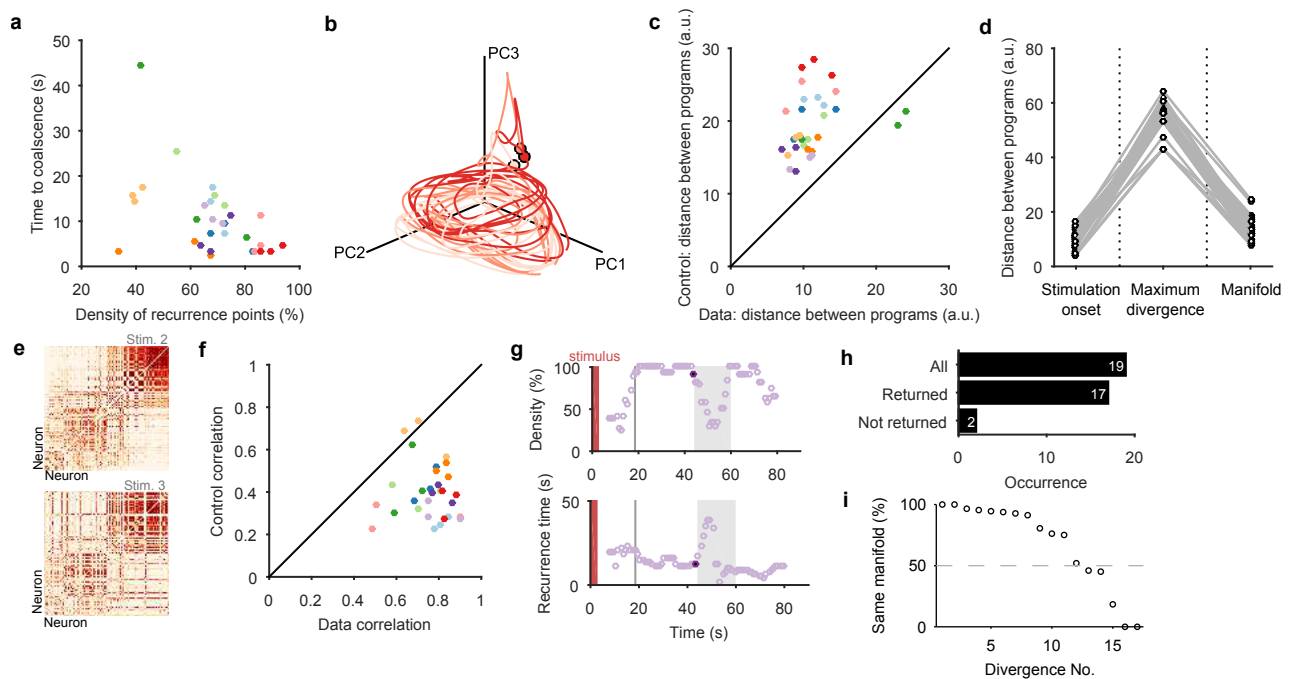


Figure 3: Low dimensional population dynamics are consistent with a periodic attractor. **a** Distribution of the time the population dynamics took to coalesce onto the attractor from the stimulation onset, and the subsequent stability of the attractor (measured by the proportion of recurrent points). 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 later. **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 (Supplementary Fig. 4). **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 (Supplementary Fig. 4). **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. 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.

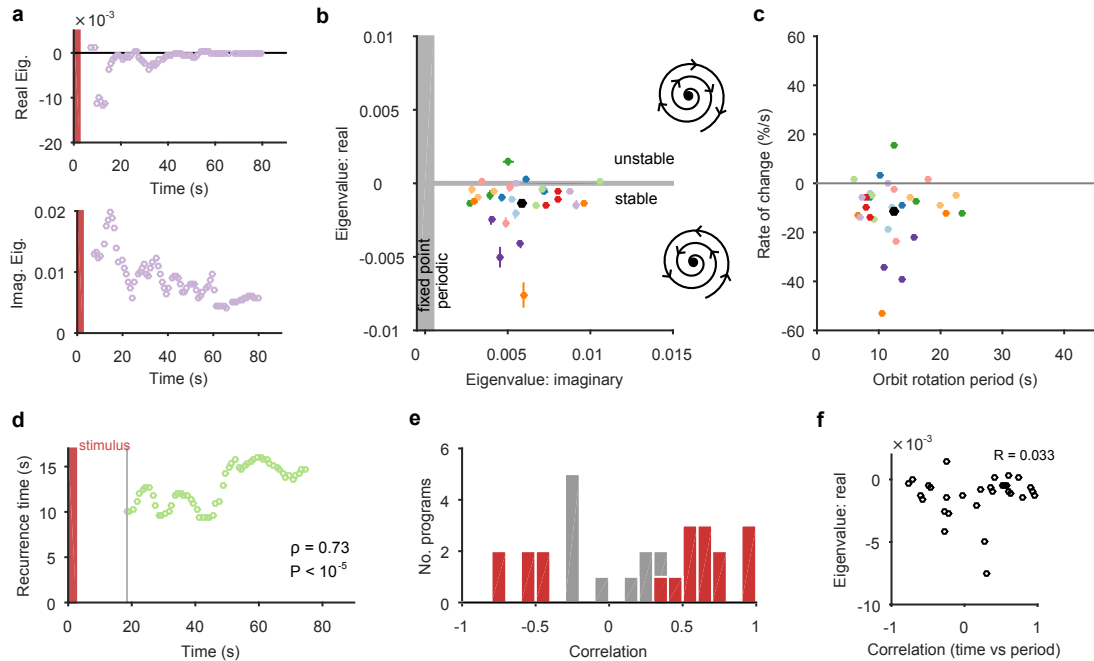


Figure 4: Uncovering consistent periodic dynamics in heterogenous activity. **a**, Example time-series from population response of the real (top) and imaginary (bottom) component of the maximum eigenvalue for the local linear model. Points are averages over a 5s sliding window. Red bar indicates stimulus duration. **b** Dominant dynamics for each evoked population response. Dots and lines give means ± 2 s.e.m. of real and imaginary components of 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 **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 **b**.) **d** Example of a change in the local estimate of the periodic orbit over a population response (here, slowing over time; Spearman's rank correlation). **e** Histogram of correlations between time elapsed and local estimate of the periodic orbit (positive: slowing; negative: speeding up). Red bars correspond to population responses with $P < 0.01$ (permutation test). **f** Relationship between the change in periodic orbit over time and the rate of contraction for each population response (Pearson's R ; $n = 30$ responses).

decreasing population spike rate in most evoked responses (Fig. 1e). However, a stable spiral indicates only a decrease in magnitude; it does not mean the orbital period is also slowing. Consequently, the presence of a stable spiral attractor predicts that the magnitude and period of the orbit are dissociable properties in the pedal ganglion network.

We checked this prediction using the linear model. The linear model estimated a mean orbital period of around 10 s (Fig. 4c), consistent with the directly-derived estimate from the recurrent points (Fig. 2f). This indicated the linear model was correctly capturing the local dynamics of each program. But our linear model also gave us a time-series of estimates of the local orbital period (Fig 4d), which we could use to check whether the orbital period was changing during each evoked response. We found that the population responses included all possible changes in periodic orbit: slowing, speeding up, and not changing (Fig 4e). As predicted there was no relationship between the contraction of the periodic orbit and its change in period (Fig 4f).

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

Collectively, these periodic, decaying dynamics are ethologically consistent with locomotion that comprises a repeated sequence of movements that decays in intensity over time^{31,36,37}. If this putative low-dimensional periodic attractor is the “motor program” for locomotion, then we should be able to decode the locomotion muscle commands from its trajectory. In 3 of the 10 preparations we were able to simultaneously record activity from the P10 nerve that projects to the neck muscles³⁸ for all three evoked population responses. The spiking of axons in this nerve should thus correspond to the specific neck contraction portion of the cyclical escape locomotion response. We thus sought to decode the spiking of P10 directly from the low-dimensional population trajectory (Fig. 5a).

We first confirmed that each recorded neural population did not appear to contain any motoneurons with axons in P10, which could make the decoding potentially trivial (Supplementary Fig. 6). To then decode P10 activity, we used a statistical model that predicts the firing rate of nerve P10 at each time point, by weighting and summing the recent history (up to 100 ms) of the trajectory in the low dimensional space, and using a non-linearity to convert this weighted sum into a firing rate. We controlled for over-fitting using cross-validation forecasting: we fit the model to a 40 s window of trajectory data, and predicted the next 10 s of P10 activity (Fig. 5b). By sliding the window over the data, we could assess the quality of the forecast over the entire recording (Fig. 5c).

The model could accurately fit (Supplementary Fig. 7) and forecast (Fig. 5b) P10 activity from the low-dimensional trajectory in all 9 population responses. Emphasising the quality of the model, in Fig. 5d we plot example forecasts of the entire P10 recording based on fitting only to the first 40s window, each example taken from the extremes we obtained for the fit-quality metrics. Notably, in one recording the population response shutdown half-way through; yet despite the model being fit only to the 40s window containing strong oscillations, it correctly forecasts the collapse of P10 activity, and its slight rise in firing rate thereafter. Within the limits of this sample of data, the low dimensional trajectory of the periodic attractor appears to encode the motor program.

To confirm this, we asked whether the motor program – as represented by the P10 activity, – was truly low-dimensional. The successful decoding of future P10 activity was achieved despite needing only 3-5 embedding dimensions to account for 80% variance in the population activity for these nine recordings (Supplementary Fig. 8). Increasing the number of embedding dimensions to account for 90% variance, at least doubling the number of embedding dimensions, did not improve the forecasts of P10 activity (Supplementary

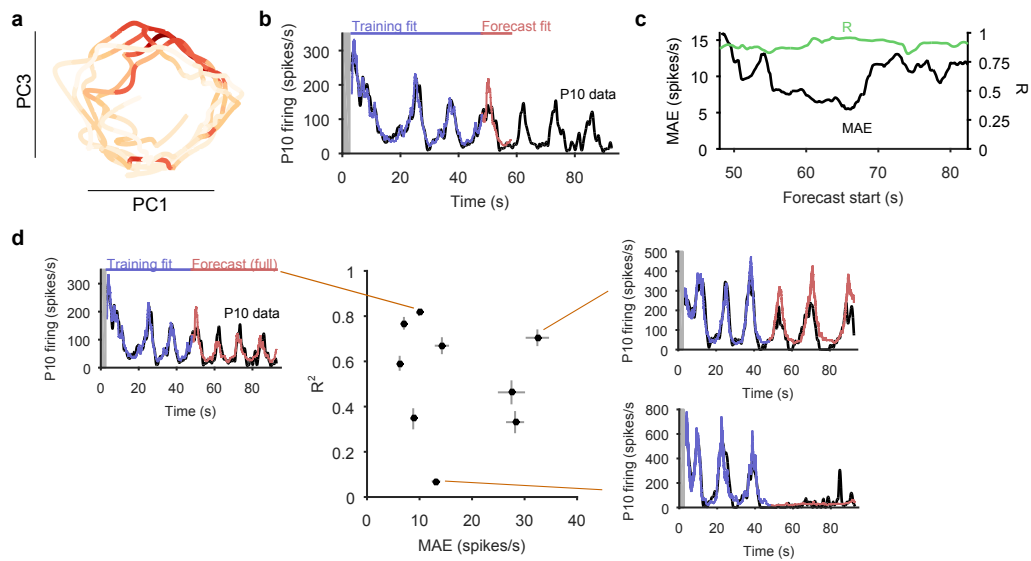


Figure 5: Motor output can be decoded directly from the low-dimensional trajectory of population activity. **a**, An example two-dimensional projection of one population's response trajectory, color-coded by simultaneous P10 firing rate. In this example pair of dimensions, we can see nerve P10 firing 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, quality of the forecast in the 10 s after each fitted 40 s sliding window. Match between model forecast and P10 data quantified by both fit to change (R: correlation coefficient) and scale (MAE: median absolute error) of activity over the forecast window. **d** Summary of model forecasts for all 9 population responses with P10 activity. Dots and lines show means ± 2 SEM 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.

Fig. 8). These results suggest that the low dimensional population trajectory is sufficient to describe the locomotion motor program.

Variable neuron participation in stable motor programs

If the low-dimensional trajectory described by the joint activity of the population just is the locomotion motor program, then how crucial to this program are the firing of individual neurons^{9-12,39}? Having quantified the motor program as the low-dimensional activity trajectory, we could uniquely ask how much each neuron participated in each evoked program. We quantified each neuron's *participation* as the absolute sum of its weights on the principal axes (eigenvectors): large total weights indicate a dominant contribution to the low-dimensional trajectory, and small weights indicate little contribution. So quantified, participation includes both firing rate and synchrony effects on each neuron's contribution (Supplementary Fig. 9).

Every population response had a long-tailed distribution of participation (Fig. 6a), indicating that a minority of neurons dominated the dynamics of any given response. Nonetheless, these neurons were not fixed: many with high participation in one population response showed low participation in another (Fig. 6b,c). To rule out noise effects on the variability of participation (for example, due to the finite duration of recording), we fitted a noise model to the change in participation separately for each preparation (Fig. 6d,e). Every preparation's pedal ganglion contained neurons whose change in participation between responses well-exceeded that predicted by the noise model (Fig. 6f). Consequently, the contribution of single neurons was consistently and strongly variable between population responses in the same preparation.

These data show that a neuron's role within the locomotion motor program is not fixed, but leave open the question of whether single neuron variability causes variation in the program itself. In our analysis, variation between sequentially-evoked population responses is quantified by the distance between their low-dimensional projections (as in Fig. 3c). We found that the distance between a pair of population responses did not correlate with either the total change in neuron participation between the two responses (Fig. 6g) or the distance between their participation distributions (Fig. 6h). The execution of the motor program is thus robust to the participation of individual neurons.

Participation maps identify potential locations of the pattern generator network

To get some insight into the physical substrate of the attractor, we plotted maps of the participation of each neuron in each preparation. We found that neurons with strong participation across the three evoked population responses were robustly located in the caudo-lateral quadrant of the ganglion (Fig. 7a,b). Maps of the right ganglion also indicated strong participation in the rostro-medial quadrant; due to the low numbers of maps for each side, it is unclear whether this is a true asymmetry of the ganglia or simply reflects sampling variation. Neurons with highly variable participation between population responses (Fig. 7c,d) were similarly found in the caudo-lateral quadrants of both ganglia. Strongly participating neurons were thus confined to specific distributed regions of the pedal ganglion's network.

These data are consistent with a network-level distribution of the attractor, with a particularly strong contribution from the caudo-lateral quadrant. Encouragingly, from a different data-set we previously described this region as containing neural ensembles that generated a cyclical packet of neural activity, which moved in phase with activity from the

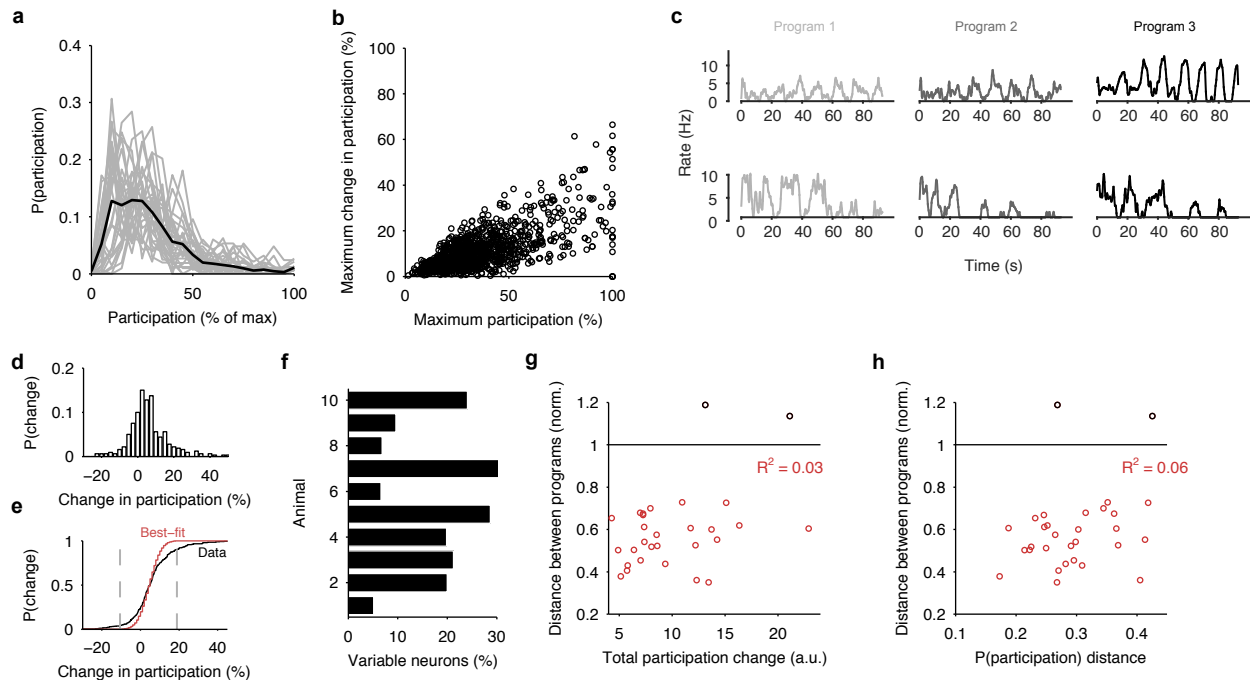


Figure 6: Single neuron participation varies within and between evoked locomotion

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 response, 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 relative to the distance between the random projections, allowing comparison between preparations (Fig. 3c). Black dots are excluded outliers, corresponding to the pairs containing response 1 in preparation 4 with apparent chaotic activity (Supplementary Fig. 4). **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 7: 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 (a,b), but plotting the range of participation across the three evoked population responses.

neck-projecting P10 nerve³². Consequently, both those data and our new data support our hypothesis that the pattern generator for locomotion is predominantly located in the caudo-lateral network.

Discussion

Locomotion networks provide a tractable basis for testing theories of neural dynamics^{6,13,14,32,40–42}, as they couple complex dynamics with clearly defined outputs. We took advantage of this to test the hypothesis that neural circuits form attractor networks. Across all 30 evoked population responses examined here, there was a remarkable heterogeneity of spike-train patterns, from visually evident widespread oscillations to noisy, stuttering oscillations in a minority of neurons. Yet our analysis shows that underpinning this heterogeneity is the same dynamical system: a low-dimensional, decaying, periodic orbit. We found a remarkably consistent periodicity and rate of orbital decay across evoked responses within a preparation and between preparations. The stability of these dynamics, and the convergence of population activity to the same manifold, suggests our recordings have captured for the first time a stable-spiral attractor neuronal network.

Treating a neural circuit as a realisation of a dynamical system takes the emphasis away from the details of individual neurons - their neurotransmitters, their ion channel repertoire - and places it instead on their collective action. This allows us to take a Marr-ian perspective⁴³, which neatly separates the computational, algorithmic, and implementation levels of motor control. The computational problem here is of how to generate rhythmic

locomotion for a finite duration; the algorithmic solution is a decaying periodic attractor; and the implementation of that attractor is the particular configuration of neurons in the pedal ganglion - one of many possible implementations⁴⁴⁻⁴⁷.

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

The wide variation of single neuron participation between evoked bouts of fictive locomotion also raises new challenges for theories of neural network attractors. While a variety of models present solutions for self-sustaining periodic activity in a network of neurons^{44,46,47}, it is unclear if they can account for the variable participation of single neurons. A further challenge is that while the variable participation of individual neurons does not affect the underlying program, clearly it takes a collective change in single neuron activity to transition between behaviours - as, for example, in the transition from galloping to crawling in *Aplysia*. What controls these transitions, and how they are realised by the population dynamics, is yet to be explored either experimentally or theoretically.

Our results nonetheless argue against a number of hypotheses for the implementation of rhythmic locomotion by the pedal ganglion. As noted above, such single neuron variability between sequential locomotion bouts argues against the generation of rhythmic activity by one or more independent neurons that are endogenous oscillators. Our results also argue against the existence of many stable periodic states in this network⁴⁵. Such meta-stability would manifest as changes in periodicity following perturbation. Our results show that spontaneous divergences from the attractor overwhelmingly returned to the same attractor. We observed only three (of 17) that were consistent with moving to a different periodic state; potentially these were transitions from the escape gallop to normal crawling³¹.

How then might the pedal ganglion network implement a periodic attractor? Our data were collected from an isolated central nervous system preparation, in which the modulatory influence of neurons outside the pedal ganglion cannot be discounted⁴⁸. Nonetheless, as the pedal ganglion contains the central pattern generator for locomotion²⁶, we can suggest how that generator is realised. Our results here support the hypothesis that the periodic activity is an emergent property of the ganglion's network. We know the pedal ganglion contains a mix of interneurons and motorneurons²⁸, and that the motorneurons are not synaptically coupled²⁷, suggesting they read-out (and potentially feedback to) the dynamics of an interneuron network. An hypothesis consistent with our results here is that the ganglion contains a recurrent network of excitatory interneurons, centred on the caudo-lateral quadrant, which feed-forward to groups of motorneurons³². This recurrent network embodies the attractor, in that stimulation of the network causes a self-sustained packet of activity to sweep around it³². We see this as the periodic trajectory of joint population activity (cf Fig. 2a, Fig. 3b).

Our data further suggest that the pedal ganglion network supports at least two stable states, the spontaneous activity and the stable-spiral attractor. Reaching the stable-spiral attractor from the spontaneous activity required long-duration, high-voltage pedal nerve stimulation (Fig 1; ref.³²). In dynamical systems terms, this suggests that the spontaneous state's basin of attraction is large: most perturbations return to that state, and it

takes a large perturbation to move into a different basin of attraction. The co-existence of stable spontaneous and periodic states in the same network suggests that something must reconfigure the network to sustain periodic activity; otherwise, irrespective of the stimulation, the network would always return to the spontaneous state. One possibility in the pedal ganglion is that serotonin alters the effective connections between neurons: escape galloping is both dramatically extended by systemic injection of serotonin alongside tail stimulation³⁷, and evoked by stimulating serotonergic command neurons CC9/CC10 in the cerebral ganglion⁴⁸. Future experimental work should thus test the stability of the spontaneous state, and test how manipulating serotonin affects reaching and sustaining the stable-spiral attractor.

Finding and quantifying the attractor required new analytical approaches. We introduce here the idea of using recurrence analysis to solve two problems: how to identify rotations of activity in a high-dimensional space (not necessarily periodic); and how to identify when the recorded system is and is not on the putative underlying attractor. By extracting the times when the population activity is on that attractor, we could then quantify and characterise it, including identifying transient perturbations, and estimating changes in orbital period. Crucially, these extracted attractor times let us further introduce the idea of using linear models as a statistical estimator, to identify the type of attractor, and compare the detected attractor's parameters within and between preparations. Our analysis approach thus offers a road-map for further studies of population dynamics.

The grand challenge of systems neuroscience is to find simplifying principles in the face of overwhelming complexity^{5,49}. We have provided evidence here that neural population activity is a high-dimensional view of the true underlying low-dimensional dynamical system. It follows that dimension reduction approaches are not just useful descriptions of population activity¹⁷, but can reveal the underlying system realised by a neuronal network^{18,32}.

Online Methods

Imaging Full details of the preparation are given in ref.³². Briefly, the cerebral, pleural and pedal ganglia were dissected out, pinned to the bottom of a chamber, and maintained at 15 – 17°C. Imaging of neural activity used the fast voltage sensitive absorbance dye RH-155 (Anaspec), and a 464-element photodiode array (NeuroPDA-III, RedShirtImaging) sampled at 1600 Hz. Optical data from the 464 elements were bandpass filtered in Neuroplex (5 Hz high pass and 100 Hz low pass Butterworth filters), and then spike-sorted with independent component analysis in MATLAB to yield single neuron action potential traces (the independent components), as detailed in⁵⁰. Rhythmic locomotion motor programs were elicited using 8V 5ms monophasic pulses delivered at 20Hz for 2.5s via suction electrode to pedal nerve 9. A separate suction electrode was attached to pedal nerve 10 to continuously monitor the locomotion rhythm³⁸. The stimulation protocol (Fig. 1b) used short (15 mins) and long (60 mins) intervals between stimulations, as the original design also sought effects of sensitisation.

Spike-train analysis Power spectra were computed using multi-taper spectra routines from the Chronux toolbox⁵¹. We computed the power spectrum of each neuron's spike-train post-stimulation, and plot means over all spectra within a recorded population, and the mean over all mean spectra. We computed the spike-density function $f(t)$ for each neuron by convolving each spike with a Gaussian, with $\sigma = \{\text{median ISI in population}\}/\sqrt{12}$,

using a time-step of 10 ms. To visualise the entire population’s spiking activity (Fig. 1c), we cluster neurons by the similarity of their firing patterns using our modular deconstruction toolbox³². Different dynamical types of ensembles were identified by the properties of their autocorrelograms - see³² for details.

Model network We used a three-neuron network to demonstrate the dynamical properties of a periodic attractor as realised by neurons (Supplementary Fig. 1). Each neuron’s membrane dynamics were given by $\tau_a \dot{a}_i = -a_i(t) + c_i(t) + \sum_{j=1}^3 w_{ji} r_j(t) - \gamma y_i(t)$, with adaptation dynamics $\tau_y \dot{y}_i = -y_i(t) + r_i(t)$, and output firing rate $r_i(t) = \max\{0, a_i(t)\}$. Weights $w_{ji} \leq 0$ give the strength of inhibitory connections between the neurons, each of which receives a driving input c_i . This model, due to Matsuoka^{52,53}, generates self-sustained oscillation of network firing rates given constant scalar inputs $c_i(t) = c$, despite each neuron not being an endogenous oscillator. The time constants of membrane τ_a and adaptation τ_y dynamics, together with the strength of adaptation γ , determine the periodicity of the oscillations^{52,53}. Here we use $\tau_a = 0.025$, $\tau_y = 0.2$, and $\gamma = 2$; input was $c_i = 3$ throughout except where noted.

Recurrence analysis Low dimensional projections of the joint population activity were obtained for each program using standard principal components analysis, applied to the covariance matrix of the spike-density functions. The d leading eigenvectors W_i of the covariance matrix define the d principal dimensions, and the d corresponding eigenvalues are proportional to the variance accounted for by each dimension. Projections (the “principal components”) onto each of the chosen dimensions is given by $p_i(t) = \sum_{k=1}^n W_i^k f^k(t)$, where the sum is taken over all n neurons in the analyzed population.

We used recurrence analysis^{33,34} to determine if the low-dimensional projection contained a stable periodic orbit. To do so, we checked if the low-dimensional projection $P(t) = (p_1(t), p_2(t), \dots, p_d(t))$ at time t recurred at some time $t + \delta$ in the future. Recurrence was defined as the first point $P(t + \delta) = (p_1(t + \delta), p_2(t + \delta), \dots, p_d(t + \delta))$ that was less than some Euclidean distance θ from $P(t)$. The recurrence time of point $P(t)$ is thus δ s. Contiguous regions of the projection’s trajectory from $P(t)$ that remained within distance θ were excluded. Threshold θ was chosen based on the distribution of all distances between time-points, so that it was scaled to the activity levels in that particular program. Throughout we use the 10% value of that distribution as θ for robustness to noise; similar periodicity of recurrence was maintained at all tested thresholds from 2% upwards (Supplementary Figure 3).

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

We measured the attractor’s stability as the percentage of all points that were in

periodic orbits. Evolving dynamics of each program were analysed using 5 s sliding windows, advanced in steps of 1 s. We defined the “coalescence” time of the attractor as the mid-point of the first window in which at least 90% of the points on the trajectory were recurrent.

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

To check the robustness of the convergence to the same manifold, we repeated this analysis starting from a common set of principal axes for the three programs, obtained using principal component analysis of their concatenated spike-density functions. We plot the results of this analysis in Supplementary Fig. 4a.

As a further robustness control, we sought evidence of the manifold convergence independent of any low-dimensional projection. We made use of the idea that if neurons are part of sequential programs on a single manifold, then the firing of pairs of neurons should have a similar time-dependence between programs^{22,35}. For each pair of programs (A, B) from the same preparation, we computed the similarity matrix⁵⁴ $S(A)$ between the spike-density functions of all neuron pairs in A , and similarly for B , giving $S(B)$. We then computed the correlation coefficient between $S(A)$ and $S(B)$: if A and B are on the same manifold, so their pairwise correlations should themselves be strongly correlated. As a control we computed a null model where each neuron has same total amount of similarity as in the data, but its pairwise similarity with each neuron is randomly distributed⁵⁴. The expected value of pairwise correlation between neurons i and j under this model is then $E_{ij} = s_i s_j / T$, where (s_i, s_j) are the total similarities for neurons i and j , and T is the total similarity in the data matrix. For comparison, we correlated $S(A)$ with E , and plot these as the control correlations in Fig. 3e.

Testing return to the same manifold after perturbation We detected divergences of the trajectory away from the putative manifold, indicating spontaneous perturbations of population dynamics. We first defined potential perturbations after coalescence as a contiguous set of 5s windows when the density of recurrent points was below 90% and fell below 50% at least once. The window with the lowest recurrence density in this divergent period was labelled the divergent point. We removed all such divergent periods whose divergent point fell within 2 oscillation cycles of the end of the recording, to rule out a fall in recurrence due solely to the finite time horizon of the recording. For the remaining 19 divergent periods, we then determined if the population activity returned to a recurrent

state after the divergent point; that is, whether the density of recurrence returned above 90% or not. The majority (17/19) did, indicating the perturbation returned to a manifold.

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

Statistical estimation of the attractor’s parameters We introduce here a statistical approach to analysing the dynamics of low-dimensional projections of neural activity time-series obtained from experiments. We first fitted a linear model around each point on the low-dimensional trajectory to capture the local dynamics. For each point $P(t)$, we took the time-series of points before and after $P(t)$ that were contiguous in time and within $2.5 \times \theta$ as its local neighbourhood; if less than 100 points met these criteria $P(t)$ was discarded. We then fitted the dynamical model $\dot{P}^* = AP^*$ that described the local evolution of the low-dimensional projection P^* by using linear regression to find the Jacobian matrix A ; to do so, we use the selected local neighbourhood time-series as P^* , and their first-order difference as \dot{P}^* . The maximum eigenvalue $\lambda = a + ib$ of A indicates the dominant local dynamics⁵⁵, whether contracting or expanding (sign of the real part a of the eigenvalue), and whether oscillating or not (existence of the complex part of the eigenvalue i.e. $b \neq 0$). We fitted A to every point $P(t)$ after the stimulation off-set, typically giving ≈ 9000 local estimates of dynamics. The dominant dynamics for the whole program were estimated by averaging over the real a and the complex b parts of the maximum eigenvalues of the models fitted to all recurrent points in the dominant periodic orbit. The linear model’s estimate of the orbit rotation period was estimated from the complex part as $\omega = 2\pi b\Delta t$, with the sampling time-step $\Delta t = 0.01$ s here. The linear model’s estimate of the contraction rate is $\exp(a/\Delta t)$, which we express as a percentage.

Tracking changes in periodicity over a program We tracked changes in the oscillation period by first averaging the recurrence time of all recurrent points in a 5s sliding window. We then correlated the mean time with the time-point of the window to look for sustained changes in the mean period over time, considering only windows between coalescence and the final window with 90% recurrent points. We used a weighted version of Spearman’s rank to weight the correlation in favour of time windows in which the trajectory was most clearly on the periodic orbit, namely those with a high proportion of recurrent points and low variation in recurrence time. The weighted rank correlation is: given vectors x and y of data rankings, and a vector of weights w , compute the weighted mean $m = \sum_i w_i x_i / \sum_i w_i$ and standard deviation $\sigma_{xy} = \sum_i w_i (x_i - m_x)(y_i - m_y) / \sum_i w_i$, 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 s_i is the standard deviation of recurrence times in window i , and Q_i is the proportion of recurrent points in window i . P-values were obtained using a permutation test with 10000 permutations.

Decoding motor output We decoded P10 activity from the low-dimensional trajectory of population activity using a generalised linear model. We first ruled out that any simultaneously recorded neuron was a motoneuron with an axon in nerve P10, by checking if any neurons had a high ratio of locking between their emitted spikes and spikes

occurring at short latency in the P10 recording. Supplementary Fig. 6 shows that no neuron had a consistent, high ratio locking of its spikes with the P10 activity.

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

We tested activity histories between 50 and 200ms duration, with time-steps of 10ms, so that the largest model for a given program had $d \times 20$ coefficients. These short windows were chosen to rule out the contributions of other potential motoneurons in the population recording that would be phase offset from neck contraction (as 200 ms is 2% of the typical period). All results were robust to the choice of history duration, so we plot results using history durations that had the smallest median absolute error in forecasting for that program.

Single neuron participation We quantified each neuron’s participation in the low-dimensional projection as the L1-norm: the absolute sum of its weights on the principal 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 principal axes, $W_j^m(i)$ is the neuron’s weight on the j th axis, and λ_j^m is the axis’ corresponding eigenvalue. Within a program, participation for each neuron was normalised to the maximum participation in that program. To fit a noise model for the variability in participation between programs, we first computed the change in participation for each neuron between all pairs of programs in the same preparation. We then fit a Gaussian model for the noise, using an iterative maximum likelihood approach to identify the likely outliers; here the outliers are the participation changes that are inconsistent with stochastic noise. In this approach, we compute the mean and variance of the Gaussian from the data, eliminate the data-point furthest from the estimate of the mean, re-estimate the mean and variance, and compute the new log likelihood of the Gaussian model without that data-point. We iterate elimination, re-estimation, and likelihood computation until the likelihood decreases. The final model (mean and variance) found before the decrease is then the best-fit Gaussian model to the bulk of the data. Neurons whose maximum change in participation exceeded a threshold of the mean $\pm 3SD$ of that best-fit model were then considered “strongly variable” neurons.

We asked whether the variation in low-dimensional dynamics of sequentially-evoked programs was a consequence of the degree of variation in single neuron participation. Between a pair of consecutively evoked programs, we quantified the variation in their low dimensional dynamics as the Hausdorff distance between them, normalised by the mean distance between their random projections. This normalisation allowed us to put all programs on a single scale measuring the closeness relative to random projections, such that 1 indicates equivalence to a random projection, < 1 indicates closer than random projections, and > 1 indicates further apart than random projections. For a given pair of programs, we quantified the variability of individual neurons’ participation in two

ways: by summing the change in participation of each neuron between the programs; and by computing the Hellinger distance between the two distributions of participation (one distribution per program).

Participation maps Each neuron's (x,y) location on the plane of the photodiode array could be estimated from the weight matrix from the independent component analysis of the original 464 photodiode time-series; see³² for full details. We were able to reconstruct locations for all neurons in 8 of the 10 recorded preparations; for the other two preparations, partial corruption of the original spike-sorting analysis data prevented reconstructions of their neuron locations. We merged all left or right ganglion recordings on to a common template of the photodiode array. The marker sizes and colour codes for each neuron were proportional to the normalised maximum participation of that neuron (Fig. 7a,c) and to the range of normalised maximum participation across the three programs (Fig. 7b,d).

Acknowledgements We thank R. Petersen and B. Mensh for comments on drafts, A. Singh for suggesting the Hausdorff distance, and J. Wang for technical assistance. M.D.H was supported by a Medical Research Council Senior non-Clinical Fellowship. W.F. was supported by NIH R01NS060921 and NSF 1257923. A.B. was supported by NIH F31NS079036.

Author Contributions A.B. performed the experiments and spike-sorting; A.B. & W.F. designed the experiments; M.D.H. designed the analyses with contributions from A.B., and developed and coded the analytical tools; M.D.H analysed the data; all authors discussed the results; M.D.H. wrote the paper with contributions from A.B. & W.F.

References

1. Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal population coding of movement direction. *Science* **233**, 1416–1419 (1986).
2. Getting, P. A. Emerging principles governing the operation of neural networks. *Annu Rev Neurosci* **12**, 185–204 (1989).
3. Ahrens, M. B. *et al.* Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471–477 (2012).
4. Portugues, R., Feierstein, C. E., Engert, F. & Orger, M. B. Whole-brain activity maps reveal stereotyped, distributed networks for visuomotor behavior. *Neuron* **81**, 1328–1343 (2014).
5. Yuste, R. From the neuron doctrine to neural networks. *Nat Rev Neurosci* **16**, 487–497 (2015).
6. Petersen, P. C. & Berg, R. W. Lognormal firing rate distribution reveals prominent fluctuation-driven regime in spinal motor networks. *bioRxiv* (2016). 062257.
7. Chestek, C. A. *et al.* Single-neuron stability during repeated reaching in macaque premotor cortex. *J Neurosci* **27**, 10742–10750 (2007).
8. Hatsopoulos, N. G., Xu, Q. & Amit, Y. Encoding of movement fragments in the motor cortex. *J Neurosci* **27**, 5105–5114 (2007).

9. Carmena, J. M., Lebedev, M. A., Henriquez, C. S. & Nicolelis, M. A. L. Stable ensemble performance with single-neuron variability during reaching movements in primates. *J Neurosci* **25**, 10712–10716 (2005).
10. Hill, E. S., Vasireddi, S. K., Bruno, A. M., Wang, J. & Frost, W. N. Variable neuronal participation in stereotypic motor programs. *PLoS One* **7**, e40579 (2012).
11. Huber, D. *et al.* Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature* **484**, 473–478 (2012).
12. Hill, E. S., Vasireddi, S. K., Wang, J., Bruno, A. M. & Frost, W. N. Memory formation in tritonia via recruitment of variably committed neurons. *Curr Biol* **25**, 2879–2888 (2015).
13. Briggman, K. L., Abarbanel, H. D. I. & Kristan, W., Jr. Optical imaging of neuronal populations during decision-making. *Science* **307**, 896–901 (2005).
14. Levi, R., Varona, P., Arshavsky, Y. I., Rabinovich, M. I. & Selverston, A. I. The role of sensory network dynamics in generating a motor program. *J Neurosci* **25**, 9807–9815 (2005).
15. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387–400 (2010).
16. Churchland, M. M., Cunningham, M. T., John P. and Kaufman, Foster, J. D., Nuyujukian, P. & Ryu, K. V., Stephen I. and Shenoy. Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
17. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* **17**, 1500–1509 (2014).
18. Kato, S. *et al.* Global brain dynamics embed the motor command sequence of caenorhabditis elegans. *Cell* **163**, 656–669 (2015).
19. Brody, C. D., Romo, R. & Kepecs, A. Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Curr Opin Neurobiol* **13**, 204–211 (2003).
20. Miller, P. Dynamical systems, attractors, and neural circuits. *F1000Research* **5**, 992 (2016).
21. Seelig, J. D. & Jayaraman, V. Neural dynamics for landmark orientation and angular path integration. *Nature* **521**, 186–191 (2015).
22. Peyrache, A., Lacroix, M. M., Petersen, P. C. & Buzski, G. Internally organized mechanisms of the head direction sense. *Nat Neurosci* **18**, 569–575 (2015).
23. Mazor, O. & Laurent, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
24. Bartho, P., Curto, C., Luczak, A., Marguet, S. L. & Harris, K. D. Population coding of tone stimuli in auditory cortex: dynamic rate vector analysis. *Eur J Neurosci* **30**, 1767–1778 (2009).

25. Jahan-Parwar, B. & Fredman, S. M. Neural control of locomotion in aplysia: role of the central ganglia. *Behav Neural Biol* **27**, 39–58 (1979).
26. Jahan-Parwar, B. & Fredman, S. M. Motor program for pedal waves during aplysia locomotion is generated in the pedal ganglia. *Brain Res Bull* **5**, 169–177 (1980).
27. Hening, W. A., Walters, E. T., Carew, T. J. & Kandel, E. R. Motorneuronal control of locomotion in aplysia. *Brain Res* **179**, 231–253 (1979).
28. Fredman, S. M. & Jahan-Parwar, B. Role of pedal ganglia motor neurons in pedal wave generation in aplysia. *Brain Res Bull* **5**, 179–193 (1980).
29. Hall, J. D. & Lloyd, P. E. Involvement of pedal peptide in locomotion in aplysia: modulation of foot muscle contractions. *J Neurobiol* **21**, 858–868 (1990).
30. McPherson, D. R. & Blankenship, J. E. Neuronal modulation of foot and body-wall contractions in *Aplysia californica*. *J Neurophysiol* **67**, 23–28 (1992).
31. Flinn, J. M., Gochman, P., Wanschura, P. & Chandhoke, V. The effect of dopamine receptor blockade on motor behavior in *Aplysia californica*. *Pharmacol Biochem Behav* **69**, 425–430 (2001).
32. Bruno, A. M., Frost, W. N. & Humphries, M. D. Modular deconstruction reveals the dynamical and physical building blocks of a locomotion motor program. *Neuron* **86**, 304–318 (2015).
33. Lathrop, D. P. & Kostelich, E. J. Characterization of an experimental strange attractor by periodic orbits. *Phys Rev A* **40**, 4028–4031 (1989).
34. Marwan, N., Romano, M. C., Thiel, M. & Kurths, J. Recurrence plots for the analysis of complex systems. *Physics Reports* **438**, 237–329 (2007).
35. Yoon, K. *et al.* Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat Neurosci* **16**, 1077–1084 (2013).
36. Walters, E. T. *Sensitization and classical conditioning in Aplysia: Behavioral and neuronal studies*. Ph.D. thesis, Columbia University, New York (1980).
37. Marinesco, S., Wickremasinghe, N., Kolkman, K. E. & Carew, T. J. Serotonergic modulation in *Aplysia*. II. Cellular and behavioral consequences of increased serotonergic tone. *J Neurophysiol* **92**, 2487–2496 (2004).
38. Xin, Y., Weiss, K. R. & Kupfermann, I. An identified interneuron contributes to aspects of six different behaviors in *Aplysia*. *J Neurosci* **16**, 5266–5279 (1996).
39. Katz, P. S., Sakurai, A., Clemens, S. & Davis, D. Cycle period of a network oscillator is independent of membrane potential and spiking activity in individual central pattern generator neurons. *J Neurophysiol* **92**, 1904–1917 (2004).
40. Lewis, J. E. & Kristan, W., Jr. A neuronal network for computing population vectors in the leech. *Nature* **391**, 76–79 (1998).
41. Briggman, K. L. & Kristan, W. B., Jr. Imaging dedicated and multifunctional neural circuits generating distinct behaviors. *J Neurosci* **26**, 10925–10933 (2006).

42. Berg, R. W., Alaburda, A. & Hounsgaard, J. Balanced inhibition and excitation drive spike activity in spinal half-centers. *Science* **315**, 390–393 (2007).
43. Marr, D. *Vision* (W. H. Freeman, San Francisco, 1982).
44. Kleinfeld, D. & Sompolinsky, H. Associative neural network model for the generation of temporal patterns. Theory and application to central pattern generators. *Biophys J* **54**, 1039–1051 (1988).
45. Pasemann, F. Characterization of periodic attractors in neural ring networks. *Neural Networks* **8**, 421 – 429 (1995).
46. Eliasmith, C. A unified approach to building and controlling spiking attractor networks. *Neural Comput* **17**, 1276–1314 (2005).
47. Rokni, U. & Sompolinsky, H. How the brain generates movement. *Neural Comput* **24**, 289–331 (2012).
48. Jing, J., Vilim, F. S., Cropper, E. C. & Weiss, K. R. Neural analog of arousal: persistent conditional activation of a feeding modulator by serotonergic initiators of locomotion. *J Neurosci* **28**, 12349–12361 (2008).
49. Koch, C. Systems biology. modular biological complexity. *Science* **337**, 531–532 (2012).
50. Hill, E. S., Moore-Kochlacs, C., Vasireddi, S. K., Sejnowski, T. J. & Frost, W. N. Validation of independent component analysis for rapid spike sorting of optical recording data. *J Neurophysiol* **104**, 3721–3731 (2010).
51. Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S. & Mitra, P. P. Chronux: a platform for analyzing neural signals. *J Neurosci Methods* **192**, 146–151 (2010).
52. Matsuoka, K. Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biol Cybern* **52**, 367–376 (1985).
53. Matsuoka, K. Mechanisms of frequency and pattern control in the neural rhythm generators. *Biol Cybern* **56**, 345–353 (1987).
54. Humphries, M. D. Spike-train communities: finding groups of similar spike trains. *J Neurosci* **31**, 2321–2336 (2011).
55. Strogatz, S. H. *Nonlinear Dynamics and Chaos* (Westview Press, Cambridge, MA, 1994).

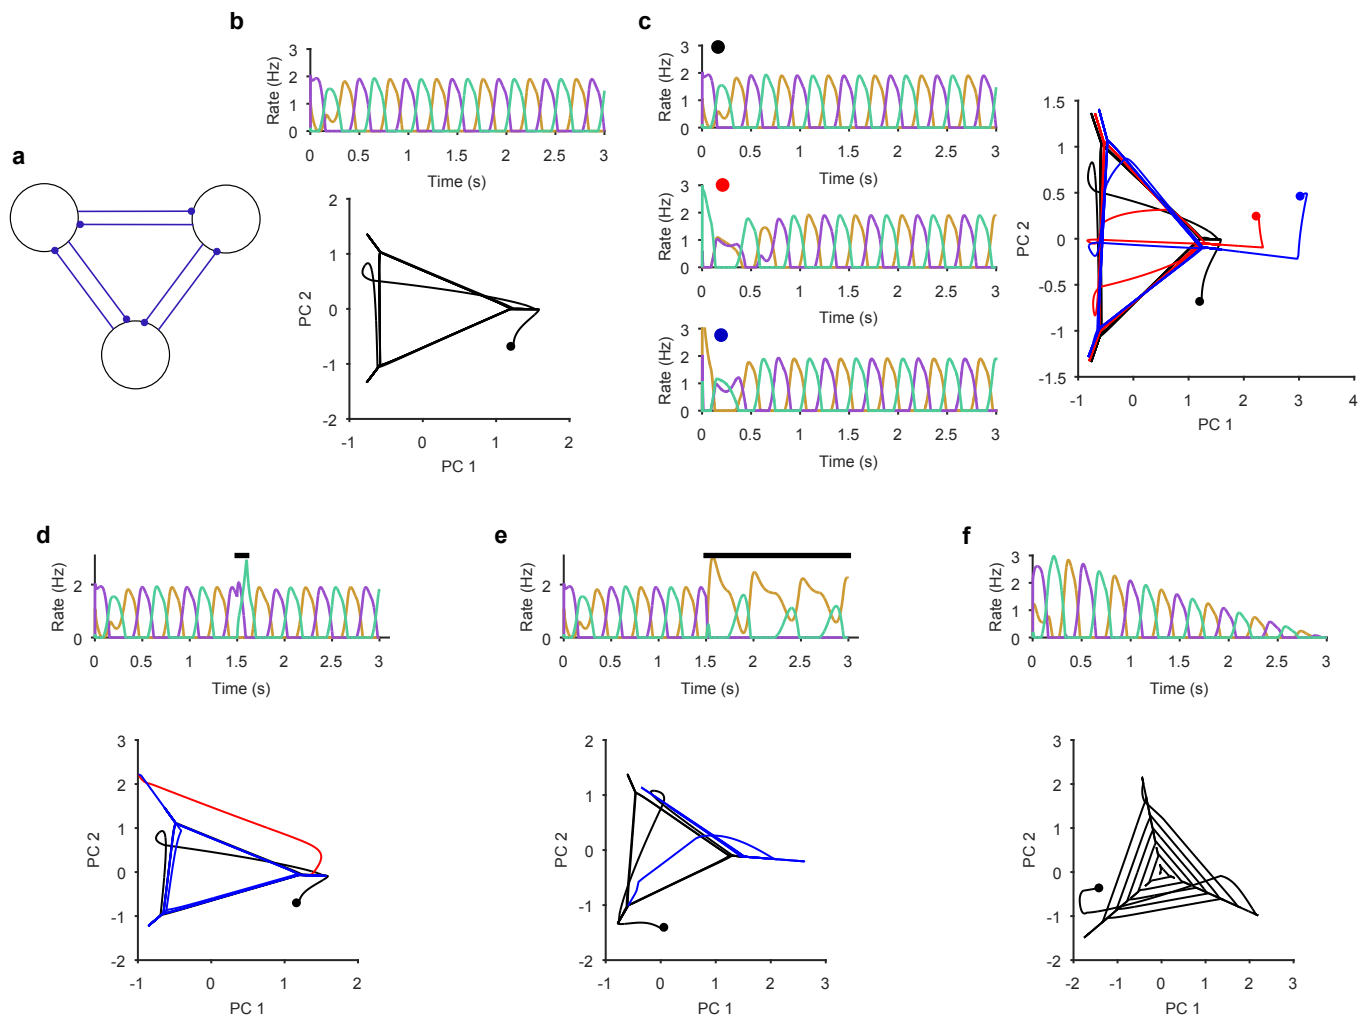


Figure S1: Necessary and sufficient conditions of a periodic attractor. Here we demonstrate the essential properties of a periodic attractor network. **a** A simple neuronal implementation of a periodic attractor network using three mutually inhibitory neurons. Model neurons are not endogenous bursters: oscillatory activity arises from the network. **b** Underlying periodic attractor is revealed by the recurrence of the neurons' joint activity. Top: given sustained input, the three neurons rapidly settle into an alternating oscillatory rhythm. Bottom: projection of the three neurons' firing rates onto the first two principal components (PCs). Starting from time zero (black dot), the joint activity of the network rapidly converges to a recurrent trajectory, where each corner of the triangle - extending into the short diagonal line - corresponds to the maximum rate of one neuron. The space containing the repeated trajectory is the attractor's manifold. **c** Convergence to the manifold from different initial conditions. From three different initial conditions for the neurons (coloured dots), each corresponding joint activity converges to the attractor manifold. **d** Transient perturbations cause a divergence and return to the attractor's manifold. Top: A brief increase in input to one neuron (black bar) perturbs the ongoing oscillation, which is rapidly recovered. Bottom: in the PC projection, the initial recurrent trajectory (in black) is moved away from the manifold by the perturbation (red); removing the perturbing input rapidly moves the joint activity back to the manifold (blue). **e** Sustained perturbations move the network to a new attractor. Top: A sustained change in input (black bar) causes one neuron (purple) to become silent and others to change their oscillation patterns. Bottom: in the PC projection, the perturbation causes the joint activity of the network to move from the initial stable manifold (black) to a new stable manifold (blue: note this is a repeating line). **f** Decaying input causes a spiral. Top: linearly decaying input to all neurons changes the amplitude, but not the frequency of the oscillatory activity. Bottom: in the PC projection, this decay is shown by the spiral of the joint activity's trajectory to the stable point at (0,0) - which corresponds to no activity.

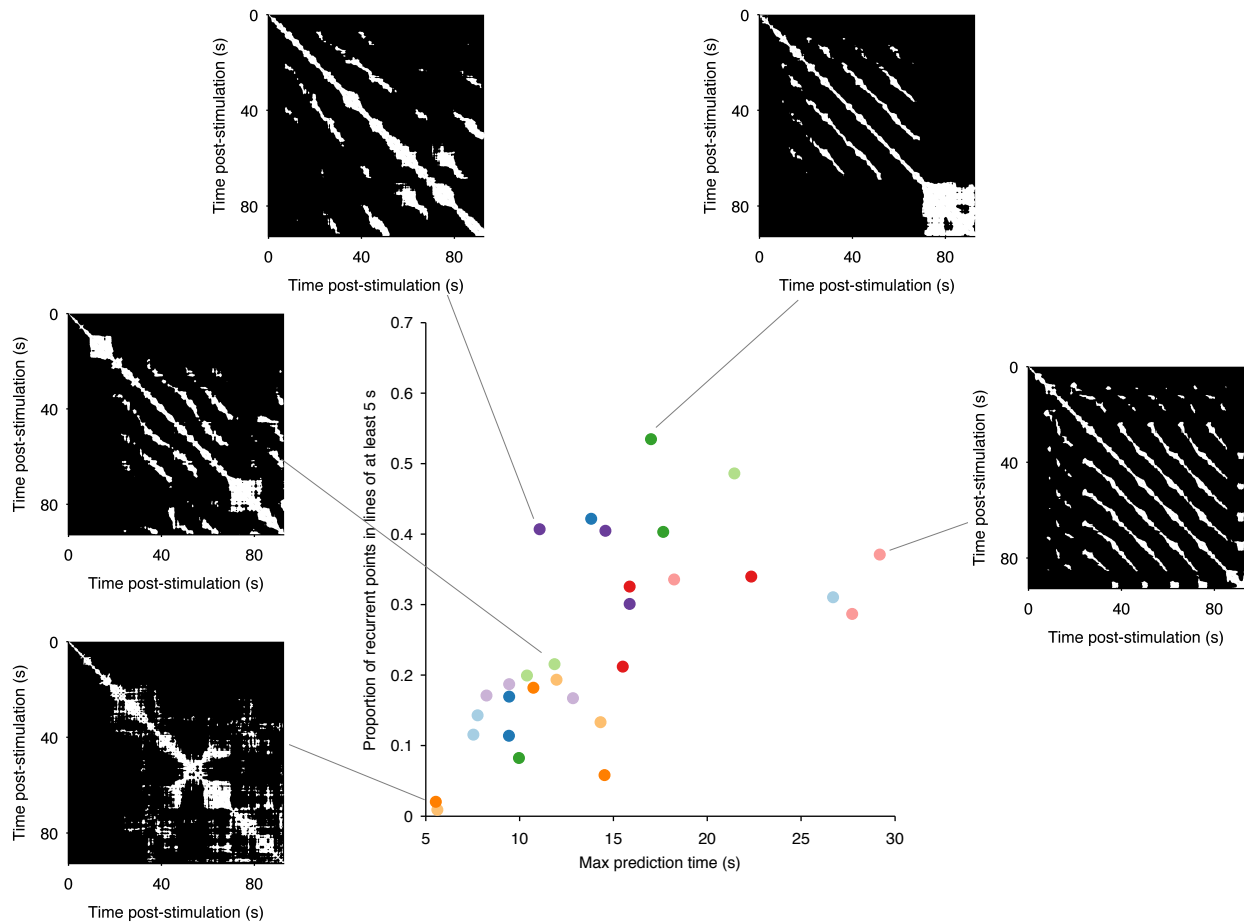


Figure S2: Range of dynamics in the evoked locomotion programs. Here we quantify the complexity of the dynamics in each evoked program using the structure of the recurrence plots (5 example plots are shown). Each white dot is a pair of time-points at which the low-dimensional trajectory recurred. A diagonal line indicates a contiguous period of recurrence of the population dynamics. We thus quantified the complexity of each program by the duration of the longest diagonal line (measuring the predictability of the system) and by the proportion of recurrent points falling in long (> 5 s) diagonal lines (measuring the cleanness of the system's dynamics). We plot the joint distribution of these two quantities here, one point per program; colours indicate programs from the same preparation. These show how the range of dynamics captured by our dataset stretches from quasi-noise (bottom left), though stuttering/pausing oscillations (top left), to continuous oscillation (right). Despite this heterogeneity, almost all evoked programs revealed the same underlying dynamical system (Figs. 2-4, main text)

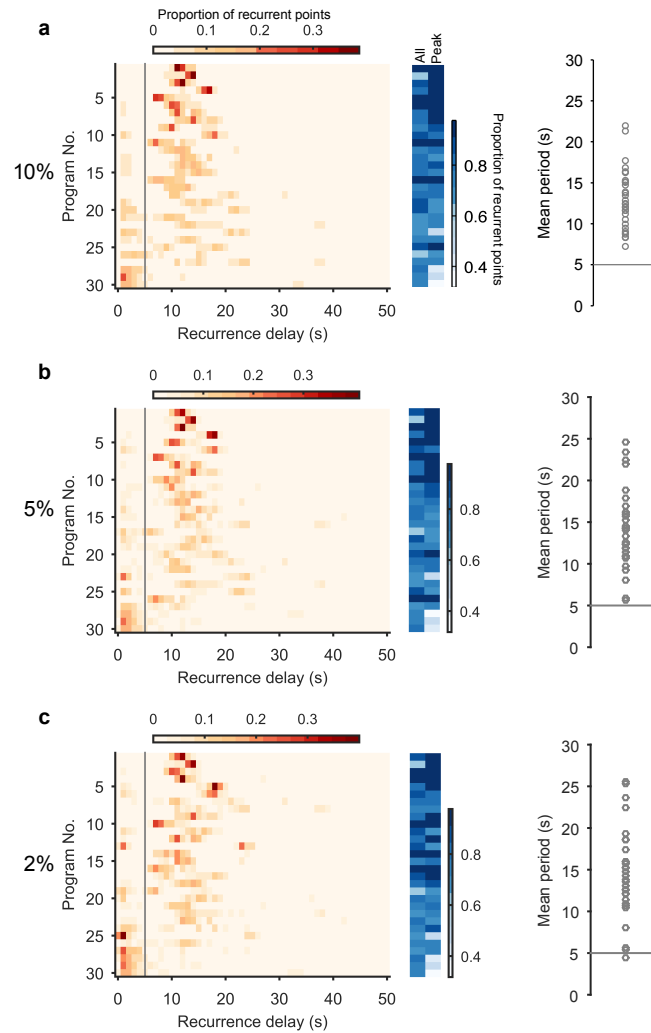


Figure S3: Robustness of the periodic orbits. Here we show that the periodic orbits, defined by recurrence, are robust to the choice of threshold for defining “close”. We defined the threshold as a percentile of all distances between all points in the post-stimulation trajectory. We show here the effects of choosing a threshold of 10% (a), 5% (b), and 2% (c). Left column: histograms of all recurrence times in each program, 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 program. Middle column: density of time-points that were recurrent, and density of recurrence points with times in the dominant period. Right column: Periodic orbit of each program, estimated as the mean recurrence time from the dominant period. We use a threshold of 10% in the main text as it allows a denser sampling of potential trajectory points that are noisy realisations of the underlying attractor.

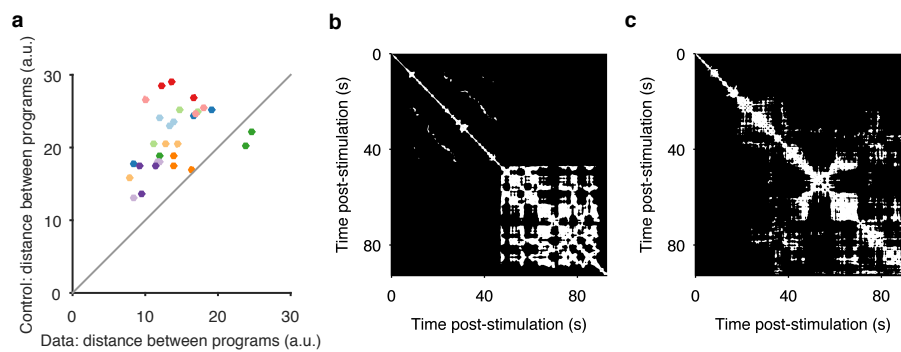


Figure S4: Convergence and non-convergence to the same manifold. **a** Robustness of convergence to the same manifold. Here we repeat the analysis of Figure 3c, but using different low-dimensional projections. In the main text Figures 2-4, we analyse low-dimensional projections of each individual program; the comparison of distances between sequentially evoked programs was then done by projecting all three programs onto the axes of the first program. Here we derive a common set of axes, by first applying PCA to the concatenated time-series of the three programs. We then project each program onto those common axes, before computing distances between data and control (random projection) pairs of programs (dots). We obtain the same results: sequentially-evoked programs converge to the same manifold. **b** Apparent chaotic motion in the non-convergent program. Here we show the recurrence plot for program 1 of preparation 4, the only program that was further from the other programs in the same preparation than random projections. We see that this program seemed to reach an apparently chaotic attractor about 45s after stimulation: quasi-periodic activity, with no regular structure. **c** Apparent random-walk in the non-correlated program. Here we show the recurrence plot for program 3 of preparation 7, the only program that did not recapitulate the pair-wise correlations of neural activity of the other programs in the same preparation. We see that this program had no clear periodic activity, but seemingly random periods of locally-similar activity, consistent with a dynamical system undergoing a random walk.

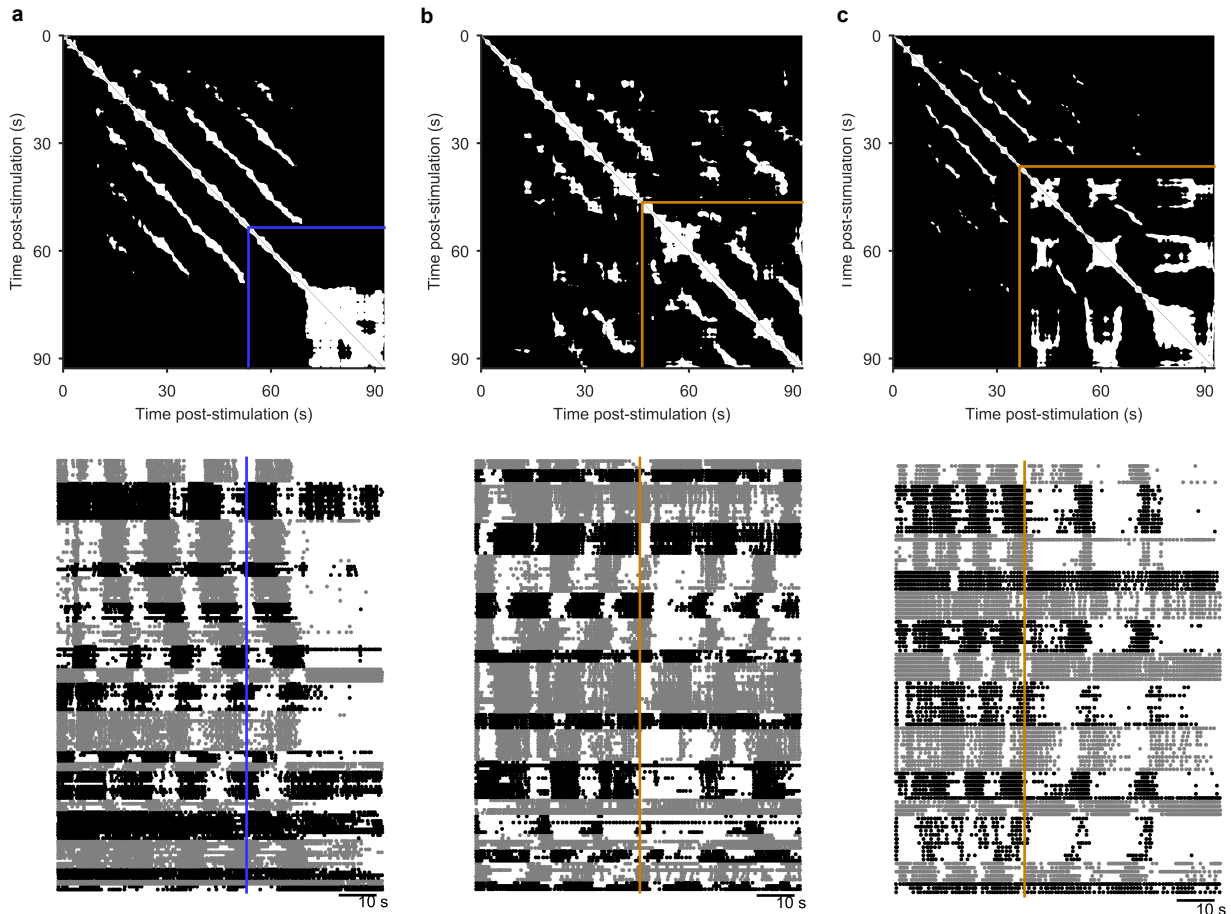


Figure S5: Spontaneous perturbations of ongoing programs. We plot here examples of each type of spontaneous perturbation we observed, visualised using the recurrence plot (top), and the raster plot of population activity (bottom) sorted into detected ensembles. **a** Spontaneous perturbation away from the periodic attractor. Blue line indicates the final 5s window with more than 50% recurrence, long before the end of the recording. **b** Spontaneous perturbation returning to the same manifold of the periodic attractor. The orange lines indicates the point of maximum divergence from the manifold. The recurrence plot shows that same trajectory was recovered thereafter (white patches in upper right corner), visible as the return of the phasically firing ensembles (raster plot). **c** Spontaneous perturbation moving to a new periodic attractor. The recurrent trajectory up to 30s post-stimulus did not recur thereafter (black region in upper right of the recurrence plot), but was replaced with a new recurrent trajectory. The raster plot shows this was an apparent transition in the oscillation period, and/or firing pattern of the ensembles.

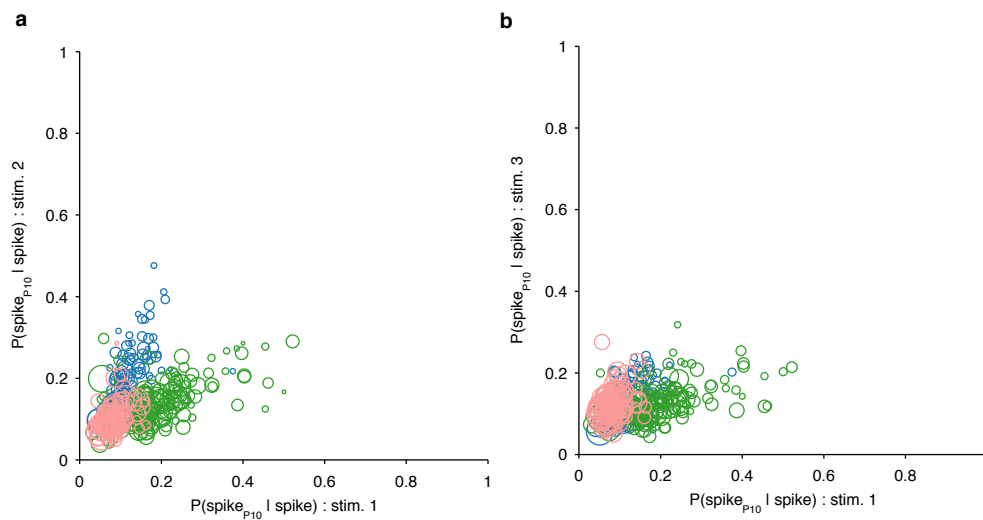


Figure S6: Ruling out P10 motorneurons in the recorded population. We aimed to decode the spiking activity of the motorneuron axons in the P10 nerve directly from the low-dimensional trajectory of population activity in the pedal ganglion (Fig. 5). If any P10 motorneurons had been captured in our population recording, these could have made a disproportionate contribution to the decoding. We thus first sought to detect them, and remove them if necessary. We reasoned that any putative P10 motorneuron should have close to a 1:1 locking between its spikes in the population activity recordings and corresponding spikes in the P10 nerve recording; in other words $P(\text{spike}_{P10}|\text{spike}) \approx 1$. Moreover, we expect a putative P10 motorneuron to show $\sim 1:1$ locking consistently across all three programs from the same preparation. For each neuron in each recorded program with simultaneous P10 activity, we computed $P(\text{spike}_{P10}|\text{spike})$ separately for a range of delays (0 to 25ms) between neuron and P10 spikes, allowing for ± 2 ms jitter at each delay. We plot the maximum probability for each neuron in programs 1 and 2 (**a**), and programs 1 and 3 (**b**); symbol size is proportional to the mean number of spikes for that neuron in the two programs (as few spikes would indicate an unlikely motorneuron). Colours indicate neurons from the three different preparations. We see that no neuron in our dataset comes close to $P(\text{spike}_{P10}|\text{spike}) \approx 1$ at any delay, and there are no neurons with consistently high locking between programs. We concluded that there are no P10 motorneurons in our population recordings.

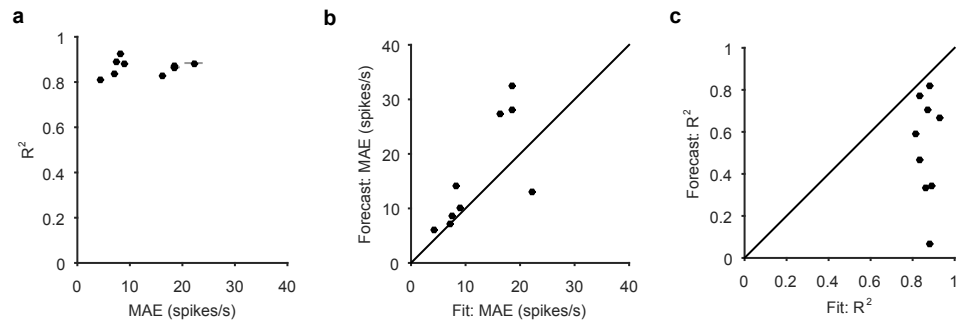


Figure S7: Training fits of the P10 decoding model. **a** Summary of model fits for all 9 programs with P10 activity. Dots and lines show means ± 2 SEM, taken over all 40s windows used to fit the GLMs ($N = 173$). MAE: median absolute error of fit between P10 activity and best-fit model. R^2 : correlation between P10 activity and best-fit model. **b** Relationship between mean model fit to and forecast of the scale of activity during the data-window (fit: 40s; forecast: next 10s). Success in fitting the scale of activity was predictive of forecasting. **c** Relationship between mean model fit to and forecast of the change of activity (measured by R^2). Success in fitting the change in activity was not always strongly predictive of forecasting the change.

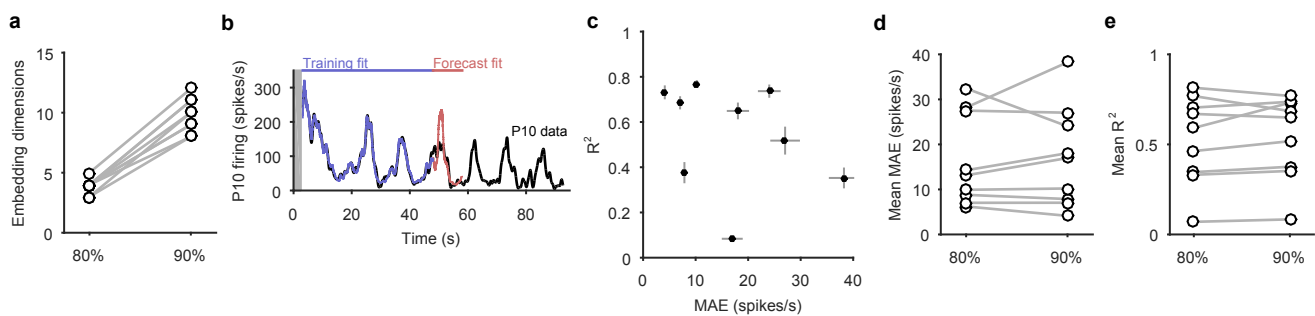


Figure S8: Increasing dimensionality of state-space did not improve the P10 decoding model. Here we show the effects of increasing the embedding dimensions for the population activity before using the decoding model for P10 activity. We used enough dimensions to capture at least 90% of the variance; we used 80% in the main text. **a** Comparison of the number of embedding dimensions needed to explain 80% and 90% variance for the nine recordings with simultaneous P10 activity. A 10% increase in variance needed at least double the number of dimensions. **b** Example fit and forecast by the statistical decoding model for P10 firing rate when using more embedding dimensions. Grey bar indicates stimulation time. **c** Summary of model forecasts for all 9 population responses with P10 activity, when using more embedding dimensions. Dots and lines show means ± 2 SEM over all forecast windows ($N = 173$). MAE: median absolute error (error in scale) of forecast; R^2 : variance explained of forecast. **d** Comparison of the mean MAE in forecasting P10 activity when using embedding dimensions explaining 80% and 90% variance. Increasing the number of dimensions did not systematically improve the forecasting of the scale of the P10 activity. **e** Comparison of the mean R^2 in forecasting P10 activity when using embedding dimensions explaining 80% and 90% variance. Increasing the number of dimensions did not systematically improve the forecasting of the variance of the P10 activity.

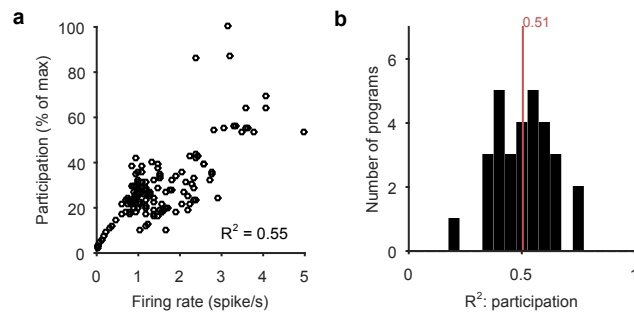


Figure S9: Participation measures rate and synchrony. **a** Example correlation between the evoked firing rates of all neurons in one program and the participation of those neurons in all three programs evoked from that preparation. **b** Histogram of firing rate and participation correlations across all 30 programs. Firing rate explains on average (R^2 , red line) 50% of the variation in participation within a preparation; the remaining variation in participation can be accounted for by differences in synchrony between individual neurons and the low-dimensional projection.