Supporting Information

Rubén J. Sánchez García[1*]

[1]Mathematical Sciences, University of Southampton, University Road, Southampton SO17 1BJ, United Kingdom
[*]R.Sanchez-Garcia@soton.ac.uk

SI Text

This Supporting Information provides additional details, including full mathematical proofs, of the statements in the Main Text (Results and Methods). It is organised into sections of the same headings as in the Main Text, to which they correspond. We assume the terminology and notation in the Main Text.

Related Work

We have kept the content as self-contained as possible, by including details on material relatively well known in the graph theoretic literature, but perhaps not so for an applied audience. The automorphism group of a graph is well studied in algebraic and spectral graph theory[1–3], and the concepts of equitable partitions, characteristic matrix and quotient graph can be found in e.g.[3,4], including the relation between quotient and parent eigenvalues, and the main ingredients of the spectral decomposition for binary graphs [3, Remark 3.9.6]. More recently, Francis et al.[5,6] have developed a general theory of equitable decompositions for automorphism compatible matrices. The geometric decomposition and symmetric motifs were originally defined in[7], and the redundant spectrum and spectral decomposition in[8], in both cases for binary adjacency matrices only.

Following the seminal work of MacArthur et al.[7,8], symmetry has been used in empirical networks, for instance to study the quotient as a coarse graining tool real-world networks[9], detect symmetric motifs via symmetry compression[10], and reduce shortest path query computations[11]. The redundant Laplacian spectrum has been explored for regular degree[12] and random geometric graphs[13], and the dynamical implications of eigenvector localisation analysed[14]. A key motivation of this article has been to develop the most general common framework for these (and other) results.

SI Symmetry in Complex Networks

On Labels and Symmetries

A network is a combinatorial object which encodes pairwise relations (edges or links) between objects (vertices or nodes). It is therefore independent of the ordering, or labelling, of the vertices. An ordering is needed to refer to, and work with, a network. We can choose an ordering simply by enumerating the vertices 1 to $n = |V|$. Such ordering is needed, for example, to define the adjacency matrix. Note that a different ordering results in a (possibly) different adjacency matrix of the same network. We always assume a chosen, and thereafter fixed, labelling 1 to $n = |V|$ of the vertices.

On the other hand, a symmetry, or automorphism, of a network is a permutation of the vertices preserving adjacency. It can also be thought of as a vertex relabelling, but one that results on the *same* adjacency matrix. (This is expressed mathematically by the condition $AP = PA$, Eq. (1) in the Main Text.) Therefore, to find symmetries, we first fix an initial labelling of the vertices (and in particular an adjacency matrix), so we can unequivocally refer to the vertices, and then look for further relabelling preserving the adjacency matrix.

However, as a relabelling, a symmetry or automorphism σ produces the same graph, and, necessarily, vertices i and $\sigma(i)$, or edges (i, j) and $(\sigma(i), \sigma(j))$, are indistinguishable from one another and therefore structurally equivalent. In particular, for a vertex, respectively pairwise, network measure depending on structure alone, we have $G(i) = G(\sigma(i))$, respectively $F(i, j) = F(\sigma(i), \sigma(j))$, for all $\sigma \in \mathrm{Aut}(\mathscr{G})$, and for all $i, j \in V$, that is, Eqs. (5) and (3) in the Main Text.

Visually, automorphisms still correspond to symmetries, as perceived by the human eye, in a (suitable) geometric representation of the network (cf. Figure 1 in the Main Text).

Geometric Decomposition

We now explain the geometric decomposition introduced in Methods, in more detail, following MacArthur et al.[7]. We start with some preliminary notions. Each automorphism $\sigma \in \mathrm{Aut}(\mathscr{G})$ is a permutation of the vertices of the graph, and automorphisms can be composed by applying one permutation after the other, forming a mathematical structure called a *group*[15]. The *support* of the permutation σ is the set of vertices moved by σ,

$$\mathrm{supp}(\sigma) = \{i \in V \text{ such that } \sigma(i) \neq i\}. \tag{1}$$

Two automorphisms σ and τ are *support-disjoint* if the intersection of their supports is empty, $\text{supp}(\sigma) \cap \text{supp}(\tau) = \emptyset$. In particular, σ and τ commute, that is, the order in which they are applied does not affect the result, mathematically $\sigma\tau = \tau\sigma$. Similarly, two arbitrary subsets of automorphisms $S, T \subset \text{Aut}(\mathscr{G})$ are *support-disjoint* if σ and τ are support-disjoint for every pair $\sigma \in S$ and $\tau \in T$.

The geometric decomposition is obtained by partitioning a *set of generators* $X = \{\sigma_1, \ldots, \sigma_k\}$ of $\text{Aut}(\mathscr{G})$ (that is, every automorphism $\sigma \in \text{Aut}(\mathscr{G})$ is the product, or composition, of elements in X), which the additional property of being *essential* (explained below), into pairwise support-disjoint subsets

$$X = X_1 \cup \ldots \cup X_k. \tag{2}$$

(How to obtain in practice an essential set of generators, and the support-disjoint partition, is explained in the Network Symmetry Computation section below). Let us call M_i the set of vertices moved by generators in X_i, that is,

$$M_i = \{i \in V \text{ such that } \sigma(i) \neq i \text{ for some } \sigma \in X_i\} = \bigcup_{\sigma \in S_i} \text{supp}(\sigma), \tag{3}$$

and $\mathscr{M}_i$ the subgraph induced by M_i (that is, the graph with vertex set M_i, and all edges in $\mathscr{G}$ between vertices in M_i). Further, define $H_i = \langle X_i \rangle$, the subgroup generated by X_i (that is, all the permutations obtained by composing elements in X_i) and call it a *geometric factor*. The support-disjoint decomposition of X above, (2), gives a direct product[15] decomposition of $\text{Aut}(\mathscr{G})$ into geometric factors

$$\text{Aut}(\mathscr{G}) = H_1 \times \ldots \times H_k, \tag{4}$$

called the *geometric decomposition* of $\text{Aut}(\mathscr{G})$. In other words, every automorphism $\sigma \in \text{Aut}(\mathscr{G})$ can be uniquely decomposed as a product (composition)

$$\sigma = \tau_1 \tau_2 \ldots \tau_k, \quad \tau_i \in H_i = \langle X_i \rangle, \tag{5}$$

of permutations τ_i of vertices of M_i only. Hence understanding the symmetries of each subgraph $\mathscr{M}_i$, we recover all automorphisms of $\mathscr{G}$. (Of course not every automorphism of $\mathscr{M}_i$, considered as a graph on its own, is an automorphism of $\mathscr{G}$, as it also depends on how $\mathscr{M}_i$ is embedded in $\mathscr{G}$.) The subgraphs $\mathscr{M}_i$ are called *symmetric motifs* since they generate all the network symmetry. In real-world networks, symmetric motifs are typically small, however, as the network automorphisms consists on all possible combinations of these localised symmetries, their presence explains the large size, in absolute terms, of $\text{Aut}(\mathscr{G})$ observed empirically (SI Table 1,[7]).

Proposition 1. *The support-disjoint decomposition* (2) *implies the direct product decomposition* (4).

Proof. This is an immediate consequence of the characterisation of direct product of groups[15], shown here for $m = 2$: a group G is isomorphic to $K \times L$ for subgroups K and L if (i) their intersection $K \cap L$ is trivial (the identity element only); (ii) every element in G is a product kl with $k \in K$ and $l \in L$; (iii) the elements in K and L commute $kl = lk$ for all $k \in K$, $l \in L$. Since X is a generating set, condition (ii) is satisfied by $H_1 = \langle X_1 \rangle$ and $H_2 = \langle X_2 \rangle$, and the support-disjoint condition immediately implies (i) and (iii). The case $m > 2$ is a simple generalisation. □

Not every vertex participates in a network symmetry, and we write M_0 for the set of such *fixed* vertices, that is,

$$M_0 = \{i \in V \text{ such that } \sigma(i) = i \text{ for all } \sigma \in \text{Aut}(\mathscr{G})\}. \tag{6}$$

We then have a partition of the vertex set

$$V = M_0 \cup M_1 \cup \ldots M_k. \tag{7}$$

That is, each network can be partitioned into an *asymmetric core* (the subgraph, typically connected, induced by the fixed vertices M_0) and the symmetric motifs. The asymmetric core is related, but not equal, to the quotient (as a vertex set, the quotient equals the asymmetric core plus one vertex per orbit, see the Quotient Network section below).

Every symmetric motif can be further decomposed into one or more orbits of structurally equivalent vertices. The *orbit* of a vertex i is the set of vertices to which i can be moved to by an automorphism, that is,

$$\{\sigma(i) \text{ such that } \sigma \in \text{Aut}(\mathscr{G})\}. \tag{8}$$

Every vertex belongs to an orbit (made of just the vertex itself if it is fixed) and if a symmetric motif contains a vertex in an orbit, then it must contain all the vertices in the orbit.

Proposition 2. *If vertices i and j belong to the same orbit, then they also belong to the same symmetric motif.*

Name	$n_{\mathscr{G}}$	gen	aut	sm	bsm_1	bsm_2	bsm_{3+}	opm	vpm	$\tilde{n}_{\mathscr{Q}}$	$\tilde{n}_{\mathscr{Q}}^{\text{basic}}$	$\tilde{m}_{\mathscr{Q}}$	$\tilde{m}_{\mathscr{Q}}^{\text{basic}}$
HumanDisease	1,419	713	10^{407}	272	259	2	0	1.01	3.68	48.3	52.2	50.4	53.8
Yeast	1,647	380	10^{218}	149	138	10	0	1.07	3.46	76.3	76.4	83.5	83.7
OpenFlights	3,397	732	10^{373}	321	290	10	0	1.03	3.33	77.3	79.7	94.4	95.0
USPowerGrid	4,941	414	10^{153}	302	266	22	6	1.15	2.36	90.2	91.6	91.4	92.5
HumanPPI	9,270	972	10^{528}	437	434	3	0	1.01	3.22	89.5	89.5	96.9	97.0
Astro-Ph	17,903	3,232	$10^{1,588}$	1,682	1672	0	0	1.00	2.93	81.9	82.1	80.5	80.5
InternetAS	34,761	15,587	$10^{15,403}$	3,189	3,153	32	1	1.01	5.85	55.0	55.1	78.2	78.2
WordNet	145,145	52,152	$10^{24,064}$	28,456	25,759	375	52	1.03	2.98	60.1	70.2	58.0	63.8
Amazon	334,863	32,098	$10^{12,495}$	23,302	22,964	286	13	1.01	2.37	90.3	90.4	89.0	89.0
Actors	374,511	182,803	$10^{39,950}$	36,703	36,683	4	0	1.00	5.98	51.2	51.2	66.4	66.4
InternetAS-skitter	1,694,616	319,738	$10^{258,835}$	84,675	81,537	2,226	169	1.03	3.82	85.4	85.8	92.7	92.8
CaliforniaRoads	1,957,027	36,430	$10^{11,228}$	35,210	33,220	1,255	328	1.08	2.04	97.9	98.1	98.5	98.6
LiveJournal	5,189,808	410,575	$10^{197,552}$	245,211	243,743	1,182	26	1.01	2.67	92.1	92.1	97.5	97.5

Table 1. Symmetry in some real-world networks (continued). In addition the summary statistics shown in Table 1 in the Main Text, here we report the size of the automorphism group to the closest power of 10 (aut), the number of BSMs with one, two, or more than two orbits (bsm_1, bsm_2, bsm_{3+}), the average number of orbits per motif (opm) and vertices per orbit (vpo), and the proportion of vertices $\tilde{n}_{\mathscr{Q}}^{\text{basic}}$ and edges $\tilde{m}_{\mathscr{Q}}^{\text{basic}}$ in the basic quotient.

Proof. We can assume $i \neq j$. As they belong to the same orbit, there is an automorphism σ with $\sigma(i) = j$. Write σ as a product of generators $\sigma = x_1 \ldots x_l$, $x_i \in X_i$ and argue by induction on $l \geq 1$. If $l = 1$, then $x_1(i) = j$ and $i \neq j$ hence $i \in \operatorname{supp}(x_1)$. In addition, $j \in \operatorname{supp}(x_1)$: otherwise $x_1(j) = j$ implies $j = x_1^{-1}(j) = i$, a contradiction. As $i, j \in \operatorname{supp}(x_1)$, they are in the same symmetry motif, by (3). Suppose now $l > 1$ and the induction hypothesis. Write $k = x_2 \ldots x_l(i)$ so that $j = x_1(k)$. Assume $j \neq k$ (otherwise remove x_1 so that $j = x_2 \ldots x_l(i)$). By the induction hypothesis, k and i belong to the same symmetric motif, and, by the same argument as in the case $l = 1$, $j \neq k$ also belong to the same symmetric motif. □

This proposition, implicit in[7], proves that the vertices in each symmetric motif can be partitioned into orbits,

$$M_i = V_i^{(1)} \cup \ldots \cup V_i^{(m_i)}. \tag{9}$$

Note that any set of generators can be partitioned into support-disjoint subsets (2), and this gives the direct product decomposition (4) when $H_i = \langle X_i \rangle$. However, this decomposition is not unique (a geometric factor could be further decomposed) and depends on the choice of generators (for example, adding a generator $s_1 s_2$ to X, where $s_1 \in X_1$ and $s_2 \in X_2$ would force $X_1' = X_1 \cup X_2$ and $H_1' = H_1 \times H_2$, a coarser geometric decomposition). By requiring X to be essential (see below), the geometric decomposition (4) is unique (up to permutation of the factors H_i), and the finest possible (each H_i cannot be further decomposed into a direct product of non-trivial support-disjoint subgroups), see Proposition 2.1 in[7]. In particular, the geometric decomposition and symmetric motifs are well-defined.

A set of generators X is *essential* if (1) it does not contain the identity 1 (trivial permutation); (2) $s = gh \in X$ with g and h support-disjoint implies $g = 1$ or $h = 1$; and (3) if $X' \subset X$ generates $H_1 \times H_2$ support-disjoint (as sets), then $X' = X_1 \cup X_2$ (necessarily support-disjoint) with X_i generating H_i. Graph automorphism algorithms such as NAUTY produce essential sets of generators (Theorem 2.34 in[16]) and so does SAUCY, at least in practice (see the Network Symmetry Computation section below). Having said this, all the results in the Main Text are independent of whether the geometric decomposition is indeed the finest possible or not, and hence, any decomposition into a (possibly large) number of small symmetric motifs can be used to exploit network symmetry as explained in this article.

In real-world networks, symmetry is localised at small subgraphs (the symmetric motifs) and thus generated by low degree vertices. The universality of power-law degree distributions[17] guarantees the abundance of such low degree vertices and justifies the large (in absolute terms) automorphism groups observed in real-world networks. Indeed, the authors in[18] show how a Barabasi-Albert model reproduces the characteristic symmetry found in real-world networks. On the other hand, without a power-law degree distribution, we can find real-world networks with trivial automorphism group, such as the Blue Brain connectome[19], which has hardly any low degree vertices (neurons). Even with a power-degree distribution, we found networks with relatively low symmetry, such as CaliforniaRoads, where only 4% of all the vertices participate in any symmetry (Table 1 in the Main Text).

Basic Symmetric Motifs

In theory, any finite group can be the automorphism group of a graph[1]. In real-world networks, however, $\mathrm{Aut}(\mathscr{G})$ is the direct (occassionally, semidirect) product of symmetric groups S_n[7]. Moreover, most symmetric motifs are made of orbits of the same size with the geometric factor realising every possible permutation of the vertices in each orbit. Formally, a symmetric motif $\mathscr{M}_i$ is called *basic* if it consists of one or more orbits of n vertices, and the symmetric factor H_i is the symmetric group S_n realising all the $n!$ permutations of each orbit. (We then call $\mathscr{M}_i$ a basic symmetric motif of *type* n.) Most (over 90%) of symmetric motifs in our test networks are basic (Table 1 in the Main Text), similar to the results in[7].

Basic symmetric motifs (BSMs) have a very constrained structure:

- every orbit is a complete or an empty graph;
- each pair of orbits in the same symmetric motif can only be connected in one of four possible ways (each vertex in one orbit connected to either all, none, one, or $n-1$ vertices in the other orbit);
- every vertex not in the symmetric motif joins either all or none of the vertices in an orbit, for each orbit.

(For a proof, see[8], or the more general Theorem 2 below.) It is easy to show (or see Theorem 2) that the third property holds for any symmetric motif, basic or not. Also note that for a BSM with only two orbits, the 'all-to-all' or 'none-to-none' connectivity need not be considered, since in that case the orbits can be classified as two separate BSMs of one orbit each.

Non-basic symmetric motifs are called *complex*; they are rare in real-world networks and can be studied on a case-by-case basis. Typical complex symmetric motifs are branched trees (M_7 in Figure 1 in the Main Text), among others[7].

Weighted and Directed Networks

The adjacency matrix of a network can encode arbitrary weights and directions, as explained in the Main Text, making a general $n \times n$ real matrix A the adjacency matrix of some (weighted, directed) network. The definition of automorphism group, geometric decomposition, geometric factor, symmetric motif and orbit, and their properties, as they are defined only in terms of A, carry verbatim to arbitrarily weighted and directed networks. In this setting, a symmetry (automorphism), respects not only adjacency, but weights and directions. In particular, the automorphism group is smaller than (a subgroup of) the automorphism group of the underlying undirected, unweighted network. By introducing edge weights or directions, some symmetries will disappear, removing (and occasionally subdividing) geometric factors, symmetric motifs and orbits.

Theorem 1. *Let $A_{\mathrm{w}} = (w_{ij})$ be the adjacency matrix of an arbitrarily weighted and directed network $\mathscr{G}_{\mathrm{w}}$, and $A = (a_{ij})$ the adjacency matrix of the underlying undirected and unweighted network $\mathscr{G}$, that is, $a_{ij} = \mathrm{sgn}(|w_{ij}| + |w_{ji}|)$. Consider the geometric decomposition*

$$\begin{aligned} \mathrm{Aut}(\mathscr{G}) &= H_1 \times \ldots \times H_m, \quad \textit{respectively} \\ \mathrm{Aut}(\mathscr{G}_{\mathrm{w}}) &= H'_1 \times \ldots \times H'_{m'}, \end{aligned}$$

with symmetric motifs with vertex sets $M_1, \ldots, M_m$, respectively $M'_1, \ldots, M'_{m'}$. Then for every $1 \le i \le m'$ there is a unique $1 \le j \le m$ such that $H'_i \subseteq H_j$ and, consequently, $M'_i \subseteq M_j$. In other words, the geometric decomposition of $\mathscr{G}_{\mathrm{w}}$ is a refinement of that of $\mathscr{G}$. Similarly, each vertex orbit in $\mathscr{G}_{\mathrm{w}}$ is a subset of a vertex orbit in $\mathscr{G}$.

Proof. First we show that the automorphism group of $\mathscr{G}_{\mathrm{w}}$ is a subgroup of the automorphism group of $\mathscr{G}$. If $\sigma\colon V \to V$ is a permutation of the vertices, then

$$w_{\sigma(i)\sigma(j)} = w_{ij} \implies a_{\sigma(i)\sigma(j)} = a_{ij}$$

by considering two cases: $w_{ij} \neq 0$ implies $w_{\sigma(i)\sigma(j)} \neq 0$ which gives $a_{ij} = a_{\sigma(i)\sigma(j)} = 1$; $w_{ij} = 0$ implies $w_{\sigma(i)\sigma(j)} = 0$ which gives $a_{ij} = a_{\sigma(i)\sigma(j)} = 0$ (note $w_{ij} \neq 0 \iff a_{ij} = 1$). Hence $\mathrm{Aut}(\mathscr{G}_{\mathrm{w}}) \subset \mathrm{Aut}(\mathscr{G})$, which immediately gives the result on orbits.

For the geometric decomposition, we choose essential sets of generators S, respectively S', of $\mathrm{Aut}(\mathscr{G})$, respectively $\mathrm{Aut}(\mathscr{G}_{\mathrm{w}})$, with support-disjoint partitions

$$X = X_1 \cup \ldots \cup X_m, \text{ respectively } X' = X'_1 \cup \ldots \cup X'_{m'}.$$

It is enough to prove the statement for these sets: given i, there is unique j such that $X'_i \subseteq X_j$. Let $x' \in X'_i \subseteq \mathrm{Aut}(\mathscr{G}_{\mathrm{w}}) \subseteq \mathrm{Aut}(\mathscr{G})$ thus we can write $x' = h_1 \cdot \ldots \cdot h_m$ with $h_k \in H_k = \langle X_k \rangle$. Since X' is an essential set of generators, there is an index j such that $h_k = 1$ (the identity, or trivial permutation) for all $k \neq j$, so that $x' = h_j$. Given any other $y' \in X'_i$, the same argument gives $y' = h_l$ for some $1 \le l \le m$. We claim $j = l$, as follows. The partition of X, respectively X', above are the equivalence classes of the equivalence relation generated by $\sigma \sim \tau$ if σ and τ are not support-disjoint permutations. Since x', y' are in the same equivalence class, so are h_j and h_l and thus $j = l$. □

The same result applies to networks with other additional structure, not necessarily expressed in terms of the adjacency matrix, such as arbitrary vertex or edge labels, by restricting to automorphisms preserving the additional structure. We obtain fewer symmetries, and a refinement of the geometric decomposition, symmetric motifs, and orbits as above. The results in this paper, although applicable in theory, become less useful in practice as further restrictions are imposed, reducing the number of available network symmetries.

SI Structural Network Measures

As explained in the Main Text, a (pairwise) structural network measure F applied to a network $\mathscr{G}$ inherits all the symmetries of $\mathscr{G}$. It is possible that further symmetries will appear in $F(\mathscr{G})$, the network representation of F (a trivial example would be $F(i,j)=c$, a constant, for all i,j), but so rare in practical situations that we only consider symmetries directly inherited from $\mathscr{G}$. (A few additional symmetries would only result on a slightly finer geometric decomposition of $F(\mathscr{G})$.) Consequently, we can assume that $F(\mathscr{G})$ has the same automorphism group as $\mathscr{G}$, $\mathrm{Aut}(F(\mathscr{G}))=\mathrm{Aut}(\mathscr{G})$, and, in particular, the same geometric decomposition and geometric factors H_i. The symmetric motifs in $F(\mathscr{G})$ have the same vertex sets and orbits, but with edges weighted by $F(i,j)$, and possibly directed if $F(i,j)\neq F(j,i)$. As $F(\mathscr{G})$ inherits all the symmetries of $\mathscr{G}$, the basic symmetric motifs in $F(\mathscr{G})$ still have a very constrained structure, explained below (cf. Theorem 1 in Methods).

Theorem 2. *Let M be the vertex set of a BSM of a network $\mathscr{G}$ (that is, a symmetric motif made of one or more orbits of size n with geometric factor the symmetric group S_n realising all the permutation in each orbit), and F a structural network measure. Then the graph induced by M in $F(\mathscr{G})$ is a BSM of $F(\mathscr{G})$, and satisfies:*

(i) for each orbit $\Delta=\{v_1,\dots,v_n\}$, there are constants α and β such that the orbit internal connectivity is given by $\alpha=F(v_i,v_j)$ for all $i\neq j$ and $\beta=F(v_i,v_i)$ for all i;

(ii) for every pair of orbits Δ_1 and Δ_2, there is a labelling $\Delta_1=\{v_1,\dots,v_n\}$, $\Delta_2=\{w_1,\dots,w_n\}$ and constants γ_1, γ_2, δ_1, δ_2 such that $\gamma_1=F(v_i,w_j)$, $\gamma_2=F(w_j,v_i)$, $\delta_1=F(v_i,w_i)$, and $\delta_2=F(w_i,v_i)$, for all $i\neq j$;

(iii) every vertex v not in the BSM is joined uniformly to all the vertices in each orbit $\{v_1,\dots,v_n\}$ in the BSM, that is, $F(v,v_i)=F(v,v_j)$ and $F(v_i,v)=F(v_j,v)$ for all i,j.

Moreover, property (iii) holds in general for any symmetric motif.

Note that, if $\mathscr{G}$ is undirected and F is symmetric, $\gamma_1=\gamma_2$ and $\delta_1=\delta_2$ and, in the terminology of the Main Text, each orbit is a (α,β)-uniform graph $K_n^{\alpha,\beta}$ and each pair of orbits form a (γ,δ)-uniform join, explaining Figure 2(a-b) in the Main Text. Also note that, for unweighted, undirected graphs without loops, we recover the statements in the previous section: every orbit is an empty or complete graph ($\beta=0$, $\alpha=0,1$), and every pair of orbits are joined in one of the four possible ways ($\gamma,\delta=0,1$).

Proof of Theorem. As $F(\mathscr{G})$ inherits all the symmetries of $\mathscr{G}$, M has the same orbit decomposition and the geometric factor is S_n acting in the same way, hence M induces a BSM in $F(\mathscr{G})$ too. For the internal connectivity, note that every permutation of the vertices v_i is realisable. Thus, given arbitrary $1\leq i,j,k,l\leq n$, we can find $\sigma\in Aut(\mathscr{G})$ such that $\sigma(v_k)=v_i$ and, if $j\neq i$ and $l\neq k$, additionally satisfies $\sigma(v_l)=v_j$. This gives

$$F(v_i,v_j)=F(\sigma(v_k),\sigma(v_l))=F(v_k,v_l),$$

as F is a structural network measure. The other case, $i=j$ and $k=l$, gives

$$F(v_i,v_i)=F(\sigma(v_k),\sigma(v_k))=F(v_k,v_k).$$

For the orbit connectivity result (ii), we generalise the argument in [20, p.48] to weighted directed graphs with symmetries, particularly $F(\mathscr{G})$. We assume some basic knowledge and terminology about group actions[15] and symmetric groups S_n. Given two orbits $\Delta_1=\{v_1,\dots,v_n\}$ and $\Delta_2=\{w_1,\dots,w_n\}$ and $1\leq i\leq n$, define

$$\Gamma_i=\left\{w_j\in\Delta_2\,|\,F(v_i,w_j)\neq 0\right\},$$

the vertices in Δ_2 joined to v_i in $F(\mathscr{G})$. If a finite group G acts on a set X, the stabiliser of a point $G_x=\{g\in G\,|\,gx=x\}$ is a subgroup of G of index $[G:H]=\frac{|G|}{|H|}$ equals to the size of the orbit of x. Hence the stabilisers G_{v_i} or G_{w_j} are subgroups of S_n of index n, for all i,j. The group S_n has a unique, up to conjugation, subgroup of index n if $n\neq 6$. In this case, G_{v_1} is conjugate to G_{w_1} so $G_{v_1}=\sigma G_{w_1}\sigma^{-1}=G_{\sigma w_1}$ for some $\sigma\in S_n$. Relabelling σw_1 as w_1 we have $G_{v_1}=G_{w_1}$. Similarly, we can relabel the remaining vertices in Δ_2 so that $G_{v_i}=G_{w_i}$ for all i: write $v_2=\sigma_2 v_1$, $v_3=\sigma_3 v_1$, … and relabel $w_2=\sigma_2 w_1$, $w_3=\sigma_3 w_1$, …, noticing there cannot be repetitions as $\sigma_k w_1 = \sigma_l w_1$ for $k \neq l$ implies $\sigma_k \sigma_l^{-1} \in G_{w_1} = G_{v_1}$ and thus $v_k = \sigma_k v_1 = \sigma_l v_1 = v_l$, a contradiction. Fix $1 \leq i \leq n$. The stabiliser G_{v_i} fixes v_i but it may permute vertices in Δ_2. In fact, the set Γ_i above must be a union of orbits of G_{v_i} on Δ_j: if $w \in \Gamma_i$ and $\sigma \in G_{v_i}$ then

$$0 \neq F(v_i, w) = F(\sigma v_i, \sigma w) = F(v_i, \sigma w)$$

so σw also belongs to Γ_i. The orbits of $G_{v_i} = G_{w_i}$ in Δ_2 are $\{w_i\}$ and $\Delta_2 \setminus \{w_i\}$, as G_{w_i} fixes w_i and freely permutes all other vertices in Δ_2. The case $n = 6$ is similar, except that S_6 has two conjugacy classes of subgroups of index 6, one as above, and the other a subgroup acting transitively on the 6 vertices, which gives a unique orbit Δ_2. In all cases, the set $\Delta_2 \setminus \{w_i\}$ is part of an G_{v_i}-orbit, which gives the connectivity result, as follows. Fix $1 \leq i \leq n$. For $1 \leq j, k \leq n$ different from i, the vertices w_j and w_k are in the same G_{v_i}-orbit so there is $\sigma \in G_{v_i}$ with $\sigma w_j = w_k$ and, therefore,

$$F(v_i, w_j) = F(\sigma v_i, \sigma w_j) = F(v_i, w_k).$$

The argument is general, so we have shown $a_i = F(v_i, w_j)$ is constant for all $j \neq i$. It is enough to show $a_i = a_1$ for all i. Choose $j \neq i$, then

$$a_i = F(v_i, w_j) = F(\sigma_i v_1, \sigma_j w_1) = F(v_1, \sigma_i^{-1} \sigma_j w_1) = a_1$$

as long as $\sigma_i^{-1} \sigma_j w_1 \neq w_1$, which cannot happen as otherwise $\sigma_i^{-1} \sigma_j \in G_{w_1} = G_{v_1}$ implies $\sigma_i^{-1} \sigma_j v_1 = v_1$ or $v_j = \sigma_j v_1 = \sigma_i v_1 = v_i$, a contradiction. Hence we have shown $F(v_i, w_j)$ is a constant, call it γ_1, for all $i \neq j$. In addition,

$$F(v_i, w_i) = F(\sigma_i v_1, \sigma_i w_1) = F(v_1, w_1)$$

is also a constant, call it δ_1, for all i. The cases $\gamma_2 = F(w_j, v_i)$ and $\delta_2 = F(w_i, v_i)$ are identical, reversing the roles of Δ_1 and Δ_2.

Property (iii) holds for any symmetric motif, not necessarily basic, as follows. By the definition of orbit, for each pair i, j we can find an automorphism σ in the geometric factor such that $\sigma(v_j) = v_i$. Since v is not in the support of that geometric factor, it is fixed by σ, that is, $\sigma(v) = v$. Therefore

$$F(v, v_i) = F(\sigma(v), \sigma(v_j)) = F(v, v_j),$$

and similarly $F(v_i, v) = F(v_j, v)$. □

Consequently, every orbit in a BSM is characterised by three parameters: its size n, the connectivity between any pair of distinct vertices α, and the connectivity of a vertex with itself β. In the terminology and notation of the Main Text, every orbit is a (α, β)-*uniform graph*, $K_n^{\alpha,\beta}$, the graph with adjacency matrix

$$A_n^{\alpha,\beta} = \alpha C_n + \beta I_n \tag{10}$$

where C_n is the adjacency matrix of a complete graph (0 diagonal and 1 off-diagonal entries), and I_n the identity matrix. For example, a $(1,0)$-uniform graph is a complete graph, and a $(0,0)$-uniform graph is an empty graph.

Similarly, every pair of orbits in the same BSM form a $(\gamma_1, \delta_1, \gamma_2, \delta_2)$-*uniform join* of two uniforms graphs $K_n^{\alpha_1,\beta_1}$ and $K_n^{\alpha_2,\beta_2}$, that is, the graph with adjacency matrix, possibly after a suitable reordering of the vertices,

$$A = \left(\begin{array}{c|c} A_n^{\alpha_1,\beta_1} & A_n^{\gamma_1,\delta_1} \\ \hline A_n^{\gamma_2,\delta_2} & A_n^{\alpha_2,\beta_2} \end{array}\right), \tag{11}$$

where the submatrices are given as per (10).

This constrained structure will be important when we discuss the quotient network, and the spectral signatures of symmetry.

SI Redundancy in Network Measures

Quotient Network

If A is the $n \times n$ adjacency matrix of a graph $\mathscr{G}$, the *quotient network* with respect to a partition of the vertex set $V = V_1 \cup \ldots \cup V_m$ is the graph $\mathscr{Q}$ with $m \times m$ adjacency matrix the *quotient matrix* $Q(A) = (b_{kl})$ defined by

$$b_{kl} = \frac{1}{|V_k|} \sum_{\substack{i \in V_k \\ j \in V_l}} a_{ij}, \tag{12}$$

the average connectivity from a vertex in V_k to vertices in V_l. That is, the quotient amalgamates the vertices of each V_k into a single vertex, and have edges representing average connectivity.

There is an explicit matrix equation for the quotient. Consider the $n \times m$ *characteristic matrix* S of the partition, that is, $[S]_{ik} = 1$ if $i \in V_k$, and zero otherwise, and the diagonal matrix $\Lambda = \mathrm{diag}(n_1, \dots, n_m)$, where $n_k = |V_k|$. Then

$$Q(A) = \Lambda^{-1} S^T A S. \tag{13}$$

They are other ways of taking the average in (12), such as $1/n_l$ or $1/(\sqrt{n_k}\sqrt{n_l})$ instead of $1/n_k = 1/|V_k|$, giving slightly different, but spectrally equivalent, quotient matrices $S^T A S \Lambda^{-1}$ respectively $\Lambda^{-1/2} S^T A S \Lambda^{-1/2}$. We refer to our choice as the *left quotient* (matrix or network), and the other two as the *right* respectively *symmetric quotient*, writing $Q_\mathrm{l}(A)$, $Q_\mathrm{r}(A)$ or $Q_\mathrm{s}(A)$ when necessary. (Note that $Q_\mathrm{r}(A) = Q_\mathrm{l}(A)^T$ if A is symmetric.) Occasionally, it will be convenient to ignore weights and directions: we call *unweighted quotient* to the underlying unweighted, undirected quotient network, and, if we also remove self-loops (so that two orbits are connected if they are distinct and at least one vertex of one orbit is connected to a vertex of the other orbit) this is the *quotient skeleton* in Figure 1 in the Main Text (the s-skeleton in[9]). Finally, in the context of quotient networks, we call $\mathscr{G}$ the *parent* network of $\mathscr{Q}$.

Note that the (left) quotient $\mathscr{Q}$ is a directed and weighted network even if the parent network $\mathscr{G}$ is not: $b_{ij} \neq b_{ji} \notin \{0,1\}$ in general. However, $Q_\mathrm{l}(A)$ is spectrally equivalent to the symmetric matrix $Q_\mathrm{s}(A)$, hence in particular is has real eigenvalues if $\mathscr{G}$ is undirected (SI Spectral Signatures of Symmetry).

A natural quotient in the context of symmetries is given by the partition of vertices into orbits, Eqs. (7) and (9). Note that this partition is finer than the one associated to the geometric decomposition, as each symmetric motif consists of one or more orbits, and that each fixed point in V_0 (the asymmetric core) becomes its own orbit, that is, it is not identified with any other vertex in the quotient. From now on, we will implicitly refer to this quotient unless stated otherwise. Occasionally, we will consider the quotient with respect to a subgroup of $\mathrm{Aut}(\mathscr{G})$ (cf. Theorem 1), or only certain symmetric motifs (e.g. basic quotient, explained below).

A crucial property of the quotient with respect to orbits is that it removes all the symmetries from the network: if $\sigma(i) = j$, then i and j are in the same orbit and hence represented by the same vertex in the quotient network, which is then fixed by σ. (Although new symmetries might appear in the quotient, these are rare and would not be symmetries of the parent network, and hence of no interest to us.) We can, consequently, infer and quantify properties arising from redundancy alone by comparing a network with its quotient. This is the approach taken in[7] for spectral properties, which we will generalise to undirected, arbitrarily weighted, networks with symmetries, such as the network representation of a network measure (SI Spectral Signatures of Symmetry).

As we explain in the the Main Text, we can also exploit the quotient in the context of network structural measures for compression (storage savings), and for computational reduction (time and memory savings).

Quotient Compression and Recovery

The network representation of a structural network measure $F(\mathscr{G})$ inherits all the symmetries of $\mathscr{G}$, which present themselves as redundancies, namely as repeated values. For instance, for each symmetric motif $\mathscr{M}$, the values $F(u,v)$ are constant, for each u not in $\mathscr{M}$ and each v in the same orbit of $\mathscr{M}$. If the network has n vertices, this means $(n-k)l \approx nl$ repeated values for each orbit of size l in a symmetric motif motif of size k (typically k and l are very small). The internal connectivity of a symmetric motif can also be efficiently encoded, for instance each orbit in a BSM can be recovered from three values, its size n and constants α and β, and the connectivity between pairs of orbits in the same BSM from two or four values (undirected/directed case), and a permutation of the second orbit (Theorem 2). All this can be exploited in a compression algorithm that eliminates redundancies induced by symmetries in an arbitrary network structural measure, as we explain next. We first observe that this is most useful for full measures, since for sparse structural measures, the values $F(u,v)$ as above are mostly 0 and hence a sparse representation of $F(\mathscr{G})$ will account for most of the aforementioned redundancies.

Average Compression and Recovery

If we are not interested in the internal symmetric motif connectivity, or it can be recovered easily by other means (e.g. locally one motif at a time), a simple algorithm (Algorithm 1) compresses and recovers all but the internal symmetric motif connectivity, which is replaced by the average connectivity, as the next result (Theorem 2 in Methods) shows.

Theorem 3. *Let $A = (a_{ij})$ be the $n \times n$ adjacency matrix of a (possibly directed and weighted) network with vertex set V. Let S be the $n \times m$ characteristic matrix of the partition of V into orbits of the automorphism group of the network, and Λ the diagonal matrix of column sums of S. Define $B = S^T A S$ and $A_{\mathrm{avg}} = R B R^T = (\bar{a}_{ij})$ where $R = S\Lambda^{-1}$. Then,*

(i) if $i, j \in V$ belong to different symmetric motifs, $\bar{a}_{ij} = a_{ij}$.

(ii) if $i, j \in V$ belong to orbits $i \in \Delta_1$ and $j \in \Delta_2$ in the same symmetric motif,

$$\bar{a}_{ij} = \frac{1}{|\Delta_1|}\frac{1}{|\Delta_2|}\sum_{\substack{u\in\Delta_1\\ v\in\Delta_2}} a_{uv}. \tag{14}$$

Before proving this statement, we make a few observations. The column sums of S equal the sizes of the vertex partition sets, that is, the matrix Λ in the statement is the same as in the definition of quotient matrix, (13), and can be obtained easily from S. The matrix S is very sparse and can be stored very efficiently, as it has at most n non-zero elements (each row has a unique non-zero entry). Case (i) covers the vast majority of vertex pairs (external edges) for a network measure (see ext_s and int_f in Table 1 in the Main Text). In (ii), the case $\Delta_1 = \Delta_2$ is allowed. The matrix $B = S^T AS$ is symmetric with integer entries if A is too, hence generally easier to store than $Q_l(A) = \Lambda^{-1}S^T AS$. However, the theorem still holds for the left, right and symmetric quotient, as

$$A_{\text{avg}} = S\Lambda^{-1}S^T AS\Lambda^{-1}S^T = S\,Q_l(A)\,\Lambda^{-1}S^T = S\Lambda^{-1}\,Q_r(A)\,S^T = S\Lambda^{-1/2}\,Q_s(A)\,\Lambda^{-1/2}S^T. \tag{15}$$

Proof of Theorem. Let $V = \Delta_1 \cup \ldots \cup \Delta_m$ be the partition into orbits, and write $n_k = |\Delta_k|$. Clearly, the row sums of S equals $n_1, \ldots, n_m$. Writing $[M]_{ij}$ for the (i, j)-entry of a matrix M, matrix multiplication gives

$$[R]_{ik} = \sum_l [S]_{il}[\Lambda^{-1}]_{lk} \overset{l=k}{=} [S]_{ik}\frac{1}{n_k} = \begin{cases} \frac{1}{n_k} & \text{if } i \in \Delta_k, \\ 0 & \text{otherwise.} \end{cases}$$

Similarly, assuming $i \in \Delta_k$ and $j \in \Delta_l$, we have

$$\bar{a}_{ij} = [RBR^T]_{ij} = \sum_{\alpha,\beta}[R]_{i\alpha}[B]_{\alpha\beta}[R]_{j\beta} = \frac{1}{n_k}\frac{1}{n_l}[B]_{kl} = \frac{1}{n_k}\frac{1}{n_l}\sum_{\substack{u\in\Delta_k\\ v\in\Delta_l}} a_{uv}.$$

This expression reduces to a_{ij} if the orbits belong to different symmetric motifs, since in this case all the summands in $\sum_{u\in\Delta_k, v\in\Delta_l} a_{uv}$ are equal to one another. Indeed, given $i_1, i_2 \in \Delta_k$ and $j_1, j_2 \in \Delta_l$, we can find, by the definition of orbit and symmetric motif, automorphisms σ and τ such that $\sigma(i_1) = i_2$ while fixing j_1, and $\tau(j_1) = j_2$ while fixing i_1. This gives

$$a_{i_1 j_1} = a_{\tau\sigma(i_1)\tau\sigma(j_1)} = a_{\tau(i_2)\tau(j_1)} = a_{i_2 j_2}.$$

□

A similar proof to case (i) above shows that, if $Q_l(A) = (b_{kl})$ is the left quotient, then

$$a_{ij} = \frac{1}{n_l}b_{kl} = \frac{1}{n_k}b_{lk},$$

where vertex i, respectively j, belongs to an orbit of size n_k, respectively n_l. As in the proof, each summand in $\sum_{u\in\Delta_k, v\in\Delta_l} a_{uv}$ is constant and, therefore,

$$b_{kl} = \frac{1}{|\Delta_k|}\sum_{u\in\Delta_k, v\in\Delta_l} a_{uv} = \frac{1}{|\Delta_k|}|\Delta_k||\Delta_l|a_{ij} = |\Delta_l|a_{ij},$$

which gives the first equality. The second equality follows from observing that, although $Q_l(A) = (b_{kl})$ is non-symmetric, $b_{kl} = \frac{n_l}{n_k} b_{lk}$ for all k, l.

As in the Main Text, we quantify the redundancy in an arbitrary sparse, respectively full, network measure using

$$c_{\text{sparse}} = \frac{m_{\mathcal{Q}}}{m_{\mathcal{G}}} = \widetilde{m}_{\mathcal{Q}} \text{ and } c_{\text{full}} = \frac{n_{\mathcal{Q}}^2}{n_{\mathcal{G}}^2} = \widetilde{n}_{\mathcal{Q}}^2.$$

These are therefore the compression ratios obtained by using the quotient to represent a sparse, respectively full, network measure on $\mathcal{G}$. The values of $c_{\text{sparse}} =$ and c_{full} for our test networks are given in Table 1 in the Main Text.

In summary, we have a simple compression/decompression procedure (Algorithms 1 and 2 below, or 2 and 3 in Methods) that eliminates the symmetry-induced redundancies in a network measure, and achieves exact recovery for external edges (the vast majority of edges, or vertex pairs), and average recovery for internal edges.

Input: adjacency matrix A, characteristic matrix S
Output: quotient matrix B

$B \leftarrow S^T AS$

Algorithm 1: Average symmetry compression.

Input: quotient matrix B, characteristic matrix S
Output: adjacency matrix A_{avg}

$\Lambda \leftarrow \texttt{diag}(\texttt{sum}(S))$
$R \leftarrow S\Lambda^{-1}$
$A_{\text{avg}} \leftarrow RBR^T$

Algorithm 2: Average symmetry decompression.

Lossless Compression and Recovery

We can achieve lossless compression if we recover the exact internal motif connectivity as well. This can be done by exploiting the structure of BSMs, which account for most of the symmetry in real-world networks. If the motif is basic, we can preserve the exact parent network connectivity in an annotated quotient, as follows. Each orbit in a BSM is a uniform graph $K_n^{\alpha,\beta}$ which appears in the quotient as a single vertex with a self-loop weighted by $(n-1)\alpha+\beta$ (Figure 2 (c, top) in Main Text). Hence if we annotate this vertex in the quotient by not only n but also α, or β, we can recover the internal connectivity. Similarly, the connectivity between two orbits in the same symmetric motif is given by two parameters γ, δ (undirected case) and appears in the quotient as an edge weighted $(n-1)\gamma+\delta$ (Figure 2 (c, bottom) in Main Text) and thus can also be recovered from a quotient with edges annotated by γ, or δ.

There is no general formula for an arbitrary non-basic symmetric motif, and for those we can either record their internal connectivity separately, or, alternatively, leave them unchanged in the quotient, by taking a partial quotient, the *basic quotient*, written $\mathscr{Q}_{\text{basic}}$, with respect to the partition of the vertex set into orbits in BSMs only (vertices in non-basic symmetric motifs become fixed points and become part of the asymmetric core). The annotated (as above) basic quotient achieves most of the symmetry reduction in a typical empirical network ($\tilde{n}_{\mathscr{Q}}^{\text{basic}} \approx \tilde{n}_{\mathscr{Q}}$, $\tilde{m}_{\mathscr{Q}}^{\text{basic}} \approx \tilde{m}_{\mathscr{Q}}$, see SI Table 1) while retaining all the parent network connectivity. However, to maintain the same vertex labelling as in the parent network, we also need to record, for each pairs of orbits in the same symmetric motif, the corresponding permutation of the second orbit (Theorem 2(ii)), as otherwise we loose vertex identity on each orbit.

Similarly to average compression, we introduce compression ratios for lossless compression with respect to the basic quotient

$$c_{\text{sparse}}^{\text{basic}} = \frac{m_{\mathscr{Q}_{\text{basic}}}}{m_{\mathscr{G}}} = \tilde{m}_{\mathscr{Q}}^{\text{basic}} \quad \text{and} \quad c_{\text{full}}^{\text{basic}} = \frac{n^2_{\mathscr{Q}_{\text{basic}}}}{n^2_{\mathscr{G}}} = (\tilde{n}_{\mathscr{Q}}^{\text{basic}})^2, \tag{16}$$

and note that typically $c_{\text{sparse}}^{\text{basic}} \approx c_{\text{sparse}}$ and $c_{\text{full}}^{\text{basic}} \approx c_{\text{full}}$ for a real-world network (SI Table 1).

Algorithms for lossless compression and recovery with vertex identity based on the basic quotient are described below (Algorithms 3 and 4), and MATLAB implementations for BSMs up to two orbits are available at[21]. The results reported in Figure 3 in the Main Text are with respect to these implementations, and the actual compression ratios reported include the size of the annotation data for lossless compression with vertex identity (a very small fraction of the size of the quotient in practice, adding at most 0.02% to the basic full compression ratio in all our test cases).

SI Computational Reduction

Network symmetries can also be used to reduce the computational time of a network measure F. Recall that a sparse, respectively full, network measure consists on at most $m_{\mathscr{G}}$, respectively $n^2_{\mathscr{G}}$, non-zero values $F(i,j)$. Since $F(i,j) = F(\sigma(i),\sigma(j))$ for each $\sigma \in \text{Aut}(\mathscr{G})$, we essentially need to evaluate F on $m_{\mathscr{Q}}$, respectively $n^2_{\mathscr{Q}}$, orbit representatives only. This amounts to a computational reduction ratio between c_{sparse} and c_{full}. Of course, this assumes that the calculation on each pair of vertices is independent of one another, which is often not the case. Moreover, the calculation on each pair $F(i,j)$ is still performed on the whole network $\mathscr{G}$. Next we investigate whether we could perform the calculation on the quotient instead, which would lead to greater computational gains by evaluating F on a smaller graph.

We call a network measure F *partially quotient recoverable* if it can be applied to a quotient network and all the external edges of $F(\mathscr{G})$ can be recovered from $F(\mathscr{Q})$, for all networks $\mathscr{G}$. (Here $\mathscr{Q}$ is a quotient of $\mathscr{G}$, possibly basic or annotated.)

Input: adjacency matrix A, characteristic matrix for the basic quotient S, list of BSMs motifs
Output: quotient matrix B, annotation structure a

$B \leftarrow S^T AS$

extract orbits from S
foreach orb *in* orbits **do**
  rep $\leftarrow$ `min`(orb)
  $\beta \leftarrow A(\mathsf{rep}, \mathsf{rep})$
  store β in annotation structure a
end

$k_{\max} \leftarrow$ `max`(`size`(motifs)) maximal number of orbits in a motif
for $k \leftarrow 2$ **to** k_{max} **do**
  extract k-BSM (list of BSMs with k orbits) from motifs
  foreach bsm *in* k-BSM **do**
    foreach *pairs of distinct orbits* V_1, V_2 *in* bsm **do**
      compute δ and permutation of V_2 perm such that $A(k, \mathsf{perm}(k)) = \delta$ for all $k \in V_1$
      store orbit numbers (with respect to S), δ and perm in annotation structure a
    end
  end
end

Algorithm 3: Lossless symmetry compression.

Input: quotient matrix B, characteristic matrix S, annotation structure a
Output: adjacency matrix A

$\Lambda \leftarrow \mathtt{diag}(\mathtt{sum}(S))$
$R \leftarrow S\Lambda^{-1}$
$A \leftarrow RBR^T$

extract orbits from S
foreach orb *in* orbits **do**
  $n \leftarrow \mathtt{size}(\mathsf{orb})$
  extract β from a
  compute α from B, β and n (using $[B]_{\mathsf{orb},\mathsf{orb}} = n((n-1)\alpha + \beta)$)
  construct adjacency matrix of the orbit $A_n^{\alpha,\beta}$
  $A(\mathsf{orb}, \mathsf{orb}) \leftarrow A_n^{\alpha,\beta}$
end

extract pairs of orbits in the same BSM from a
foreach (V_1, V_2) *in* pairs **do**
  $n \leftarrow \mathtt{size}(V_1)$
  extract δ, perm from a
  compute γ from B, δ and n (using $[B]_{V_1,V_2} = n((n-1)\gamma + \delta)$)
  construct matrix $A_n^{\gamma,\delta}$
  $A(V_1, \mathsf{perm}) \leftarrow A_n^{\gamma,\delta}$
  $A(\mathsf{perm}, V_1) \leftarrow A_n^{\gamma,\delta}$
end

Algorithm 4: Lossless symmetry decompression.

Since the quotient averages the network connectivity, we can often recover the average values of F between orbits as well. We call F *average quotient recoverable* if, in addition to external edges, the average intra-motif edges can be recovered from $F(\mathscr{Q})$. A typical situation is when $F(\mathscr{Q})$ equals the quotient representation of F, $Q(F(\mathscr{G}))$, that is, in symbols,

$$F(\mathscr{Q}) = Q(F(\mathscr{G})). \tag{17}$$

(Recall that we can recover the external edges, and the average internal edges, of $F(\mathscr{G})$ from $Q(F(\mathscr{G}))$, for any version of the quotient, by Theorem 3 and (15)). We will show that communicability is average quotient recoverable (see the Communicability section below), and shortest path distance is partially, but not average, quotient recoverable (see the Shortest Path Distance below). Not every measure can be (partially) recovered from the quotient, for example the number of distinct paths between two vertices. Note that the word 'partially' can be misleading: typically almost all edges are external (see ext_s and int_f in Table 1 in the Main Text).

Finally, we call F *fully quotient recoverable* if the external and internal edges intra-motif edges, that is, the whole network representation $F(\mathscr{G})$, can be obtained from $F(\mathscr{Q})$, for every network $\mathscr{G}$. Technically, the parent network $\mathscr{G}$ can be fully recovered from an annotated basic quotient (see e.g. lossless compression above), so by full recoverability we mean, beyond evaluating $F(\mathscr{Q})$, a local (hence parallelizable) computation on each symmetric motif (typically a very small graph). For example, communicability is fully quotient recoverable (see Communicability), and the shortest path distances by reconstructing each motif at a time (see Shortest Path Distance). When quotient recoverability holds, there is a substantial computational reduction by evaluating F on a smaller graph. For instance, if F has time complexity $O(f(n,m))$, then we can evaluate F on $\mathscr{Q}$ on a fraction $f(\tilde{n}_{\mathscr{Q}}, \tilde{m}_{\mathscr{Q}})$ of the time. In Figure 4 in the Main Text, we report the computational time reduction of evaluating F on the quotient for the network measures, and the test networks, considered in this article.

SI Spectral Signatures of Symmetry

Preliminaries

Symmetry naturally produces high-multiplicity eigenvalues: if v is an eigenvector of the adjacency matrix A of a (possibly weighted, directed) network $\mathscr{G}$ with eigenvalue λ, so is Pv,

$$APv = PAv = \lambda Pv, \tag{18}$$

for any P permutation matrix representing an automorphism of $\mathscr{G}$, and v and Pv will generally be linearly independent. In[8], the authors formalise and quantify the effects of network symmetry on the spectrum and eigenvectors of real-world networks, predicting the observed 'peaks' in spectral density and showing that symmetry explains most of their multiplicity. Here we explain how the spectral results in[8] generalise to weighted networks with symmetries, such as the network representation $F(\mathscr{G})$ of a structural network measure F.

Let $A = (a_{ij})$ be the $n \times n$ adjacency matrix of an arbitrarily weighted and directed network, $B = (b_{kl})$ its $m \times m$ (left) quotient matrix with respect to the orbit partition $V_1 \cup \ldots \cup V_m$, and S the $n \times m$ characteristic matrix of the partition. A vector $v \in \mathbb{R}^n$, respectively $w \in \mathbb{R}^m$, can be seen as a vector on (the vertices of) the parent, respectively quotient, network. Then Sw is the vector w *lifted* to the parent network by repeating the entries on each V_j. Similarly, $S^T v$ (where S^T is the matrix transpose) is the vector v *projected* to the quotient by adding all its entries on each V_j. Note that $Sw \neq 0$ if $w \neq 0$. We call the vector v *orthogonal to the partition* if its projection $S^T v$ is zero, that is, the sum of the entries of v on each orbit is zero.

The partition into orbits satisfy important regularity conditions. A partition of the vertex set $V = V_1 \cup \ldots \cup V_m$ is *left equitable* if

$$\sum_{j \in V_l} a_{i_1 j} = \sum_{j \in V_l} a_{i_2 j} \quad \text{for all } i_1, i_2 \in V_k, \text{ for all } 1 \leq k,l \leq m, \tag{19}$$

that is, if the connectivity from a node in V_i to all nodes in V_j is independent of the chosen node in V_i. Similarly, the partition is *right equitable* if

$$\sum_{i \in V_k} a_{i j_1} = \sum_{i \in V_k} a_{i j_2} \quad \text{for all } j_1, j_2 \in V_l, \text{ for all } 1 \leq k,l \leq m. \tag{20}$$

Clearly, if A is symmetric, being left and right equitable are equivalent properties. Next we show matrix characterisations of left and right equitability, and that the partition into orbits is both left and right equitable.

Proposition 3. *Let $V = V_1 \cup \ldots V_m$ be a partition of the vertex set of a graph with adjacency matrix $A = (a_{ij})$, and let S be the characteristic matrix of the partition. Write $Q_l(A)$, respectively $Q_r(A)$ for the left, respectively right, quotient with respect to the partition.*

(i) The partition is left equitable partition if and only if $AS = SQ_l(A)$.

(ii) The partition is right equitable partition if and only if $A^T S = SQ_r(A)^T$.

(iii) The partition into orbits of the automorphism group $V = \Delta_1 \cup \ldots \cup \Delta_m$ is left and right equitable.

Proof. (i) Fix $1 \le i \le n$ and $1 \le k \le m$, and suppose $i \in V_l$. Then

$$[AS]_{ik} = \sum_{j \in V_k} a_{ij},$$

and, using the left equitable condition,

$$[SQ_\mathrm{l}(A)]_{ik} = [Q_\mathrm{l}(A)]_{lk} = \frac{1}{|V_l|} \sum_{\substack{i_1 \in V_l \\ j \in V_k}} a_{i_1 j} = \frac{1}{|V_l|} |V_l| \sum_{j \in V_k} a_{ij} = \sum_{j \in V_k} a_{ij}.$$

For the converse, note that $[AS]_{il}$ does not depend on i but on the orbit of i. Namely, given $i_1, i_2 \in V_k$,

$$\sum_{j \in V_l} a_{i_1 j} = [AS]_{i_1 l} = [Q_\mathrm{l}(A)]_{kl} = [AS]_{i_2 l} = \sum_{j \in V_l} a_{i_2 j}.$$

(ii) Similarly, with i, k and l as above,

$$[A^T S]_{ik} = \sum_{j \in V_k} a_{ji},$$

and, using the right equitable condition,

$$[SQ_\mathrm{r}(A)^T]_{ik} = [Q_\mathrm{r}(A)]_{kl} = \frac{1}{|V_l|} \sum_{\substack{i_1 \in V_l \\ j \in V_k}} a_{j i_1} = \frac{1}{|V_l|} |V_l| \sum_{j \in V_k} a_{ji} = \sum_{j \in V_k} a_{ji}.$$

For the converse, let $j_1, j_2 \in V_l$, then

$$\sum_{i \in V_k} a_{i j_1} = [A^T S]_{j_1 k} = [Q_\mathrm{r}(A)]_{kl} = [A^T S]_{j_2 k} = \sum_{i \in V_k} a_{i j_2}.$$

(iii) Given i_1 and i_2 in the same orbit Δ_k, choose an automorphism σ such that $\sigma(i_1) = i_2$. Then, since automorphisms respect the adjacency matrix,

$$\sum_{j \in \Delta_l} a_{i_1 j} = \sum_{j \in \Delta_l} a_{\sigma(i_1)\sigma(j)} = \sum_{j \in \Delta_l} a_{i_2 \sigma(j)} = \sum_{j \in \Delta_l} a_{i_2 j},$$

where the last equality follows from the fact that an element in a group permutes orbits, in this case, $\{j : j \in \Delta_l\} = \{\sigma(j) : j \in \Delta_l\}$. Hence the partition into orbits is left equitable. A similar argument shows that it is right equitable as well. □

Note that (iii) holds for any subset of orbits or any subgroup of the automorphism group (in particular, for the basic quotient).

Spectral Decomposition Theorem

The key spectral property of the quotient into orbits of the automorphism group (or any subgroup) is that its eigenvalues are a subset of the eigenvalues of the parent network. Namely, if v is a right (respectively left) eigenvector of $Q_\mathrm{l}(A)$ (respectively $Q_\mathrm{r}(A)$) with eigenvalue λ, then Sv (respectively vS^T) is a right (respectively left) eigenvector of A with the same eigenvalue. (These two statement are obviously equivalent if A is symmetric.) This follows immediately from the matrix characterisation of equitability above:

$$Q_\mathrm{l}(A)v = \lambda v \implies A(Sv) = SQ_\mathrm{l}(A)v = \lambda Sv, \text{ and} \tag{21}$$

$$vQ_\mathrm{r}(A) = \lambda v \implies A^T(Sv^T) = SQ_\mathrm{r}(A)^T v^T = S(vQ_\mathrm{r}(A))^T = \lambda Sv^T \iff (vS^T)A = \lambda vS^T. \tag{22}$$

(Recall that $Sv \neq 0$ if $v \neq 0$.) All in all, the spectrum of the quotient is a subset of the spectrum of the graph, with eigenvectors lifted from the quotient by repeating entries on orbits. Moreover, we can complete an eigenbasis with eigenvectors orthogonal to the partition (adding up to zero on each orbit), at least in the symmetric case.

Theorem 4. *Suppose that A is an $n \times n$ real symmetric matrix and B the $m \times m$ (left) quotient matrix with respect to an equitable partition $V_1 \cup \ldots \cup V_m$ of the set $\{1,2,\ldots,n\}$. Let S be the characteristic matrix of the partition. Then A has an eigenbasis of the form*

$$\{Sv_1,\ldots,Sv_m,w_1,\ldots,w_{n-m}\},$$

where $\{v_1,\ldots,v_m\}$ is any eigenbasis of B, and $S^T w_i = 0$ for all i.

Proof. Recall that $Sv \neq 0$ if $v \neq 0$ (S lifts the vector v from the quotient by repeating entries on each orbit) so the linear map

$$\mathbb{R}^m \to \mathbb{R}^n, v \mapsto Sv$$

has trivial kernel and hence it is an isomorphism onto its image. In particular, $\mathscr{B} = \{Sv_1,\ldots,Sv_m\}$ is also a linearly independent set, and they are all eigenvectors of A, since $AS = SB$ as the partition is equitable. To finish the proof we need to complete $\mathscr{B}$ to a basis $\{Sv_1,\ldots,Sv_m,w_1,\ldots,w_{n-m}\}$ such that each w_j is an A-eigenvector orthogonal to all Sv_i. As $\mathscr{B}$ is a basis of $\mathrm{Im}(S)$, this would imply $w_i \in \mathrm{Im}(S)^\perp = \mathrm{Ker}(S^T)$, giving $S^T w_i = 0$ for all i, as desired. Since A is diagonalisable, $\mathbb{R}^n$ decomposes as an orthogonal direct sum of eigenspaces, $\mathbb{R}^n = \bigoplus_\lambda E_\lambda$. In each E_λ, we can find vectors w_j such that they complete $V_\lambda = \{Sv_i \in \mathscr{B} \,|\, v_i\ \lambda\text{-eigenvector}\}$ to a basis of E_λ and that are orthogonal to all vectors in V_λ (consider the orthogonal complement of the subspace generated by V_λ in E_λ). Repeating this procedure on each E_λ, we find vectors $\{w_1,\ldots,w_{n-m}\}$ as needed. □

The statement and proof above holds for arbitrary matrices A by replacing 'eigenbasis' by 'maximal linearly independent set' and removing the condition $S^T w_i = 0$. It would be interesting to know whether the condition $S^T w_i = 0$ holds for motif eigenvectors in the directed case as well (the proof above is no longer valid).

We call $\{Sv_1,\ldots,Sv_m\}$ *quotient eigenvectors* of $\mathscr{G}$: they arise from a quotient eigenbasis by repeating values on each orbit; and we call $\{w_1,\ldots,w_{n-m}\}$ *redundant eigenvectors* of $\mathscr{G}$: they arise from the symmetries in the network (and they 'disappear' in the quotient), and add up to zero on each orbit ($S^T w_i = 0$). Similarly, we use the terminology *quotient* and *redundant* eigenvalues for their associated spectrum.

Further to the spectral decomposition theorem, we can give an even more precise description of the redundant spectrum: it is made of the contributions from the spectrum of each individual symmetric motif, as we explain next.

Redundant Spectrum of Symmetric Motifs

As stated in the Main Text (Theorem 3 in Methods), the redundant spectrum of a graph $\mathscr{M}$ is a subset of the spectrum of any (undirected) network $\mathscr{G}$ containing $\mathscr{M}$ as a symmetric motif. This essentially follows from the condition $S^T w = 0$ for redundant eigenvectors of $\mathscr{M}$.

Theorem 5. *Let $\mathscr{M}$ be a symmetric motif of a (possibly weighted) undirected graph $\mathscr{G}$. If (λ, w) is a redundant eigenpair of $\mathscr{M}$ then $(\lambda, \widetilde{w})$ is a eigenpair of $\mathscr{G}$, where $\widetilde{w}$ is equal to w on (the vertices of) $\mathscr{M}$, and zero elsewhere.*

Proof. Since (λ, v) is an $\mathscr{M}$-eigenpair,

$$\sum_{j \in V(\mathscr{M})} [A_{\mathscr{M}}]_{ij} w_j = \lambda w_i \quad \forall i \in V(\mathscr{M}),$$

where $A_{\mathscr{M}}$ is the adjacency matrix of $\mathscr{M}$. We can decompose $\mathscr{M}$ into orbits,

$$V(\mathscr{M}) = V_1 \cup \ldots \cup V_m,$$

and, by the spectral decomposition theorem above applied to $\mathscr{M}$, w is orthogonal to each orbit, that is,

$$\sum_{j \in V_i} w_j = 0 \quad \forall 1 \leq i \leq m.$$

We need to show that $(\lambda, \widetilde{w})$ is a $\mathscr{G}$-eigenpair. Let us write A for the adjacency matrix of $\mathscr{G}$ (recall $\mathscr{M}$ is a subgraph so A restricts to $A_{\mathscr{M}}$ on $\mathscr{M}$). We need to show $A\widetilde{w} = \lambda \widetilde{w}$. Given $i \in V(\mathscr{G})$, we have two cases. First, if $i \in V(\mathscr{M})$,

$$\begin{aligned}\sum_{j \in V(\mathscr{G})} [A]_{ij} \widetilde{w}_j &= \sum_{j \in V(\mathscr{M})} [A]_{ij} \widetilde{w}_j + \sum_{j \in V(\mathscr{G}) \setminus V(\mathscr{M})} [A]_{ij} \widetilde{w}_j \\ &= \sum_{j \in V(\mathscr{M})} [A]_{ij} w_j = \lambda w_i = \lambda \widetilde{w}_i,\end{aligned}$$

since $\widetilde{w}$ equals w on $\mathscr{M}$, and is zero outside $\mathscr{M}$. The second case, when $i \in V(\mathscr{G}) \setminus V(\mathscr{M})$, gives

$$\sum_{j \in V(\mathscr{G})} [A]_{ij} \widetilde{w}_j = \sum_{j \in V(\mathscr{M})} [A]_{ij} w_j ,$$

as before, and then we use the decomposition of $\mathscr{M}$ into orbits,

$$\sum_{j \in V(\mathscr{M})} [A]_{ij} w_j = \sum_{k=1}^{m} \sum_{j \in V_k} [A]_{ij} w_j = \sum_{k=1}^{m} \alpha_k \sum_{j \in V_k} w_j .$$

Here we have used that the vertex i, outside the motif, connects uniformly to each orbit (see Main Text), that is, $A_{ij_1} = A_{ij_2}$ for all $j_1, j_2 \in V_k$, and we call this quantity α_k. Finally, recall that w is orthogonal to each orbit, to conclude

$$\sum_{j \in V(\mathscr{M})} [A]_{ij} w_j = \sum_{k=1}^{m} \alpha_k \sum_{j \in V_k} w_j = 0 = \lambda \widetilde{w}_i . \qquad \square$$

Therefore, the redundant spectrum of $\mathscr{G}$ is the union of the redundant eigenvalues of the symmetric motifs, together with their redundant eigenvectors localised on them. Since most symmetric motifs in real-world networks are basic, most symmetric motifs in the network representation of a network measure will be basic too. Given their constrained structure, one can in fact determine the redundant spectrum of BSMs with up to few orbits, for arbitrary undirected networks with symmetry. This is what we do next.

Redundant Spectrum of a 1-orbit BSM

A BSM with one orbit is an (α, β)-uniform graph $K_n^{\alpha,\beta}$ with adjacency matrix $A_n^{\alpha,\beta} = (a_{ij})$ given by $a_{ij} = \alpha$ and $a_{ii} = \beta$ for all $i \neq j$. Then $K_n^{\alpha,\beta}$ has eigenvalues $(n-1)\alpha + \beta$ (non-redundant), with multiplicity 1, and $-\alpha + \beta$ (redundant), with multiplicity $n-1$. The corresponding eigenvectors are $\mathbf{1}$, the constant vector 1 (non-redundant), and $\mathbf{e}_i$, the vectors with non-zero entries 1 at position 1, and -1 at position i, $2 \leq i \leq n$ (redundant). This can be shown directly by computing $A_n^{\alpha,\beta}\mathbf{1}$ and $A_n^{\alpha,\beta}\mathbf{e}_i$, and noting that $\mathbf{1}$, $\mathbf{e}_2, \ldots, \mathbf{e}_n$ are linearly independent (although not orthogonal) and thus form an eigenbasis. Indeed, $A_n^{\alpha,\beta}\mathbf{1}$ is the vector of column sums of the matrix $A_n^{\alpha,\beta}$, which are constant $(n-1)\alpha + \beta$, and $A_n^{\alpha,\beta}\mathbf{e}_i$ is the constant 0 vector, except possibly at positions 1, which equals $\beta - \alpha$, and i, which equals $\alpha - \beta$.

Note that, for unweighted graphs without loops ($\beta = 0$, $\alpha \in \{0, 1\}$), we recover the redundant eigenvalues 0 and -1 predicted in[8].

Redundant Spectrum of a 2-orbit BSM

A BSM with two orbits is a uniform join of the form

$$K_n^{\alpha_1,\beta_1} \overset{\gamma,\delta}{\longleftrightarrow} K_n^{\alpha_2,\beta_2} . \tag{23}$$

Define $a = \alpha_1 - \beta_1$, $b = \alpha_2 - \beta_2$, $c = \gamma - \delta$, and note that $c \neq 0$: otherwise $\gamma = \delta$ and we can freely permute one orbit while fixing the other, that is, this would not be a BSM with two orbits but rather two BSMs with one orbit each. As above, let $\mathbf{e}_i$ be the vector with non-zero entries 1 at position 1, and -1 at position i, for any $2 \leq i \leq n$.

Lemma 1. *The following set of vectors is linearly independent*

$$\{(\kappa_1 e_i \,|\, e_i), (\kappa_2 e_i \,|\, e_i) \mid 2 \leq i \leq n\}$$

for all $\kappa_1 \neq \kappa_2 \in \mathbb{R}$.

Proof. Define the $(n-1) \times n$ matrix

$$B_n = \left(\mathbf{1} \mid -\mathrm{Id}_{n-1}\right)$$

where $\mathbf{1}$ is a constant 1 column vector, and Id_{n-1} the identity matrix of size $n-1$. The set of vectors in the statement can be arranged in block matrix form as

$$\left(\begin{array}{c|c} \kappa_1 B_n & B_n \\ \hline \kappa_2 B_n & B_n \end{array}\right) .$$

This matrix has a minor of order $2(n-1)$,

$$\det\left(\begin{array}{c|c} -\kappa_1 \operatorname{Id}_{n-1} & -\operatorname{Id}_{n-1} \\ \hline -\kappa_2 \operatorname{Id}_{n-1} & -\operatorname{Id}_{n-1} \end{array}\right).$$

Using that $\det\left(\begin{smallmatrix} A & B \\ C & D \end{smallmatrix}\right) = AD - BC$ whenever A, B, C, D are square blocks of the same size and C commutes with D[22], this minor equals

$$\det\left(-\kappa_1 \operatorname{Id}_{n-1} - \kappa_2 \operatorname{Id}_{n-1}\right) = (-1)^{n-1}(\kappa_1 + \kappa_2)^{n-1} \neq 0 \iff \kappa_1 \neq \kappa_2. \qquad \square$$

Next we derive conditions for a vector $v_i = (\kappa e_i | e_i)$ to be an eigenvector of the uniform join (23), that is, $Av_i = \lambda v_i$, for some $\lambda \in \mathbb{R}$, where A is the (symmetric) adjacency matrix of the uniform join,

$$A = \left(\begin{array}{c|c} A_n^{\alpha_1,\beta_1} & A_n^{\gamma,\delta} \\ \hline A_n^{\gamma,\delta} & A_n^{\alpha_2,\beta_2} \end{array}\right). \tag{24}$$

The jth entry of the vector Av is

$$\begin{array}{ll} \kappa\beta_1 - \kappa\alpha_1 + \delta - \gamma = -(\kappa a + c) & j = 1 \\ \kappa\alpha_1 - \kappa\beta_1 + \gamma - \delta = \kappa a + c & j = i \\ \kappa\delta - \kappa\gamma + \beta_2 - \alpha_2 = -(\kappa c + b) & j = n+1 \\ \kappa\gamma - \kappa\delta + \alpha_2 - \beta_2 = \kappa c + b & j = n+i \\ 0 & \text{otherwise.} \end{array}$$

Comparing these with the entries of the vector λv_i, we obtain

$$Av_i = \lambda v_i \iff (\kappa a + c = -\lambda\kappa \ \text{ and } \ \kappa c + b = -\lambda), \tag{25}$$

The two equations on the right-hand side are satisfied if and only if $\lambda = -\kappa c - b$ and κ is a solution of the quadratic equation

$$c\kappa^2 + (b-a)\kappa - c = 0, \tag{26}$$

which has two distinct real solutions

$$\frac{(a-b) \pm \sqrt{(a-b)^2 + 4c^2}}{2c}, \tag{27}$$

since $c \neq 0$, as explained above. Together with the lemma, we have shown the following (Theorem 4 in Methods).

Theorem 6. *The redundant spectrum of a symmetric motif with two orbits $K_n^{\alpha_1,\beta_1} \overset{\gamma,\delta}{\longleftrightarrow} K_n^{\alpha_2,\beta_2}$ is given by the eigenvalues*

$$\lambda_1 = -b - c\kappa_1 = \frac{-(a+b) + \sqrt{(a-b)^2 + 4c^2}}{2}, \text{ and,}$$
$$\lambda_2 = -b - c\kappa_2 = \frac{-(a+b) - \sqrt{(a-b)^2 + 4c^2}}{2},$$

each with multiplicity $n-1$, and eigenvectors $(\kappa_1 e_i | e_i)$ and $(\kappa_2 e_i | e_i)$ respectively, where κ_1 and κ_2 are the two solutions of the quadratic equation $c\kappa^2 + (b-a)\kappa - c = 0$, $a = \alpha_1 - \beta_1$, $b = \alpha_2 - \beta_2$ and $c = \gamma - \delta \neq 0$.

It is interesting to note that the redundant eigenvalues of uniform graphs and joins depend on the differences ($\alpha_1 - \beta_1$, $\alpha_2 - \beta_2$, $\gamma - \delta$) rather than the particular coefficients. Therefore, for redundant spectrum calculations, we can assume all BSMs to be loop-less ($\beta = 0$), and uniform joins of type $(0, \delta)$.

For unweighted graphs without loops, we recover the redundant eigenvalues for BSMs with two orbits predicted in[8], as follows. We have $\beta_1 = \beta_2 = 0$, $\alpha_1, \alpha_2, \gamma, \delta \in \{0, 1\}$ and thus $a, b \in \{0, 1\}$ and $c \in \{-1, 1\}$. If $a = b$, the quadratic equation becomes $\kappa^2 - 1 = 0$ with solutions $\kappa = \pm 1$ and thus $\lambda = -b - c\kappa \in \{-2, -1, 0, 1\}$. If $a \neq b$ we can assume $a = 1$, $b = 0$ and the quadratic $c\kappa - \kappa - c = 0$ has solutions φ and $1 - \varphi$ if $c = 1$, $-\varphi$ and $\varphi - 1$ if $c = -1$, where $\varphi = \frac{1+\sqrt{5}}{2}$ is the golden ratio. In either case, the redundant eigenvalues $\lambda = -b - c\kappa = -c\kappa$ are $-\varphi$ and $\varphi - 1$. Altogether, the redundant eigenvalues for 2-orbit BSMs are $\{-2, -\varphi, -1, 0, \varphi - 1, 1\}$, which equals the redundant eigenvalues RSpec_2 in the notation of[8].

We omit the calculation of the redundant spectrum of BSMs with three (or more) orbits, as it becomes much more elaborate, and its relevance in real-world networks is less justified (SI Table 1).

For real-world networks, we predict symmetry to explain most of the discrete part of the spectrum (observed as 'peaks' in the spectral density) of the network representation of any network measure F (this can be quantified by comparing high-multiplicity eigenvalues in the parent versus the quotient network). This is shown for the Laplacian eigenvalues on six of our test networks in Figure 5 in the Main Text, with 89% to 97% of the discrete spectrum explained by the underlying network symmetry.

Let us define RSpec_1^F and RSpec_2^F as the sets of redundant eigenvalues of $F(\mathscr{G})$ associated to BSMs of one, respectively two, orbits, for any network $\mathscr{G}$, given by formulae above (or Table 2 in the Main Text). Our results predict most of the discrete part of the spectrum of $F(\mathscr{G})$ to occur at the values of these sets (cf. Figure 5 in the Main Text). For specific choices of F (communicability, Laplacian, shortest path), we will be able to describe these sets in more detail, and for the adjacency matrix, as explained above, we recover the sets RSpec_1 and RSpec_2 in[8].

Eigendecomposition Algorithm

We can use the spectral decomposition theorem above to compute the spectrum and eigenbasis of a weighted undirected network with symmetries (equivalently, a diagonalisation of the symmetric adjacency matrix $A = UDU^T$), such as $F(\mathscr{G})$, from those of the quotient, and the redundant ones of the symmetric motifs. The algorithm (Algorithm 5 below, or 4 in Methods) computes the spectral decomposition (eigendecomposition) of the quotient matrix, then, for each motif, the redundant eigenpairs.

```
Input: adjacency matrix A, characteristic matrix S, list of motifs
Output: spectral decomposition A = UDU^T

initialise U, D to zero matrices
Λ ← diag(sum(S))
B_sym ← Λ^{-1/2} S^T A S Λ^{-1/2}
[U_q, D_q] ← eig(B_sym) so that B_sym = U_q D_q U_q^{-1}
U_q ← Λ U_q
U ← ( S U_q | 0 )
D ← ( D_q | 0 )
    (  0  | 0 )
foreach motif do
    A_sm ← A(motif, motif)
    compute orbits from motif and S
    S_sm ← S(motif, orbits)
    [U_sm, D_sm] ← eig(A_sm)
    for λ ∈ unique(diag(D_sm)) do
        U_λ ← λ-eigenvectors from U_sm
        Z ← null(S_sm^T U_λ)
        d ← ncol(Z)
        if d > 0 then
            store U_λ Z in U
            store λ in D with multiplicity d
        end
    end
end
```

Algorithm 5: Eigendecomposition algorithm.

In more detail, we first compute the spectral decomposition `eig` of the symmetric quotient $B_{\mathrm{sym}} = \Lambda^{-1/2} S^T A S \Lambda^{-1/2}$ where Λ is the diagonal matrix of the orbit sizes (which can be obtained as the column sums of S). This matrix is symmetric and has the same eigenvalues as the left quotient. Moreover, if $B_{\mathrm{sym}} = U_q D_q U_q^{-1}$ then the left quotient eigenvectors are the columns of ΛU_q. These become, in turn, eigenvectors of A by repeating their values on each orbit, we can be obtained mathematically by left multiplying by the characteristic matrix S. Then, for each motif, we compute the redudant eigenpairs using a null space matrix (see below), storing eigenvalues and localised (zero outside the motif) eigenvectors.

Only redundant eigenvectors of a symmetric motif (that is, those which add up to zero on each orbit) become eigenvectors of A by extending them as zero outside the symmetric motif (Theorem 5). Therefore, we need to construct redundant eigenvectors from the ouput of `eig` on each motif (the spectral decomposition of the corresponding submatrix). If $U_\lambda = \left(v_1 \mid \ldots \mid v_k\right)$ are λ-eigenvectors of a symmetric motif with characteristic matrix of the orbit partition S_{sm}, we need to find linear combinations such that

$$S_{\text{sm}}^T\left(\alpha_1 v_1 + \ldots + \alpha_k v_k\right) = \mathbf{0} \iff S_{\text{sm}}^T U_\lambda \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_k. \end{pmatrix}. \tag{28}$$

Therefore, if the matrix $Z \neq 0$ represents the null space of $S_{\text{sm}}^T U_\lambda$, that is, $S_{\text{sm}}^T U_\lambda Z = 0$ and $Z^T Z = 0$, then the columns of $U_\lambda Z$ are precisely the redundant eigenvectors. This is implemented in Algorithm 5 within the innermost **for** loop.

SI Vertex Measures

The network representation we have exploited so far does not apply directly to a vertex measure G, unless G is defined via a pairwise network measure (this is often the case). We still have, however, symmetry-induced compression, and computational reducibility, as mentioned in the Main Text and detailed below.

As a vertex measure G is constant on orbits, we only need to store one value per orbit. Let S be the characteristic matrix of the partition of the vertex set into orbits, and Λ the diagonal matrix of orbit sizes (column sums of S). If G is represented by a vector $v = (G(i))$ of length $n_{\mathscr{G}}$, we can compress it by storing only one value per orbit, formally $w = \Lambda^{-1} S^T v$, a vector of length $n_{\mathscr{Q}}$. We recover $v = S^T w$, as the next result guarantees (Theorem 5 in Methods).

Theorem 7. *If v is a vector of length $n_{\mathscr{G}}$ constant on orbits, then $S\Lambda^{-1}S^T v = v$.*

Proof. First, note that $S^T S = \Lambda$ (this holds for any partition of the vertex set),

$$[S^T S]_{\alpha\beta} = \sum_i [S^T]_{\alpha i}[S]_{i\beta} = \sum_i [S]_{i\alpha}[S]_{i\beta} = \begin{cases} 0 & \text{if } \alpha \neq \beta, \\ |V_\alpha| & \text{if } \alpha = \beta. \end{cases}$$

As v is constant on orbits, it is already of the form $v = Sw$ for some w. Therefore

$$S\Lambda^{-1}S^T v = S\Lambda^{-1}S^T S w = Sw = v. \qquad \square$$

In terms of computational reduction, we only need to evaluate G once per orbit, achieving an $\tilde{n}_{\mathscr{Q}}$ time reduction, assuming G is evaluated at vertices independently, which is often not the case. More interesting is the case when G can be recovered from its value at the quotient network. We call G *quotient recoverable* if $G(\mathscr{G})$ can be obtained from $G(\mathscr{Q})$, where $\mathscr{Q}$ is a (possibly annotated, or basic) quotient of $\mathscr{G}$, for all networks $\mathscr{G}$. Not every vertex measure is quotient recoverable, but when one is, it can lead lead to a significant computational time reduction (Figure 4 in the Main Text).

SI Applications

Network Symmetry Computation

The results in the Main Text depend on an effective computation, storage and manipulation of the symmetries on an (unweighted, undirected, possibly very large) network $\mathscr{G}$. Here we present our approach, based in the geometric decomposition of the automorphism group of the graph. Full implementations of the algorithms outlined below are available at[21]. (For weighted or directed network, see concluding remarks in this section.)

We obtain the geometric decomposition in three steps. First, we use a graph automorphism algorithm to compute generators. Secondly, we partition the generators into disjoint-support classes (each class corresponds to a symmetric motif). Thirdly, we compute symmetric motif orbits and types from the generators and their disjoint-support partition. Below we explain each step in more detail.

We use `saucy3`[23] to compute a list of generators of the automorphism group from an edge list. Other open-source software tools are available, such as `nauty`[16,24], `traces`[24,25] or `bliss`[26]. Although `saucy` does not compute a canonical labelling (relevant to the graph isomorphism problem but not to the geometric decomposition), it is extremely fast for large but sparse networks[27] such as the ones representing real-world systems. In practice, we found a list of generators in less than two seconds for all our test networks, except our largest example LiveJournal, in just over eight seconds (Table 1 in the Main Text). Due to the similarities with `nauty`, it would be interesting to know whether the set of generators produced by `saucy` is also essential, which would guarantee that the geometric decomposition below is optimal; this seems to be the case in practice for all our test networks.

The next step is to partition the set of generators X into support-disjoint classes $X = X_1 \cup \ldots \cup X_k$. For that, we use a bipartite graph representation of vertices V and generators X. Let $\mathscr{B}$ be the graph with vertex set $V \cup X$ and undirected edges $\{i, \sigma\}$ whenever $i \in \mathrm{supp}(\sigma)$. Clearly, the finest partition into support-disjoint classes of generators correspond to the connected components of $\mathscr{B}$ (as vertex sets intersected with X). Using for instance Tarjan's algorithm, we have an efficient procedure to find the support-disjoint decomposition in linear time. In practice, this step took less than five seconds for each of our test networks. Each class in the support-disjoint partition above corresponds to (the vertex set of) a symmetric motif, by Eqs. 2 and 3.

The last step consists of computing the orbits and type of each symmetric motif. This was done in GAP[28], but any computer algebra system that can manipulate permutation groups can be used. We applied Algorithm 6 (Algorithm 1 in Methods) to each X_i in the support-disjoint decomposition above (this can be done in parallel). The pseudocode in Algorithm 6 assumes a computer algebra system that can compute the group generated by a set of permutations, its orbits, and whether the induced action on an orbit is a natural symmetric group (that is, a group acting as the symmetric group on its moved points). If that is the case, and all orbits are of the same size m, it is a basic symmetric motif of type m, that is, the corresponding geometric factor is S_m. If this is not the case, it is a complex symmetric motif and we set $m = 0$. Algorithm 6 outputs a list of orbits, and the integer m. Note that the second part of the algorithm (the outermost **if-then-else** loop) is only necessary if we need the orbit types, for example in order to later compute the basic quotient. In terms of computational time, our GAP non-parallel implementation computes about 2,000 generators per second in our small and medium networks, which suggests at most 205 seconds for our largest network (and divided by the number of processors available in a parallel implementation). Unfortunately, the generator-per-second rate decreases with the size of the network (up to 10 generators per second for the largest network) due to the internal representation of permutation groups in GAP. Fixing this issue is beyond the scope of the present article, but perhaps other choice of software should be considered for computations involving large integers.

Input: X a set of permutations of a symmetric motif
Output: $O_1, \ldots, O_k$ orbits, and type m, of the symmetric motif

```
H ← Group(X)
{O_1,...,O_k} ← Orbits(H)
m ← min(size(O_1),...,size(O_k))
if m == max(size(O_1),...,size(O_k)) then
    for i ← 1 to k do
        if not IsNaturalSymmetricGroup(Action(H,O_i)) then
            m ← 0
            break
        end
    end
else
    m ← 0
end
```

Algorithm 6: Orbits and type of a symmetric motif.

In terms of data structures, we represent vertices by integers 0 to $n_{\mathscr{G}} - 1$, and an undirected, unweighted graph by an edgelist of vertex pairs, which saucy transforms into a list of generators, each written as a list of vertex transpositions. The support-disjoint partition genpartition is simply a integer array such that genpartition$(i) = j$ if the ith generator belongs to X_j. Finally, each symmetric motif is store in motifs as a list of orbits, and their type in an integer vector motiftype. Alternatively, we could store each orbit (a list of integers) separately in a list orbits and the assignment of orbits to motifs in an integer array orbpartition. Also note that the partition of the vertex set into orbits can be also represented by its characteristic matrix S. This is an $n \times m$ matrix with at most n non-zero entries, hence it can be stored and manipulated very efficiently in sparse form.

We end this section with a few remarks. The network symmetry computation is a pre-processing step that needs to be calculated only once for each network, and can be stored efficiently as explained above (e.g. 16.3MB for our largest test network, compared to 700MB for the edge list). As mentioned before, most of the results in this paper can be applied to networks with edge weights or directions, or other edge or node labels, by restricting to the symmetries preserving the additional structure. In that situation, one can incorporate the restrictions to the automorphism group calculation (saucy admits vertex colouring, and nauty both vertex colouring and directed edges), or compute the geometric decomposition of the underlying unweighted, undirected graph and then incorporate the restrictions one symmetric motif at a time (cf. Theorem 1).

Communicability

Communicability is a measure of network connectivity between pairs of vertices which takes into account all possible walks from one vertex to the other by using the powers of the adjacency matrix A. Namely, if we choose coefficients a_k such as the matrix series

$$\sum_{k=0}^{\infty} a_k A^k \tag{29}$$

converges, then the (i, j)-entry of this limit matrix is a weighted sum of all the walks from i to j, and thus can be used as a network connectivity measure. We normally expect coefficients a_k such that we obtain positive values for the communicability, and which give less weight to longer walks. An standard choice is the factorial coefficients $a_k = \frac{1}{k!}$, which guarantees convergence for any matrix A and, in fact,

$$e^A = \sum_{k=0}^{\infty} \frac{A^k}{k!}, \tag{30}$$

the exponential matrix of A. The diagonal entry $[e^A]_{ii}$ is called the *subgraph centrality* of vertex i[29], and its sum over all the vertices the *Estrada index* of the network[30]. In general, one can define communicability for an arbitrary real analytic function (such as $f(x) = e^x$) within its radius of convergence R around 0,

$$f(x) = \sum_{k=0}^{\infty} a_k x^k \qquad |x| < R. \tag{31}$$

Given such a function f, we define the *f-communicability matrix* of a network with adjacency matrix A as

$$f(A) = \sum_{k=0}^{\infty} a_k A^k \qquad \text{if } \|A\| < R, \tag{32}$$

where $\|\cdot\|$ is a given matrix norm, and the power series convergence is with respect to that norm. (For a detailed treatment of matrix norms and convergence see[22].) From now on we will implicitly assume all calculations to be within the convergence radius of f, possibly by normalising the matrix A. The *f-communicability* from vertex i to vertex j is thus the (i, j)-entry of the communicability matrix, $[f(A)]_{ij}$. For consistent terminology, we will call the f-communicability of a vertex with itself its *f-centrality*, inherently a network centrality measure. (We consider other centrality measures, including subgraph centrality, and the effect of symmetry on them, below.) Note that matrix functions $f(A)$ can be defined in more generality[31], however the power series definition is the one with an obvious graph theoretic interpretation.

Structural Properties

We represent communicability as a network on the same set of nodes. Namely, we define the *f-communicability graph* of $\mathscr{G}$, written $f(\mathscr{G})$, as the graph with adjacency matrix $f(A)$, which we call the *communicability matrix*. This is a weighted, complete (and possibly directed, if A is not symmetric) graph with loops. For every $i \neq j$, there is an edge from vertex i to a vertex j weighted by the communicability from i to j, and a self-loop at every vertex weighted by its f-centrality. The f-communicability network, although dense, inherits all the symmetries of $\mathscr{G}$ and hence $f(\mathscr{G})$ has the same geometric decomposition, symmetric motifs, and orbits. For real-world networks, most symmetric motifs will be basic and, as induced graphs, the basic symmetric motifs are uniform joins of orbits, and each orbit is a uniform graph hence characterised by two parameters, the subgraph centrality of each vertex, and the communicability between different vertices, within the orbit. In particular, this explains what we observed in our toy example, Figure 1 in the Main Text. In terms of post-processing compression, as a full network measures, average symmetry compression with ratio c_{full} and lossless symmetry compression with ratio $c_{\text{full}}^{\text{basic}}$ apply, accounting for the symmetry-induced redundancy present on $f(\mathscr{G})$, or $f(A)$.

Quotient Recoverability

Communicability satisfies average quotient recoverability, as it 'commutes' with the quotient, that is, the communicability of the quotient is the quotient of the communicability, in symbols,

$$f(Q(A)) = Q(f(A)). \tag{33}$$

This is (17) for communicability, and implies exact recovery for external edges, and average recovery for internal edges. To prove (33), let $B = Q(A)$ be the (left) quotient matrix with respect to the partition into orbits. Since the partition is equitable, we have $AS = SB$, and hence $A^n S = SB^n$. Now, using the matrix definition of the quotient, we have

$$Q(f(A)) = \Lambda^{-1} S^T \left(\sum_{n=0}^{\infty} a_n A^n \right) S = \sum_{n=0}^{\infty} a_n \left(\Lambda^{-1} S^T A^n S \right) = \sum_{n=0}^{\infty} a_n \left(\Lambda^{-1} S^T S B^n \right) = \sum_{n=0}^{\infty} a_n B^n = f(B), \tag{34}$$

since $\Lambda^{-1}S^T S$ is the identity matrix.

Exact recovery for vertices within the same symmetric motif cannot be done in general, as the internal connectivity, replaced by the average connectivity, is lost in the quotient. However, as mentioned in the Main Text, we can use the spectral decomposition algorithm to reduce the computation of the communicability, as $A = UDU^T$ clearly implies $f(A) = Uf(D)U^T$. The computational time reduction of this approach is displayed in Figure 4 in the Main Text, for the exponential function.

Spectral Properties

The f-communicability network has eigenvalues $f(\lambda)$, where λ is an eigenvalue of the original network, and same eigenvectors: $Av = \lambda v$ implies $f(A)v = f(\lambda)v$ (equivalently, $A = UDU^T$ implies $f(A) = Uf(D)U^T$). In particular, its redundant spectrum consists of the function f applied to the redundant spectrum of $\mathscr{G}$, together with the redundant eigenvectors localised on the symmetric motifs. In particular, for undirected, unweighted networks,

$$\mathrm{RSpec}_1^{f\text{-comm}} = \{f(\lambda) \mid \lambda \in \mathrm{RSpec}_1\} = \{f(0), f(-1)\}, \text{ and} \tag{35}$$
$$\mathrm{RSpec}_2^{f\text{-comm}} = \{f(\lambda) \mid \lambda \in \mathrm{RSpec}_2\} = \{f(-2), f(-\varphi), f(-1), f(0), f(\varphi-1), f(1)\}, \tag{36}$$

account for most of the discrete part of the spectrum of the matrix $f(A)$, for the adjacency matrix A of a typical real-world network.

Shortest Path Distance

Let $A = (a_{ij})$ be the adjacency matrix of an unweighted, but possibly directed, network $\mathscr{G}$. A *path* of *length* n is a sequence $(v_1, v_2, \dots, v_{n+1})$ of distinct vertices, except possibly $v_1 = v_{n+1}$, such that v_i is connected to v_{i+1} for all $1 \le i \le n-1$. The *shortest path distance* $d^{\mathscr{G}}(u,v)$ is the length of the shortest (minimal length) path from u to v. (Technically, it is only a distance, or metric, if $\mathscr{G}$ is undirected.) If $p = (v_1, v_2, \dots, v_n)$ is a path and $\sigma \in \mathrm{Aut}(\mathscr{G})$, we define $\sigma(p) = (\sigma(v_1), \sigma(v_2), \dots, \sigma(v_n))$, also a path since σ is a bijection. A *subpath* of p is a path of the form $(v_k, v_{k+1}, \dots, v_l)$ for some $1 \le k \le l \le n$. A path p is a *shortest path* if it is of minimal length between its endpoints.

The following result contains the claims in the Main Text (Theorem 6 in Results).

Theorem 8. *Let $A = (a_{ij})$ be as above. Then*

(i) if $(v_1, v_2, \dots, v_n)$ is a shortest path from v_1 to v_n and $\sigma \in Aut(\mathscr{G})$, then $(\sigma(v_1), \sigma(v_2), \dots, \sigma(v_n))$ is a shortest path from $\sigma(v_1)$ to $\sigma(v_n)$;

(ii) if $(v_1, v_2, \dots, v_n)$ is a shortest path from v_1 to v_n, and v_1 and v_n belong to different symmetric motifs, then v_i and v_{i+1} belong to different orbits, for all $1 \le i \le n-1$;

(iii) if u and v belong to orbits U, respectively V, in different symmetric motifs, then the distance from u to v in $\mathscr{G}$ equals the distance from U to V in the unweighted (or skeleton) quotient $\mathscr{Q}$.

These statements mean that (i) automorphisms preserve shortest paths and their lengths; (ii) shortest paths do not contain intra-orbit edges; and (iii) shortest path distance is a partially quotient recoverable structural measure.

Proof of Theorem. (i) Since automorphisms are bijections and preserve adjacency, $(\sigma(v_1), \sigma(v_2), \dots, \sigma(v_n))$ is a path from $\sigma(u)$ to $\sigma(v)$ of the same length. If there were a shorter path $(\sigma(u) = w_1, w_2, \dots, \sigma(v) = w_m)$, $m < n$, the same argument applied to σ^{-1} gives a shorter path $(u = \sigma^{-1}(w_1), \sigma^{-1}(w_2), \dots, v = \sigma^{-1}(w_m))$ from u to v, a contradiction.

(ii) Any subpath of a minimal length path is also of minimal length between its endpoints. Arguing by contradiction, there exists a subpath $p = (w_1, w_2, \dots, w_n)$ (or $p = (w_n, w_{n-1}, \dots, w_1)$), such that w_1 and w_2 belong to the same orbit, and w_n belongs to a different symmetric motif. Hence, we can find $\sigma \in \mathrm{Aut}(\mathscr{G})$ with $\sigma(w_2) = w_1$ and fixing w_n. This implies $\sigma(p) = (\sigma(w_1), \sigma(w_2) = w_1, \sigma(w_3), \dots, \sigma(w_n) = w_n)$, a shortest path by (i), of length $n-1$. The subpath $(w_1, \sigma(w_3), \dots, w_n)$ has length $n-2$, contradicting p being a minimal length path from w_1 to w_n. (The case $p = (w_n, w_{n-1}, \dots, w_1)$ is analogous.)

(iii) Let $p = (u = v_1, v_2, \dots, v_{n+1} = v)$ be a shortest path from u to v, so that $d^{\mathscr{G}}(u,v) = n$. Let V_k be the orbit containing v_k, for all k. By (ii), $V_k \neq V_{k+1}$ for all $1 \le k \le n$ thus $q = (U = V_1, V_2, \dots, V_{n+1} = V)$ is a path in $\mathscr{Q}$ and $d^{\mathscr{Q}}(U,V) \le n$. By contradiction, assume there is a shorter path in $\mathscr{Q}$ from U to V, that is, $(U = W_1, W_2, \dots, W_{m+1} = V)$ with $m < n$. The we can construct a path in $\mathscr{G}$ from u to v of length m (a contradiction), as follows. For each $1 \le i \le m$, W_i is connected to W_{i+1} in $\mathscr{Q}$, hence there is a vertex in W_i connected to at least one vertex in W_{i+1}. Since vertices in an orbit are structurally equivalent, *any* vertex in W_i is then connected to at least one vertex in W_{i+1} (formally, if $w \in W_i$ is connected to $w' \in W_{i+1}$ then $\sigma(w) \in W_i$ is connected to $\sigma(w') \in W_{i+1}$). This allows us to construct a path in $\mathscr{G}$ from u to v of length $m < n$, a contradiction. $\square$

Distances between points within the same motif cannot in general be directly recovered from the quotient, not even for BSMs. (Consider for instance the double star, motif M_1, in Figure 1 in the Main Text. The distance from the top red to the bottom blue vertex is three, while in the quotient is one.) In general, therefore, the shortest path distance is partially, but not average, quotient recoverable. Intra-motif distances, if needed, could still be recovered one motif at a time. (If we are only interested in large distances, computing them in the quotient suffices as one can show that $d(i,j) \leq 2k$ for i,j vertices in a symmetric motif with k orbits.)

Note that these results can be exploited for other graph-theoretic notions defined in terms of distance, for example eccentricity (and thus radius or diameter), which only depends on maximal distances and thus it can be computed directly in the quotient (see section Eccentricity below).

In terms of post-calculation compression, the quotient compression ratio c_{full} applies, accounting for the amount of structural redundancy due solely to symmetries. The spectral results, although perhaps less relevant, still apply for $d(\mathscr{G})$, the graph encoding pairwise shortest path distances. The adjacency matrix $d(A) = (d^{\mathscr{G}}(i,j))$ is nonzero outside the diagonal, hence $d(\mathscr{G})$ is a all-to-all weighted network without self-loops and integer weights, and so is each symmetric motif. Using the formula in Theorem 6, we can compute values of the most significant part of the discrete spectrum (redundant eigenvalues) of $d(A)$,

$$\text{RSpec}_1^d = \{-2,-1\}, \text{ and } \text{RSpec}_2^d = \left\{-3,-2,-1,0,-2\pm\sqrt{2},-3\pm\sqrt{2},\frac{-3\pm\sqrt{5}}{2},\frac{-5\pm\sqrt{5}}{2},\frac{-5\pm\sqrt{13}}{2}\right\}. \tag{37}$$

The shortest path distance can be generalised to positively weighted matrices, and then (i) and (ii) above still hold, with similar proofs, but not (iii), as weights along a path are not preserved in the quotient. In would be interesting to incorporate network symmetry to current shortest path algorithms, although this is outside the scope of the present article.

Laplacian Matrix

The Laplacian matrix $L = D - A$ can be seen as the adjacency matrix of a Laplacian network $\mathscr{L}$, which inherits the symmetries of $\mathscr{G}$ as explained in the Main Text. The symmetric motifs are almost identical in $\mathscr{L}$, except that all edges are weighted by -1, and all vertices have self-loops weighted by their degrees in $\mathscr{G}$. In particular, the motif structure depends on the how it is embedded in the network. For example, an orbit in a BSM in $\mathscr{G}$, originally a complete, respectively empty, graph in $\mathscr{G}$, becomes a uniform graph $K_n^{\varepsilon,d}$ where $\varepsilon = -1$, respectively $\varepsilon = 0$, and d is the degree in $\mathscr{G}$ of a (and hence any) vertex in the orbit. In particular, the symmetric motifs in $\mathscr{L}$ are not quite the Laplacian of the original motifs, as per the next result. Define the *external degree* of a vertex as the number of adjacent vertices outside the motif it belongs to. The next result is Theorem 7 in Results.

Theorem 9. *Let M be the vertex set of a symmetric motif $\mathscr{M}$ in a graph $\mathscr{G}$. Then M induces a symmetric motif in the Laplacian network $\mathscr{L}$ with adjacency matrix*

$$L_{\mathscr{M}} + \left(d_1 I_{m_1} \oplus \ldots \oplus d_k I_{m_k}\right),$$

where $L_{\mathscr{M}}$ is the ordinary Laplacian matrix of $\mathscr{M}$ considered as a graph on its own, and $d_1,\ldots,d_k$ are the external degrees of the k orbits of $\mathscr{M}$ of sizes $m_1,\ldots,m_k$. (Here I_n is the identity matrix of size n and we use $\oplus$ to construct a block diagonal matrix.)

(The proof of this theorem should be clear from the comments above.) Note that, for a motif with one orbit, this is the Laplacian of the motif translated by a multiple of the identity. In particular, the redundant eigenvalues of a BSM with one orbit are the redundant (high multiplicity) Laplacian eigenvalues of an empty or complete graph of size n plus the external degree d, that is, d, respectively $d+n$. In particular,

$$\text{RSpec}_1^L = \mathbb{Z}^+, \tag{38}$$

and we expect 'peaks' in the Laplacian spectral density at (small) positive integers, which indeed agrees what we observed in our test networks (Figure 5 in the Main Text). Using the formula in Theorem 6, we could compute the redundant spectrum for 2-orbit BSMs, and for other versions of the Laplacian (e.g. normalised, vertex weighted), but we believe this is out of the scope of the present article.

Observe that spectral decomposition applies, since L inherits all the symmetries of A, so Algorithm 5 provides an efficient way of computing the Laplacian eigendecomposition with an expected $sp = \tilde{n}_{\mathscr{Q}}^3$ (see Table 1 in the Main Text) computational time reduction.

Degree Centrality

The degree of a node (in- or out-degree if the network is directed) is a natural measure of vertex centrality. As expected, the degree is preserved by any automorphism σ, which can also be checked directly,

$$d_i = \sum_{j\in V} a_{ij} = \sum_{j\in V} a_{\sigma(i)\sigma(j)} = \sum_{j\in V} a_{\sigma(i)j} = d_{\sigma(i)}, \tag{39}$$

as automorphisms permute orbits (so $j \in V$ and $\sigma(j) \in V$ are the same elements but in a different order). In particular, the degree is constants on orbits.

We can recover the degree centrality from the quotient, as the out-degree of the left quotient (or the in-degree of the right quotient), as follows. (This is for illustration purposes rather than a worthwhile computational gain in using the quotient for degree calculations.) Let $B = (b_{\alpha\beta})$ be the adjacency matrix of the left quotient, and $V = V_1 \cup \ldots \cup V_m$ the partition of the vertex set into orbits. If $i \in V_\alpha$, then

$$d_i^{\mathscr{G}} = \sum_{j\in V} a_{ij} = \sum_{j\in V_1} a_{ij} + \ldots + \sum_{j\in V_m} a_{ij} = \frac{1}{n_1}\sum_{\substack{j\in V_1\\ i\in V_i}} a_{ij} + \ldots + \frac{1}{n_m}\sum_{\substack{j\in V_m\\ i\in V_i}} a_{ij} = b_{i1} + \ldots + b_{im} = d_\alpha^{\mathscr{Q},\mathrm{out}}. \tag{40}$$

Eccentricity

Although not discussed in the Main Text, the reciprocal of the eccentricity is a natural centrality measure. The eccentricity of a vertex is the maximal (shortest path) distance to any vertex in the network. As shortest path distance is partially quotient recoverable, and eccentricity depends on large distances only, so in practice we can ignore intra-motif distances and therefore recover eccentricity directly from the (unweighted, or skeleton) quotient as

$$\mathrm{ecc}^{\mathscr{G}}(i) = \mathrm{ecc}^{\mathscr{Q}}(\alpha) \quad \text{if } i \in V_\alpha. \tag{41}$$

In particular, the diameter (maximal eccentricity) and radius (minimal eccentricity) of a network coincides with that of the (unweighted, or skeleton) quotient.

Closeness Centrality

The closeness centrality of a node i in a graph $\mathscr{G}$, $cc^{\mathscr{G}}(i)$, is the average shortest path length to every node in the graph. As symmetries preserve distance, they also preserve closeness centrality, explicitly,

$$cc(i) = \frac{1}{n_{\mathscr{G}}}\sum_{j\in V} d(i,j) = \frac{1}{n_{\mathscr{G}}}\sum_{j\in V} d(\sigma(i),\sigma(j)) = \frac{1}{n_{\mathscr{G}}}\sum_{j\in V} d(\sigma(i),j) = cc(\sigma(i)), \tag{42}$$

and centrality is constant on each orbit, as expected. Moreover, closeness centrality can be recovered from the quotient (shortest paths does not contain intra-orbit edges, except between vertices in the same symmetric motif, see Theorem 8), as

$$cc^{\mathscr{G}}(i) = \sum_{l\neq k} \frac{n_l}{n_{\mathscr{G}}} d^{\mathscr{Q}}(V_k, V_l) + \frac{n_i}{n_{\mathscr{G}}} d_k \tag{43}$$

if i belongs to the orbit V_k and d_k is the average intra-motif distance, that is, the average distances of a vertex in V_k to any vertex in $\mathscr{M}$, the motif containing V_k. By annotating each orbit by d_k, we can recover betweenness centrality exactly. Alternatively, as $d_k \ll n$ (note that $d_k \leq m$ if $\mathscr{M}$ has m orbits), we can approximate $cc^{\mathscr{G}}(i)$ by the first summand, or simply by the quotient centrality $cc^{\mathscr{Q}}(\alpha)$, in most practical situations.

Eigenvector Centrality

Since the Perron-Frobenius eigenvalue is always simple, it cannot be a redundant eigenvalue. Hence it is a quotient eigenvalue, and, as those are a subset of the parent eigenvalues, it must still be the largest (hence the Perron-Frobenius) eigenvalue of the quotient. Its eigenvector can then be lifted to the parent network, by repeating entries on orbits. That is, if (λ, v) is the Perron-Frobenius eigenpair of the (left) quotient, then (λ, Sv) is the Perron-Frobenius eigenpair of the parent network ((21)). In practice, we use the symmetric quotient $B_{\mathrm{sym}} = \Lambda^{-1/2} S^T A S \Lambda^{-1/2}$ for numerical reasons (Algorithm 7), obtaining significant reductions in computational times (Figure 4 in the Main Text). If A is not symmetric (but irreducible), Algorithm 7 (Algorithm 5 in Methods) gives the right Perron-Frobenius eigenpair of A, and replacing B_{sym} by $\Lambda^{-1/2} S^T A^T S \Lambda^{-1/2}$ and Rw by its transpose, we obtain the left Perron-Frobenius eigenpair of A.

Input: adjacency matrix A, characteristic matrix S
Output: (right) Perron-Frobenius eigenpair (λ, v) of A

$\Lambda \leftarrow \texttt{diag}(\texttt{sum}(S))$
$R \leftarrow S\Lambda^{-1/2}$
$B_{\text{sym}} \leftarrow R^T A R$
$(\lambda, w) \leftarrow \texttt{eig}(B_{\text{sym}}, 1)$ eigenpair of the largest eigenvalue
$v \leftarrow Rw$

Algorithm 7: Eigenvector centrality from the quotient network.

SI Weighted and Directed Networks

We have presented our results on undirected, unweighted networks and symmetric network measures in the Main Text. For networks with edge (or node) weights, labels or directions, we can restrict to symmetries respecting those, giving a smaller geometric decomposition, and fewer symmetric motifs (Theorem 1), but otherwise our results either directly apply, or can be easily adapted. This is also the case for an undirected network measure ($F(i,j) \neq F(j,i)$) as its network reprentation $F(\mathscr{G})$ is directed even if $\mathscr{G}$ is not. However, $F(\mathscr{G})$ still inherits all the symmetries of $\mathscr{G}$ and therefore has the same geometric decomposition, orbits, and symmetric motifs (as vertex sets) — see the Weighted and Directed Networks subsection in SI Symmetry in Complex Networks.

FURTHER ACKNOWLEDGEMENTS

Thanks to Yamir Moreno and Emanuele Cozzo, whose 'simple' question inadvertently prompted this lengthly answer more than a year later, and to Gareth Jones for reminding the author that S_6 has two conjugacy classes of subgroups of index 6. Special thanks to my wife Kate for her unwavering support.

References

1. Biggs, N. *Algebraic graph theory* (Cambridge University Press, 1993).

2. Godsil, C. & Royle, G. F. *Algebraic graph theory* (Springer, 2013).

3. Cvetkovic, D. M., Rowlinson, P. & Simic, S. *An introduction to the theory of graph spectra* (Cambridge University Press, 2010).

4. Brouwer, A. E. & Haemers, W. H. *Spectra of graphs* (Springer, 2011).

5. Barrett, W., Francis, A. & Webb, B. Equitable decompositions of graphs with symmetries. *Linear Algebra Appl* **513**, 409–434 (2017).

6. Francis, A., Smith, D., Sorensen, D. & Webb, B. Extensions and applications of equitable decompositions for graphs with symmetries. *Linear Algebra Appl* **532**, 432–462 (2017).

7. MacArthur, B. D., Sánchez-García, R. J. & Anderson, J. W. Symmetry in complex networks. *Discrete Appl Math* **156**, 3525–3531 (2008).

8. MacArthur, B. D. & Sánchez-García, R. J. Spectral characteristics of network redundancy. *Phys Rev E* **80**, 026117 (2009).

9. Xiao, Y., MacArthur, B. D., Wang, H., Xiong, M. & Wang, W. Network quotients: Structural skeletons of complex systems. *Phys Rev E* **78**, 046102 (2008).

10. Wang, J., Huang, Y., Wu, F.-X. & Pan, Y. Symmetry compression method for discovering network motifs. *IEEE/ACM Trans Comput Biol Bioinf* **9**, 1776–1789 (2012).

11. Xiao, Y., Wu, W., Pei, J., Wang, W. & He, Z. Efficiently indexing shortest paths by exploiting symmetry in graphs. In *Proceedings of the 12th International Conference on Extending Database Technology: Advances in Database Technology*, 493–504 (ACM, 2009).

12. Karalus, S. & Krug, J. Symmetry-based coarse-graining of evolved dynamical networks. *Europhys Lett* **111**, 38003 (2015).

13. Nyberg, A., Gross, T. & Bassler, K. E. Mesoscopic structures and the laplacian spectra of random geometric graphs. *J Complex Networks* **3**, 543–551 (2015).

14. Do, A.-L., Höfener, J. & Gross, T. Engineering mesoscale structures with distinct dynamical implications. *New J Phys* **14**, 115022 (2012).

15. Rotman, J. J. *An introduction to the theory of groups* (Springer, 2012).

16. McKay, B. D. Practical graph isomorphism. *Congr Numer* **30**, 45–87 (1981).

17. Newman, M. Networks: An Introduction. *Oxf. Univ. Press.* (2010).

18. Xiao, Y., Xiong, M., Wang, W. & Wang, H. Emergence of symmetry in complex networks. *Phys Rev E* **77**, 066108 (2008).

19. Gal, E. *et al.* Rich cell-type-specific network topology in neocortical microcircuitry. *Nat Neurosci* **20**, 1004 (2017).

20. Liebeck, M. W. Graphs whose full automorphism group is a symmetric group. *J Aust Math Soc* **44**, 46–63 (1988). DOI 10.1017/S1446788700031360.

21. Bitbucket repository. `https://bitbucket.org/rubenjsanchezgarcia/networksymmetry/` (2018).

22. Horn, R. A. & Johnson, C. R. *Matrix analysis* (Cambridge University Press, 1990).

23. saucy 3.0. `http://vlsicad.eecs.umich.edu/BK/SAUCY/` (2012).

24. McKay, B. D. & Piperno, A. Practical graph isomorphism, II. *J Symb Comput.* **60**, 94–112 (2014).

25. Piperno, A. Search space contraction in canonical labeling of graphs. *arXiv preprint arXiv:0804.4881* (2008).

26. Junttila, T. & Kaski, P. Engineering an efficient canonical labeling tool for large and sparse graphs. In *2007 Proceedings of the Ninth Workshop on Algorithm Engineering and Experiments*, 135–149 (SIAM, 2007).

27. Darga, P. T., Sakallah, K. A. & Markov, I. L. Faster symmetry discovery using sparsity of symmetries. In *Design Automation Conference, 2008 (45th ACM/IEEE)*, 149–154 (IEEE, 2008).

28. The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.8.7* (2017).

29. Estrada, E. & Rodriguez-Velazquez, J. A. Subgraph centrality in complex networks. *Phys Rev E* **71**, 056103 (2005).

30. de la Peña, J. A., Gutman, I. & Rada, J. Estimating the estrada index. *Linear Algebra Appl* **427**, 70–76 (2007).

31. Higham, N. J. & Al-Mohy, A. H. Computing matrix functions. *Acta Numer* **19**, 159 (2010).