Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Clustering and Degree Homogeneity Govern Reservoir Computing Performance

Ken-ichi Kitayama[1]

kitayama@ieee.org

National Institute of Communications and Information

Atsushi Uchida[3]

Saitama University

Article

Keywords:

Posted Date: August 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10670082/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Clustering and Degree Homogeneity Govern Reservoir Computing Performance

Ken-ichi Kitayama[1,2]* and Atsushi Uchida[3]

[1] National Institute of Information and Communications Technology, Tokyo 184-8795, Japan
[2] Hamamatsu Photonics K.K., Hamamatsu 434-8601, Japan
[3] Department of Information and Computer Sciences, Saitama University, 255 Shimo-okubo, Sakura-ku, Saitama City, Saitama 338-8570, Japan
* Corresponding author. email: kitayama@ieee.org

Abstract: Reservoir computing exploits fixed recurrent networks as computational substrates, and the choice of network architecture shapes performance. Because biological brains are sparse, clustered, and small-world networked, connectome-derived reservoirs are often expected to be computationally favorable. However, its computational capability has not yet been fully explored. Here we examine this expectation by benchmarking reservoirs constructed from a human structural connectome against Watts–Strogatz's small-world networks and random networks, matched for size and mean degree. The benchmark tests are two standard tasks: human-activity classification and Mackey–Glass chaotic signal prediction. On both tasks the small-world reservoir attains the highest performance, while the connectome performs comparably to a random network on prediction. Sweeping the small-world rewiring probability, we find that performance on both tasks is associated with high clustering and low degree heterogeneity and is largely insensitive to characteristic path length, with an optimum in the small-world regime. The connectome, characterized by relatively low clustering and high degree heterogeneity, sits away from this optimum, which is occupied by the small-world networks. These findings identify network properties associated with reservoir performance and indicate that connectome-based connectivity is not, in itself, advantageous for the tasks examined.

Reservoir computing (RC) is a unified computational framework[1,2] that emerged from two independently proposed recurrent neural network models, the echo state network[3,4] and the liquid state machine[5,6]. An input signal drives a fixed, randomly connected recurrent network, called the reservoir, and only a linear readout is trained. Because the recurrent weights are left untrained, the reservoir's architecture is not a learned artifact but a design choice, and a substantial body of work has asked how the topology of the reservoir shapes its computational capacity. Although reservoirs are conventionally wired at random, this reflects convenience rather than principle. Deterministic minimal topologies such as a simple cycle or a delay line with feedback perform comparably to random reservoirs while removing the randomness from their construction[7]; reservoirs with scale-free, highly clustered connectivity retain the echo state property over a wider range of spectral radius than random ones[8], and small-world wiring has been reported to enhance both the echo state property and the propagation of signals between input and readout nodes[9]. This makes RC an unusually direct setting in which to study the relationship between network structure and function.

Networks with small-world organization, that is, high clustering combined with short characteristic path length [10], are of particular interest in this context. Small-world connectivity is a hallmark of many real systems[11], and previous work has shown that reservoirs synthesized as Watts–Strogatz networks can outperform conventional random reservoirs on classification and prediction tasks, with an optimum in the small-world regime [12]. These results suggest that intermediate levels of disorder, rather than fully regular or fully random wiring, provide a favorable substrate for reservoir computation.

A natural extension of this idea is to move from synthetic small-world networks to the biological networks that motivated them. This step has recently been made practical by conn2res, an open-source

toolbox in which an empirical connectivity matrix supplies the fixed recurrent layer of a reservoir while the local dynamics, the input and readout nodes, and the task are specified independently [13]. Applied to a human connectome, this framework has been used to show that the memory capacity of an empirical brain network exceeds that of rewired null networks matched for density and degree sequence, and that the same holds for the connectomes of several model organisms; performance also varies systematically with which cortical system is read out, indicating that computational capacity is distributed unevenly across the network [14]. Such comparisons rest on the detailed maps of structural brain connectivity that connectomics now provides [15]. Because brains are sparse, clustered, small-world networks shaped by evolutionary and developmental pressures, it is an appealing hypothesis that connectome-derived reservoirs might be computationally advantageous relative to generic architectures. Whether this is so, however, has not been examined under conditions that control for the basic network statistics (size, density, and mean degree) that independently affect reservoir performance, nor against a synthetic small-world reservoir, the architecture those statistics suggest should be favorable. This is an intriguing problem and remains to be tested against benchmark tasks. The broader question of how macroscale connectivity relates to function, and of the models used to bridge the two, has been reviewed elsewhere[19]. A parallel line of work has built echo-state networks directly from mammalian connectomes and used them to ask which features of biological wiring support computation on memory tasks[20].

In this article, we test this hypothesis by benchmarking reservoirs built from a human structural connectome against Watts–Strogatz small-world networks and random networks on two standard tasks, classification of human-activity[16] and prediction of Mackey–Glass chaotic-signal[17]. All reservoirs are matched for size and mean degree, so that differences in performance reflect the arrangement of connections rather than their number. To identify the structural properties responsible for the observed differences, we further sweep the small-world rewiring probability and relate performance to clustering, degree heterogeneity and characteristic path length. We find that the small-world reservoir yields the highest performance on both tasks, and that reservoir performance is governed principally by clustering and degree homogeneity, properties on which the connectome is not optimally configured.

Results

We compared three families of reservoir: one derived from a human structural connectome [15], one synthesized as a Watts–Strogatz small-world network[10], and one generated as a random network. To ensure that any difference in computational performance reflected the arrangement of connections rather than their number, the three families were placed on an equal footing: they were matched for the number of nodes N and for the mean degree $\langle k \rangle = 2E / N$, where E is the number of edges, which was held at 8 throughout. At a fixed N this is equivalent to matching edge density, the fraction of all possible node pairs that are connected, $2E / [N(N - 1)]$, since density $= \langle k \rangle / (N - 1)$; because the comparison spans a range of sizes it is the mean degree that is held constant and the density that varies with N, from 0.123 at N = 66 to 0.020 at N = 400. All reservoirs were binary and symmetric (Table 1; Fig. 1). For each condition we generated 20 independent reservoir realizations and report the mean ± standard deviation (s.d.) across them. Reservoirs were evaluated on two standard benchmarks: six-class human-activity classification [16] and one-step prediction of the Mackey–Glass chaotic time series [17]. Throughout, d denotes Cohen's d, the difference between two condition means divided by their pooled standard deviation, so that larger values indicate greater separation, and r is the Pearson correlation coefficient, which measures how strongly two quantities co-vary and ranges from −1 to +1. For the prediction task these quantities are computed on the base-10 logarithm of the mean-squared error (MSE), on which the distributions are approximately normal.

Table 1 Differences of the properties in network topologies between Connectome, Small-world, and random.

	Property	Connectome	Small-World	Random
Differences	Clustering	Low	High at p =0.01-0.10 0	Low
	Degree heterogeneity	High	Low	Moderate
	Average path length	3.17	3.22	3.11
	Hubs	Present	Absent	Absent

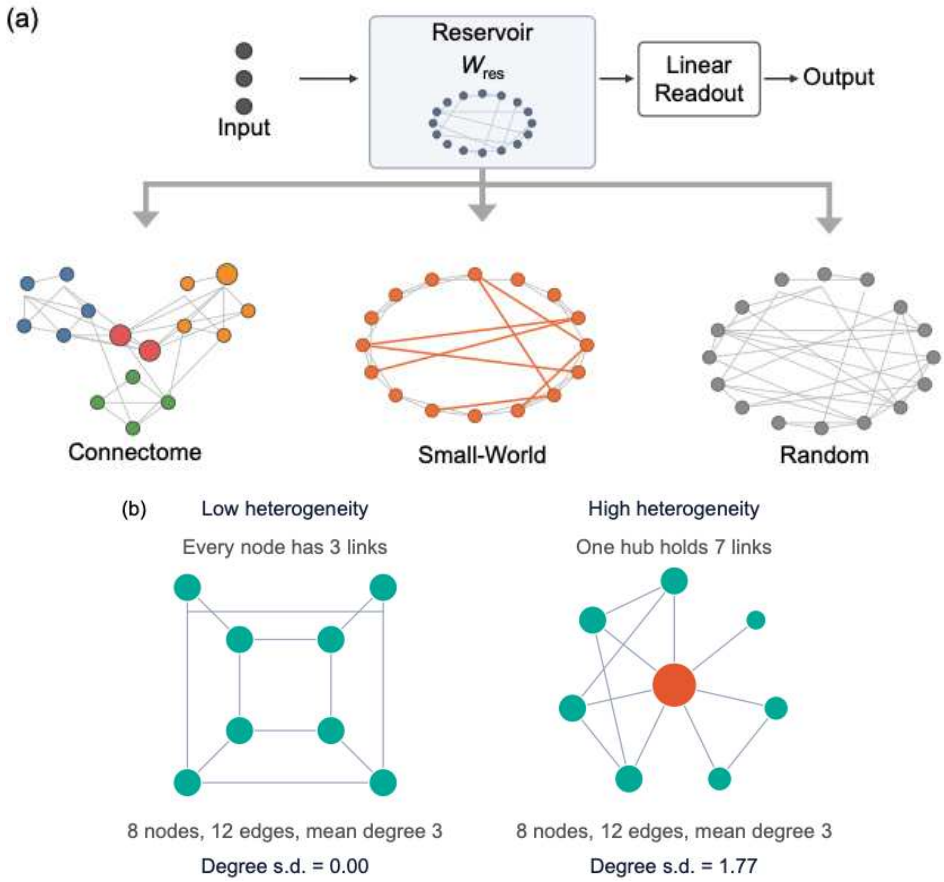


Fig. 1 The three reservoir architectures and heterogeneity of topology**.** (a) Representative examples of the connectome, small-world and random reservoirs compared in this study. All three are undirected and binary, and are matched at each size for the number of nodes and for a mean degree of 8, so that they differ only in how their connections are arranged. (b) 8-node and 12-link low-heterogeneity with degree s.d.=0.00 contrasts with high heterogeneity with degree s.d.=1.77. Every node has 3 links on the left, and one hub holds 7 links on the right.

Benchmark test of classification. The data set of the human activity16 includes six motions: walking, walking upstairs, walking downstairs, sitting, standing, and laying down. Body acceleration was captured along the x-, y- and z-axes and segmented into windows of 128 time steps per axis. The data sets of the training and the testing are 7352 and 2947, respectively. Twenty independent reservoirs were generated and evaluated for each architecture (see Methods). Across reservoir sizes spanning N = 66 to N = 400, the small-world reservoir achieved the highest classification accuracy (Fig. 2). Here N = 66 is the native resolution of the human connectome, the 66 cortical regions of the 15 parcellation, and larger reservoirs (up to N = 400) were obtained by expanding the connectivity matrix (see Methods), so that all three architectures could be compared at a common size. The results of six-class human-activity recognition at N = 400 are summarized in Table 2. The values of mean and the standard deviation (s.d.) of the degree distribution are taken for 20 reservoir realizations. All three architectures have a mean degree of 8. C is the mean clustering coefficient. The effect size d is Cohen's d, the difference between two condition means divided by their pooled standard deviation, here for the comparison against the small-world reservoir. At N = 400 the small-world reservoir reached 69.24 ± 1.05 %, exceeding the connectome (67.29 ± 1.15 %; d = 1.77) and the random network (66.28 ± 1.69 %; d = 2.10). The connectome trails the small-world reservoir by 1.95 %, while the connectome itself exceeds the random network by 1.01 %. The small-world reservoir was the highest at every size examined, N=66, 100, and 400, so its advantage held across a sixfold range of network size. The ordering of the other two architectures was not consistent. At the native resolution of N = 66 where no expansion is applied, the random network was the better than connectome, so that the

connectome ranked last. Note that the training and test partitions come from different groups of volunteers, so the accuracies reported here are lower than those obtained when the same volunteers appear in both.

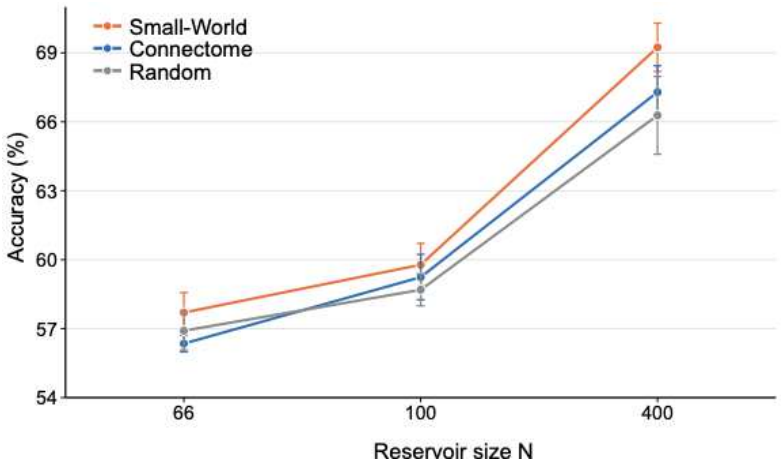


Fig.2 Classification accuracy across reservoir sizes. Performance on the six-class human-activity recognition task for the three reservoir architectures at N = 66, 100 and 400. Points show the mean over 20 independent reservoir realizations and error bars one standard deviation. The small-world reservoir is highest at every size; at N = 400 it reaches 69.24 ± 1.05 %, against 67.29 ± 1.15 % for the connectome and 66.28 ± 1.69 % for the random network. N = 66 is the native resolution of the connectome, and the larger sizes were obtained by the expansion procedure described in the Methods.

Table 2 Three-way comparison on the classification task.

Architecture	Accuracy (%)	C	Degree s.d.	*d*
Small-world (p = 0.5)	69.24 ± 1.05	0.096	1.73	—
Connectome	67.29 ± 1.15	0.049	3.59	1.77
Random	66.28 ± 1.69	0.020	2.79	2.10

To test whether this result depends on a particular choice of the rewiring probability p, we swept p across the small-world regime at N = 400 (Fig. 3). The classification accuracy varied smoothly with p, peaking at p = 0.1 (73.65 ± 0.49 %) and monotonically declining toward higher p. Crucially, the small-world reservoir exceeded the connectome (67.29 %) and the random network (66.28 %) at every value of p examined, including the least favorable case for the small-world reservoir, that is, the value of p at which its margin over the other architectures was narrowest (p = 0.7, 67.88 ± 1.38 %). The small-world advantage is therefore not an artifact of a tuned rewiring probability.

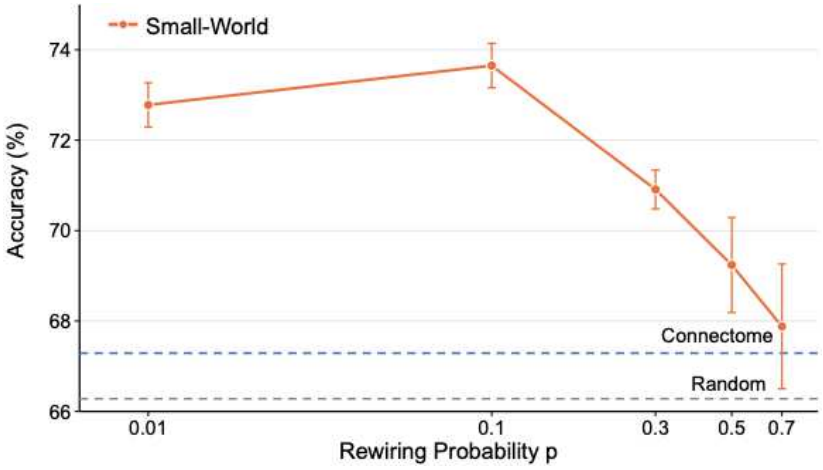


Fig. 3 Classification accuracy across the small-world family. Accuracy at N = 400 as a function of the Watts–Strogatz rewiring probability p, swept over p ∈ {0.01, 0.1, 0.3, 0.5, 0.7} with 20 realizations at each value. Accuracy peaks at p = 0.1 (73.65 ± 0.49 %) and declines monotonically as p increases. The connectome (67.29 %) and random (66.28 %) reservoirs are shown for reference; the small-world reservoir exceeds both at every value of p examined, including its weakest point at p = 0.7 (67.88 ± 1.38 %).

Benchmark test of prediction**.** Mackey–Glass chaotic-signal prediction is conducted. We used the standard benchmark instance distributed with the reference echo-state network implementation [18]. Of each run, 2000 steps were used for training, followed by 2000 test steps. Twenty independent reservoirs were generated and evaluated for each architecture (see Methods). The small-world reservoir attained the lowest prediction error (Fig. 4), with a mean-squared error (MSE) of $4.96 \times 10^{-8} \pm 1.8 \times 10^{-8}$, approximately 3.2-fold lower than both the connectome ($1.57 \times 10^{-7} \pm 1.7 \times 10^{-8}$; d = 6.27, P = 1.2×10^{-21}) and the random network ($1.64 \times 10^{-7} \pm 1.7 \times 10^{-8}$; P = 2.3×10^{-22}). Notably, the connectome and the random network were statistically indistinguishable (P = 0.17): on this task the connectome conferred no advantage over random wiring.

The results of Mackey–Glass prediction at N = 400 are summarized in Table 3, as means ± s.d. over 20 reservoir realizations. Prediction is one step ahead: at each test step the reservoir receives the true series value and predicts the next. The small-world reservoir is taken at p = 0.1 here and at p = 0.5 in Table 2, each being the value used for the corresponding three-way comparison. Effect sizes are computed on log10 MSE and refer to the comparison against the small-world reservoir; the connectome and the random network are statistically indistinguishable (d = −0.44).

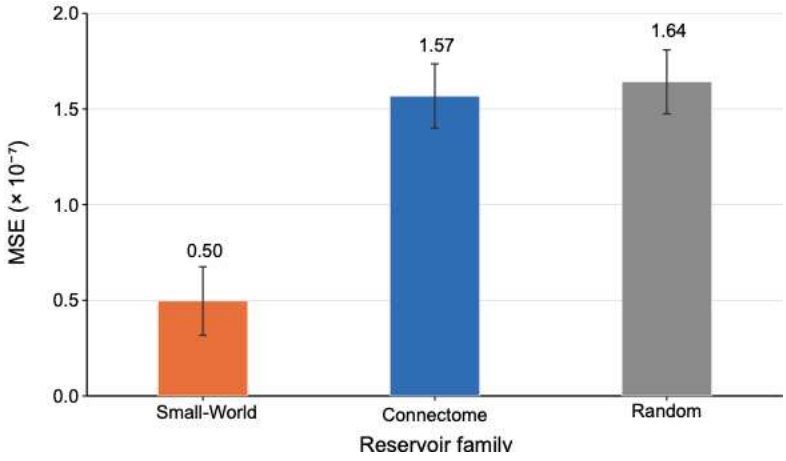


Fig. 4 Prediction error for the three architectures. Mean-squared error on one-step-ahead Mackey–Glass prediction at N = 400, over 20 reservoir realizations. The small-world reservoir attains 4.96×10^{-8}, roughly threefold lower than the connectome (1.57×10^{-7}) and the random network (1.64×10^{-7}), which do not differ from one another.

Table 3 Three-way comparison on the prediction task

Architecture	MSE (×10⁻⁸)	log₁₀ MSE	C	Degree s.d.	d
Small-world (p = 0.1)	4.96 ± 1.79	−7.32 ± 0.11	0.474	0.88	—
Connectome	15.68 ± 1.68	−6.81 ± 0.05	0.049	3.59	6.27
Random	16.42 ± 1.67	−6.79 ± 0.04	0.020	2.79	6.58

Sweeping p on the prediction task (Fig. 5) reproduces the pattern seen in classification: prediction error was lowest at p = 0.1 (4.96×10^{-8}) and the small-world reservoir outperformed the connectome (1.57×10^{-7}) at every p examined. Both benchmarks identified the same optimal rewiring probability p = 0.1.

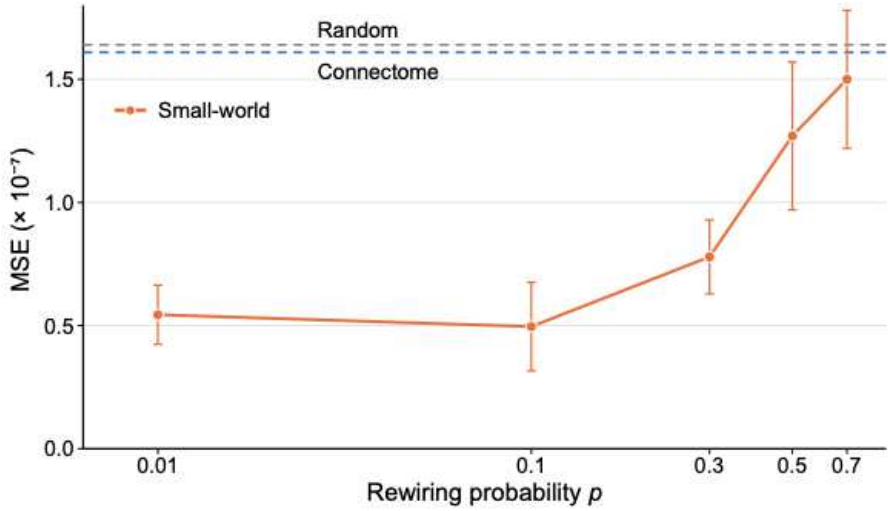


Fig. 5 Prediction error across the small-world family. Mean-squared error at N = 400 as a function of the rewiring probability, over the same five values of p and 20 realizations each. Error is lowest at p = 0.1 and the small-world reservoir outperforms the connectome at every value examined, reproducing on the prediction task the pattern seen for classification in Fig. 3.

Discussion

To identify the structural basis of these differences, we related performance to the network properties that vary across the rewiring sweep. As the rewiring probability p increases, clustering fell steeply (from C = 0.62 at p = 0.01 to C = 0.036 at p = 0.7) while degree heterogeneity rose (degree s.d. 0.29 → 1.90); characteristic path length changed comparatively little once p exceeded 0.1, the normalized value L(p)/L(0) falling from 0.17 to 0.12 across the remainder of the range at N = 400. We quantify clustering with the mean local clustering coefficient,

$$C = (1/N)\,\Sigma_i\, 2e_i / [k_i(k_i - 1)] \quad (1)$$

where k_i is the degree of node i and e_i the number of edges among its neighbors, and characteristic path length as the mean shortest-path distance between all node pairs,

$$L = 1 / [N(N - 1)]\,\Sigma_{i \neq j}\, d(i, j) \quad (2)$$

where d(i, j) is the shortest-path distance between nodes i and j.

In Fig. 6(a) the structural properties of the small-world reservoir at N = 400 are plotted against the rewiring probability. Performance on both benchmarks tracked the first two properties and not the third. Pairing the accuracies of Fig. 3 and the errors of Fig. 5 with the structure of Fig. 6(a) across the five sweep values, performance rose as clustering rose and fell as degree heterogeneity rose, strongly in both cases, the Pearson correlations lying between 0.88 and 0.94 in magnitude. Its correlation of performance with path length was not significant on either task. These correlations are computed across the five sweep values, so n = 5. The connectome sits away from the favorable region of this space on both axes at once: at N = 400 its clustering (C = 0.05) is an order of magnitude below the small-world optimum (C ≈ 0.47 at p = 0.1), and its degree heterogeneity (degree s.d. = 3.59) exceeds that of any small-world network in the sweep, and indeed that of the random network. Hubs give a brain short routes between distant regions, which serve a nervous system that must draw together information from many specialized areas; in a reservoir those same hubs concentrate the dynamics on a few nodes, which we suggest is what penalizes the connectome here.

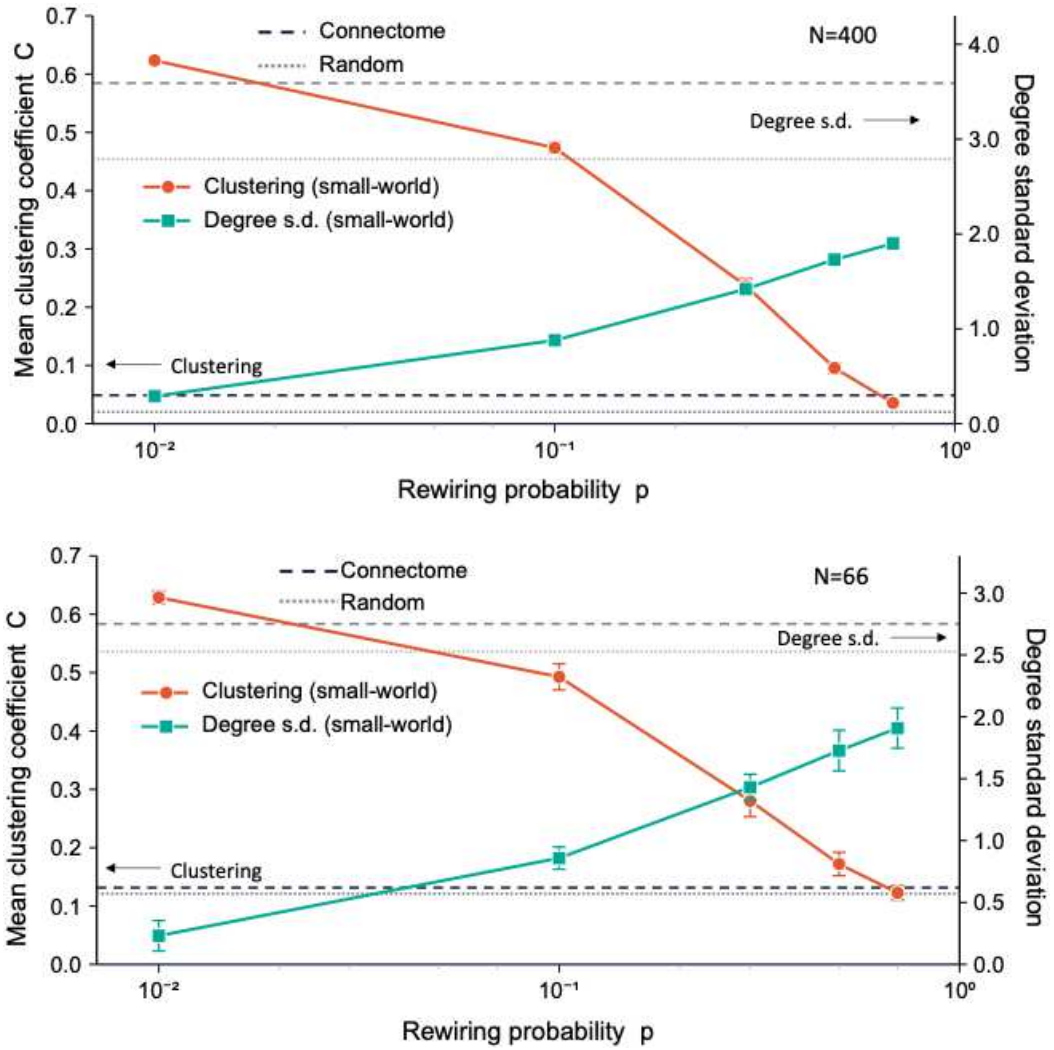


Fig. 6 The mean clustering coefficient (left axis) and the standard deviation of the degree distribution (right axis) against the rewiring probability p, at (a) N = 400 and (b) N = 66. Points are means over 20 reservoir realizations and error bars one standard deviation. The horizontal lines mark the two reference architectures at the same size and mean degree: dashed for the connectome, dotted for the random network. Each appears twice in a panel, once against each vertical axis, and the arrow beside a line indicates which axis it should be read against: a line marked with a left-pointing arrow gives that architecture's clustering coefficient, and one marked with a right-pointing arrow its degree standard deviation. Both properties of each reference architecture can therefore be compared directly with the corresponding small-world curve. A reference line lying below the clustering curve therefore indicates less clustering than any small-world network in the sweep, and one lying above the degree curve indicates greater degree heterogeneity.

Why should low heterogeneity work better than high heterogeneity? The readout is a single linear map applied to the reservoir's states, so its accuracy is limited by how many genuinely different signals those states supply. Where a few nodes hold most of the connections, the remaining nodes are driven almost entirely by those few, and their activity is largely a scaled and delayed copy of the hubs'; the reservoir then behaves as though it had fewer nodes than it does, because the readout has fewer independent signals to combine. When degrees are even, no node is redundant in this way and the same number of units yields more distinct signals. Normalization would be expected to amplify the difference: rescaling the recurrent matrix to a fixed spectral radius places every node at a comparable operating point when degrees are even, but concentrates the gain on the hubs when they are not, leaving peripheral nodes weakly driven. Clustering

raises the same quantity from the other side, since densely interconnected neighborhoods sustain local recurrent loops that hold information over several time steps, giving the readout a longer history to draw on.

This placement does not depend on the expansion procedure. No expansion is applied at N = 66: the connectome is the empirical 66-region matrix and the small-world networks are generated directly at that size, so Fig. 6(b) is independent of the construction described in the Methods. There the connectome again falls well below the small-world clustering optimum (C = 0.13 against 0.49 at p = 0.1) and is the most heterogeneous of the three architectures, with a degree standard deviation of 2.75 against 2.53 for the random network and at most 1.91 anywhere in the sweep. Absolute clustering is higher at N = 66 than at N = 400 because a fixed mean degree of 8 corresponds to an edge density of 0.123 there against 0.020, so the comparison rests on the ordering of the three architectures rather than on the absolute values. The margins are narrower at this size, but the ordering is unchanged, and it is at this size that the connectome also fails to exceed the accuracy of the random network, so structure and performance agree where no construction step intervenes.

This result sits alongside rather than against earlier work reporting that connectome-informed reservoirs outperform null networks [14], that is, synthetic comparison networks built by scrambling the connectome while holding one or more of its properties fixed. Those comparisons held the degree sequence fixed and scrambled only the topology, so the null carried the same degree heterogeneity as the connectome, whereas our random networks fix the number of connections alone and is therefore degree-homogeneous as well as unstructured. The two findings are consistent: against a null that shares its degree distribution the connectome's topology is an advantage, while against one that does not share the degree distribution, that advantage is offset by the heterogeneity of its degrees. In addition, this result is consistent with the report on photonic reservoir computing with complex networks[23].

Several limitations bound these conclusions. We examined a single connectome at one coarse parcellation of 66 regions and two benchmarks, so other parcellations, species or task families may behave differently. All reservoirs were binary, so the connectome's connection strengths were discarded, and matching to a fixed mean degree required removing a fraction of its connections. The structure–performance relationships are correlational and rest on the five values of the rewiring sweep, and across that sweep clustering and degree heterogeneity co-vary almost perfectly, so the present design establishes that performance tracks the two together but cannot separate their individual contributions; a network family that varies them independently would be needed for that. The account we offer of why degree homogeneity should help is likewise an interpretation rather than a measurement: it is consistent with the observed ordering, but we did not measure the effective dimensionality of the reservoir states directly, which would test it. Finally, the spectral radius and leaking rate were fixed rather than tuned per architecture, although a sensitivity check on the prediction task found the chosen value near-optimal.

A further question follows from these limitations. Our benchmarks ask the reservoir to separate six activity classes and to predict one step of a chaotic series, and for these the connectome is not advantageous. It does not follow that its architecture is unsuited to computation in general. Modular, hub-dominated connectivity may instead favor tasks that call on the properties those features serve in the brain, such as integrating several input streams that arrive at different sites, holding information over longer intervals, or performing more than one task at once without interference. Identifying the class of tasks for which connectome-derived connectivity is genuinely advantageous, rather than merely not disadvantageous, seems to us the natural next step.

Taken together, these results indicate that the computational capability of a reservoir in these tasks tracks generic network statistics, clustering and degree homogeneity, more closely than it tracks biological provenance. Because the two statistics co-vary across the rewiring sweep, the present design establishes their joint association with performance rather than the separate contribution of either. For the design of reservoir computing, this is an encouraging conclusion: the relevant properties are ones that synthetic generators control directly and cheaply. For neuroscience, it sharpens the question. Whether connectome topology confers a computational advantage may well depend on the tasks and dynamical regimes considered, and on the choice of null model; comparisons against nulls that separate topology from degree would help to settle the question. The converse approach points the same way: when artificial recurrent networks are optimized under biological cost constraints rather than given a connectome outright, they converge on modular small-world organization[22].

Methods

Overview. We implemented each network as the fixed recurrent layer (reservoir) of an echo-state network and evaluated its computational capacity on two supervised benchmarks. Only the linear readout was trained; the recurrent reservoir remained fixed throughout. This design isolates the contribution of connectivity itself: because the recurrent weights are never adapted, any systematic difference in task performance between reservoir families can be attributed to the arrangement of their connections rather than to learning. We used this property to ask whether the topology of the human structural connectome confers a computational advantage over synthetic small-world and random architectures matched for size and edge density.

Network construction. We compared three families of reservoir network: a human structural connectome, a Watts–Strogatz small-world network, and a random network, all undirected, binary, and matched for the number of nodes N and edge density. Structural connectivity was taken from a human consensus connectome of 66 cortical regions (33 per hemisphere) derived from the Lausanne dataset [15], distributed with the conn2res toolbox [13]. The dataset was acquired from 70 healthy adults by diffusion spectrum imaging and deterministic streamline tractography. This 66-region matrix defines the native resolution of the connectome reservoir (N = 66) shown in Fig. 7(a). The distributed matrix is symmetric and weighted, containing 311 connections (mean degree 9.42). It was symmetrized, then converted to a binary reservoir of the required size by the expansion and sparsification procedure described below; the reservoir therefore encodes the connectome's wiring pattern only so that the comparison between architectures rests on topology alone.

Small-world reservoirs were generated with the Watts–Strogatz model [10]: starting from a ring lattice in which each node is joined to its K nearest neighbors on either side, every edge was rewired with probability p to a randomly chosen target, excluding self-loops and duplicate edges. Here K counts neighbors on each side, so the mean degree is 2K; we used K = 4 throughout, giving a mean degree of 8. Fig. 7(b) shows the resulting topology at N = 400 with $(k, p) = (4, 0.1)$. The three-way comparison against the connectome and random architectures used a rewiring probability of p = 0.5 for the classification task and p = 0.1 for the prediction task, each value being one of the five covered by the sweep described below. For classification this is not the value at which small-world performance peaks, so that comparison is a conservative test of the small-world advantage; for prediction it is the best-performing value of the sweep. The rewiring probability was swept over five values, $p \in \{0.01, 0.1, 0.3, 0.5, 0.7\}$.

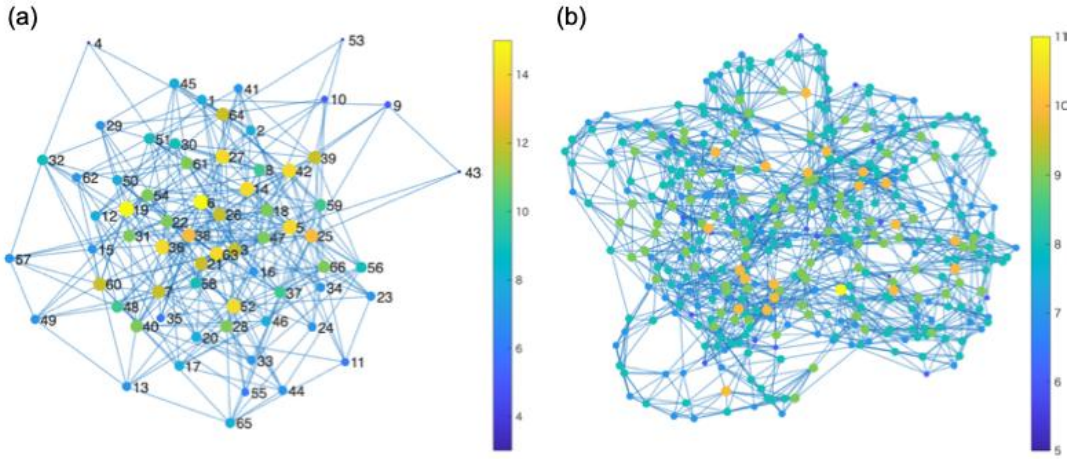


Fig. 7 Structural connectome used as the reservoir. (a) The66-region human consensus connectome drawn as a network. Each node is one cortical region and each link a structural connection retained after sparsification. The layout is force-directed, so densely interconnected regions are placed close together and the spatial arrangement carries no anatomical meaning. Node size and node colour both encode degree, the number of connections a region makes, so hub regions appear both larger and brighter; the colour scale spans degrees 3 to 15. (b) Small-world reservoir topology. A Watts–Strogatz reservoir at N = 400. The

network was built from a ring lattice in which each node joins its K = 4 nearest neighbors on either side, giving a mean degree of 8, and every edge was then rewired with probability p = 0.1. The colour scale spans degrees 5 to 11, a narrower range than the connectome of (a), reflecting the far more homogeneous degree distribution of this architecture; because the two figures use independent colour scales, colours are not comparable between them.

Random reservoirs were generated by placing the required number of connections uniformly at random among all node pairs, excluding self-connections, and symmetrizing the result; the connection count was fixed at $E = N\langle k\rangle/2$, giving 1600 connections and a mean degree of 8 at N = 400, matching the other two architectures. Because placement is uniform, node degrees follow a Poisson distribution about the mean, with a standard deviation of approximately 2.8 at N = 400, intermediate between the small-world reservoirs, whose degrees are nearly uniform, and the connectome. Connectivity was not enforced: realizations were used as generated, and a node of degree zero arises in roughly one realization in seven at this size, consistent with the Poisson expectation. Twenty independent realizations were generated.

Network expansion and sparsification. Connectome reservoirs of any size N were built from the 66-region matrix in three steps: allocation of neurons to regions, expansion to a neuron-level network in the spirit of the bio2art approach [21], and sparsification to the target mean degree as illustrated in Fig. 8. Each region i was assigned m_i neurons, with a base allocation of $\lfloor N / 66 \rfloor$, where $\lfloor \cdot \rfloor$ denotes rounding down to the nearest whole number, to every region and the remaining $N - 66\lfloor N / 66 \rfloor$ neurons given one each to the regions of greatest total connection strength, so that hub regions are represented by the largest blocks. At N = 400 this gives 62 regions of 6 neurons and 4 regions of 7; at N = 66 every region reduces to a single neuron.

Within each region, neurons were connected all-to-all, excluding self-connections, with a uniform weight equal to the median of the non-zero connection weights of the empirical matrix; this imposes the modular structure absent from a region-level matrix. Between two regions joined in the empirical connectome, every pair of their neurons was connected and assigned the empirical weight of that region pair, so the coarse-grained connectivity of the expanded network reproduces the connectome and regions with many or strong connections give rise to correspondingly many neuron-level projections. Region pairs unconnected in the empirical matrix remained unconnected. The result was symmetrized by averaging with its transpose and the diagonal set to zero.

The expanded network is far denser than the target, so it was reduced to exactly $E = N\langle k\rangle / 2$ connections (264 at N = 66 and 1600 at N = 400) by sampling that many of its connections without replacement, each with probability proportional to its expanded weight. Strong region pairs are therefore more likely to retain links than weak ones, so the modular and hub structure of the connectome survives the reduction. The retained connections were finally binarized. Because the sampling is stochastic, each of the 20 realizations at a given size is a different draw; this is the sole source of variability in the connectome arm at the native resolution, where the expansion step is an identity operation and the procedure reduces to sampling 264 of the connectome's 311 connections. In Fig.8 the process of expansion and sparsification of the connectome are illustrated.

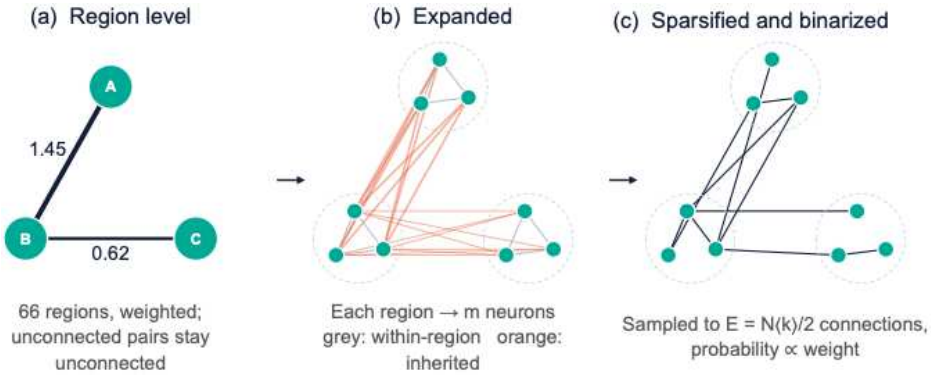


Fig. 8 Expansion and sparsification of the connectome. Schematic of the procedure, illustrated with three regions A, B, and C and three neurons per region. (a) The region-level connectome is weighted, and pairs of regions A-C that are unconnected in the empirical matrix remain unconnected throughout. (b) Each

region is instantiated as a block of m neurons. Neurons within a block are connected all-to-all at the median empirical weight (grey), and every pair of neurons belonging to two connected regions is joined at the weight of that region pair (orange, thicker for stronger connections), so that the coarse-grained connectivity of the expanded network reproduces the connectome. (c) The expanded network is reduced to the target edge count E = N⟨k⟩/2 by sampling connections without replacement with probability proportional to their weight, and the survivors are binarized; strongly connected region pairs therefore retain more links than weakly connected ones.

Reservoir dynamics. Each reservoir served as the recurrent layer of an echo-state network. The reservoir state was updated at each time step as [18]

$$x(t) = (1 - a)\, x(t-1) + a\, f(\, W_{in}\, u(t) + W\, x(t-1)\,) \quad (3)$$

$$y(t+1) = W_{out} x(t) \quad (4)$$

where $u(t)$ is the input, $x(t)$ the vector of reservoir states, W the fixed recurrent weight matrix, W_{in} the input weights, f the activation function applied element-wise, and a the leaking rate. The activation function was the hyperbolic tangent. Input weights were drawn once from a uniform distribution on [−0.5, 0.5] and applied to every reservoir node, together with a constant bias term. The same input weights were used for every realization and every architecture, so that the variability reported across realizations reflects the reservoir connectivity alone. Before each simulation the recurrent matrix was rescaled to a target spectral radius ρ by dividing by its largest eigenvalue magnitude and multiplying by ρ. The two tasks were run with different settings: the classification reservoirs used a leaking rate of 0.2 and ρ = 0.9, the prediction reservoirs a leaking rate of 0.6 and ρ = 1.2. Because the update is leaky, the operator governing the linearized dynamics is (1 − a)I + aW, whose spectral radius is (1 − a) + aρ: 0.98 for classification and 1.12 for prediction. The classification reservoirs therefore operated just below the echo-state bound, and the prediction reservoirs slightly above it. Both values were fixed in advance for their respective task rather than tuned during the experiments. A sensitivity check on the prediction task, varying ρ from 0.6 to 1.4 with every other setting held constant, showed a shallow optimum near ρ = 1.1 (effective radius 1.06), with prediction error varying by a factor of 2.5 across the range; the value used lies within 10% of that optimum, and adopting the effective radius of the classification task would have given marginally higher error. Prediction on this benchmark therefore favors dynamics slightly beyond the echo-state bound.

Readout and training. A linear readout was trained on the reservoir states to predict the task target; the recurrent and input weights were never updated. At each time-step the readout received the concatenation of a constant bias, the current input and the activations of all N reservoir nodes, so no subset of nodes was designated for readout. Its weights were obtained in closed form by ridge regression, as set out below, with regularization parameter $\lambda = 1 \times 10^{-8}$ for classification and 0.75×10^{-8} for prediction. For the prediction task the first 100 time-steps of each run were discarded as washout before states were collected; the classification task, in which states are collected within each trial, applied no washout. Prediction was one step ahead, the reservoir receiving the true series value at each step rather than its own output.

The readout weights are obtained as the unique solution of $W_{out} x_{tr} \cong y$ in the least-squares sense. Applying Tikhonov (ridge) regularization gives the closed-form solution [18]

$$W_{out} = y x_{tr}^{T} (x_{tr} x_{tr}^{T} + \lambda I_{N \times N})^{-1} \quad (5)$$

where x_{tr} is the matrix of reservoir states collected over the training period, y the corresponding matrix of targets, λ the regularization parameter, and I the identity matrix, whose dimension equals the number of reservoir states passed to the readout. Because training reduces to a single closed-form linear solve, its computational cost is low. The regularization term λI penalizes large readout weights and keeps the matrix inverse well conditioned when reservoir states are strongly correlated, as they generally are in a recurrent network. The value of λ therefore sets the balance between fitting the training data closely and generalizing to unseen data: larger values yield smoother, more conservative readouts at the cost of training accuracy, while values that are too small make the solution sensitive to noise in the reservoir states.

Benchmark tests: Classification.

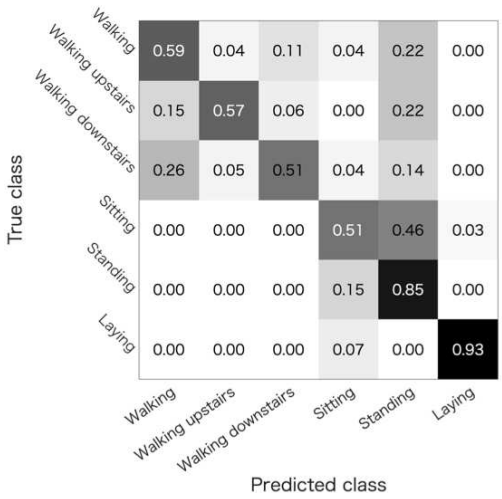


Fig. 9 Confusion matrix for the six-class human-activity recognition task at N = 400, averaged over the 20 reservoir realizations. Rows give the true activity and columns the predicted activity, and each row is normalized to sum to one, so the diagonal entries are per-class recall and darker cells indicate a larger proportion of that class. Activities are ordered as in the source dataset: the three ambulatory classes first, then the three static postures.

The data set of the human activity is available at http://archive.ics.uci.edu/ml/machi ne- learning databases/ 00240/UCI HAR Dataset.zip. It includes six motions: walking, walking upstairs, walking downstairs, sitting, standing, and laying down. The captured temporal waveforms along x-, y-, and z-axes for 2.56 s were sampled into 128-sample at the rate of 20 ms. The data sets of the training and the testing are 7352 and 2947, respectively. The reservoir state reached at the end of a window was passed to the readout, which had six outputs trained against the one-hot encoded activity label. The dataset's standard subject-wise partition was used without modification: 7352 windows for training and 2947 for testing, drawn from disjoint sets of volunteers. The raw signals were used as supplied, without rescaling or normalization. In Fig. 9 a confusion matrix for the six-class human-activity recognition task at N = 400 is shown.

Benchmark tests: Prediction (Mackey–Glass). One-step-ahead prediction of the Mackey–Glass chaotic time series, which follows $dx/dt = \beta x(t-\tau) / [1 + x(t-\tau)^n] - \gamma x(t)$ [17]. We used the conventional chaotic parameter set $\beta = 0.2$, $\gamma = 0.1$, $n = 10$ with delay $\tau = 17$. We took the standard benchmark instance distributed with the reference echo-state network implementation [18]. The series was supplied to the reservoir as a single input channel. Of each run, 2000 time-steps were used for training, the first 100 of which were discarded as washout, followed by 2000 test steps; at each test step the reservoir received the true series value and predicted the next, and the mean-squared error was computed over the first 500 test steps. Representative target and predicted waveforms at N = 400 are shown in Fig. 10.

Rewiring-probability sweep. To test whether the small-world advantage depended on a particular choice of the rewiring probability, we swept p at N = 400 and evaluated both benchmarks at each value, holding the node count, mean degree and all reservoir and readout parameters fixed. Five values were used, $p \in \{0.01, 0.1, 0.3, 0.5, 0.7\}$, with 20 independent realizations at each, giving 100 small-world reservoirs.

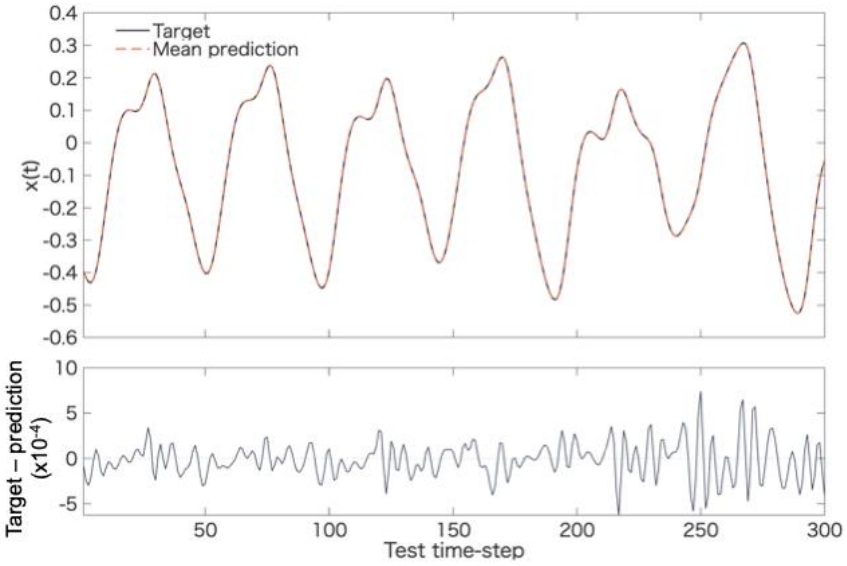


Fig. 10 One-step-ahead prediction over the held-out test set at N = 400, averaged across the 20 reservoir realizations. (a) Temporal waveforms of target (solid) and the mean prediction (dashed) over the first 300 test time-steps. The two are visually indistinguishable at this scale, and the shaded band showing one standard deviation across realizations is narrower than the line width. (b) The corresponding residual, target minus mean prediction, plotted on a scale of 10^{-4}. Note that the residual is the prediction error itself, whereas the errors quoted in the text are mean squared errors: a root-mean-square residual of 2.2×10^{-4} corresponds to a mean squared error of 5.0×10^{-8}. The residual is roughly three orders of magnitude smaller than the amplitude of the series, which is why the two traces in (a) overlap, and it fluctuates about zero without systematic drift, indicating that the reservoir tracks the series rather than accumulating error.

Performance metrics and statistical analysis. Classification was scored with balanced accuracy, the mean of the per-class recalls, which weights each of the six activities equally regardless of how often it occurs. Prediction was scored as the mean-squared error between predicted and true values on the held-out test set; for the effect sizes, tests and correlations reported for that task the error was transformed to its base-10 logarithm, on which the distributions are approximately normal. For each condition 20 independent reservoir realizations were generated and the mean ± standard deviation reported. Pairwise differences between architectures were assessed with two-sided independent-samples t-tests assuming equal variances, with effect size given by Cohen's d; no correction for multiple comparisons was applied. Structure–performance relationships across the rewiring sweep were quantified with Pearson correlations between network measures and performance, computed on the five sweep means.

Software. All network construction, reservoir simulation and analysis were carried out in MATLAB R2026a. The Statistics and Machine Learning Toolbox was used for weighted sampling without replacement in the connectome expansion and for correlation and hypothesis testing; no other toolbox was required, and the graph-theoretic measures were computed with functions included in the analysis code.

Data availability. The human structural connectome used in this study is the 66-region consensus matrix distributed with the conn2res toolbox [13], derived from the diffusion spectrum imaging dataset of Griffa, Alemán-Gómez and Hagmann, which is publicly available at https://doi.org/10.5281/zenodo.2872624. The UCI Human Activity Recognition dataset is available from the UCI Machine Learning Repository [16]. The Mackey–Glass series is the standard $\tau = 17$ benchmark instance distributed with the reference echo-state network implementation [18].

Code availability. All code used to generate the reservoir networks, run the simulations, compute the network measures and produce the figures of this study are available upon reasonable request from the authors.

References

1. Verstraeten, D., Schrauwen, B., Stroobandt, D. & Van Campenhout, J. Isolated word recognition with the liquid state machine: a case study. Inf. Process. Lett. 95, 521–528 (2005).
2. Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. Comput. Sci. Rev. 3, 127–149 (2009).
3. Jaeger, H. The 'echo state' approach to analysing and training recurrent neural networks. GMD Report 148, German National Research Center for Information Technology (2001).
4. Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science 304, 78–80 (2004).
5. Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. Neural Comput. 14, 2531–2560 (2002).
6. Maass, W. Liquid state machines: motivation, theory, and applications. In Computability in Context: Computation and Logic in the Real World 275–296 (World Scientific, 2011).
7. Rodan, A. & Tiňo, P. Minimum complexity echo state network. IEEE Trans. Neural Netw. 22, 131–144 (2011).
8. Deng, Z. & Zhang, Y. Collective behavior of a small-world recurrent neural system with scale-free distribution. IEEE Trans. Neural Netw. 18, 1364–1375 (2007).
9. Kawai, Y., Park, J. & Asada, M. A small-world topology enhances the echo state property and signal propagation in reservoir computing. Neural Netw. 112, 15–23 (2019).
10. Watts, D. J. & Strogatz, S. H. Collective dynamics of 'small-world' networks. Nature 393, 440–442 (1998).
11. Milgram, S. The small world problem. Psychol. Today 2, 60–67 (1967).
12. Kitayama, K. Guiding principle of reservoir computing based on 'small-world' network. Sci. Rep. 12, 16697 (2022).
13. Suárez, L. E., Mihalik, A., Milisav, F., Marshall, K., Li, M., Vértes, P. E., Lajoie, G. & Misic, B. Connectome-based reservoir computing with the conn2res toolbox. Nat. Commun. 15, 656 (2024).
14. Suárez, L. E., Richards, B. A., Lajoie, G. & Misic, B. Learning function from structure in neuromorphic networks. Nat. Mach. Intell. 3, 771–786 (2021).
15. Griffa, A., Alemán-Gómez, Y. & Hagmann, P. Structural and functional connectome from 70 young healthy adults [Data set]. Zenodo https://doi.org/10.5281/zenodo.2872624 (2019).
16. Anguita, D., Ghio, A., Oneto, L., Parra, X. & Reyes-Ortiz, J. L. Human activity recognition on smartphones using a multiclass hardware-friendly support vector machine. In International Workshop on Ambient Assisted Living (IWAAL 2012) (Vitoria-Gasteiz, Spain, 2012).
17. Mackey, M. C. & Glass, L. Oscillation and chaos in physiological control systems. Science 197, 287–289 (1977).
18. Lukoševičius, M. A practical guide to applying echo state networks. In Neural Networks: Tricks of the Trade 659–686 (Springer, 2012).
19. Suárez, L. E., Markello, R. D., Betzel, R. F. & Misic, B. Linking structure and function in macroscale brain networks. Trends Cogn. Sci. 24, 302–315 (2020).
20. Damicelli, F., Hilgetag, C. C. & Goulas, A. Brain connectivity meets reservoir computing. PLoS Comput. Biol. 18, e1010639 (2022).
21. Goulas, A., Damicelli, F. & Hilgetag, C. C. Bio-instantiated recurrent neural networks: integrating neurobiology-based network topology in artificial networks. Neural Netw. 142, 608–618 (2021).
22. Achterberg, J., Akarca, D., Strouse, D. J., Duncan, J. & Astle, D. E. Spatially embedded recurrent neural networks reveal widespread links between structural and functional neuroscience findings. Nat. Mach. Intell. 5, 1369–1381 (2023).
23. Park, S., Watabe, K., Sunada, S., Yamagami, T. and Uchida, A. Photonic reservoir computing with complex networks. arXiv preprint arXiv:2607.23285 (2026).

Acknowledgements

This study was supported in part by JSPS KAKENHI (JP22H05195, JP25H01129) and JST CREST (JPMJCR24R2) in Japan. Kitayama K. thank the members of AI Working Group of Hamamatsu Photonics for useful discussions. Kitayama k. also would like to thank Inoue M. of NICT and Toyoda H. of Hamamatsu Photonics for their continuous supports.

Author contributions:
Kitayama K. conceived the study, performed the simulations and analyses, and wrote the manuscript.
Uchida A. contributed to the discussion and to revision of the manuscript.

Competing interests
The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Manuscriptrevisedhighlightedkitayama.docx