

<https://doi.org/10.1038/s42005-024-01819-y>

Quantifying metadata relevance to network block structure using description length

Lena Mangold ^{1,2} ✉ & Camille Roth ^{1,2}

Network analysis is often enriched by including an examination of node metadata. In the context of understanding the mesoscale of networks it is often assumed that node groups based on metadata and node groups based on connectivity patterns are intrinsically linked. This assumption is increasingly being challenged, whereby metadata might be entirely unrelated to structure or, similarly, multiple sets of metadata might be relevant to the structure of a network in different ways. We propose the metablox tool to quantify the relationship between a network's node metadata and its mesoscale structure, measuring the strength of the relationship and the type of structural arrangement exhibited by the metadata. We show on a number of synthetic and empirical networks that our tool distinguishes relevant metadata and allows for this in a comparative setting, demonstrating that it can be used as part of systematic meta analyses for the comparison of networks from different domains.

Block structure in networks is characterised by the grouping of nodes on the basis of shared connectivity patterns^{1,2}. Such networks can be generated by Stochastic block models (SBMs)³ which—in turn—can be used as baseline models to infer block structure from observed networks. The latter is becoming increasingly popular due to a number of considerable advances in the development of SBM variants with closer resemblance of real network structure⁴, the introduction of flexible, nonparametric inference approaches⁵, and increasingly efficient inference algorithms^{6,7}. Blocks may take the shape of commonly studied mesoscale structures, such as assortative communities or internally cohesive clusters. However, other structural arrangements on the mesoscale, such as core-periphery structures, disassortative (bipartite) structures as well as (nested) combinations of the above, are also possible, owing to the relatively general definition of similarity of the SBM. It is often assumed that blocks—whichever specific structural arrangement they may have—correspond to a latent ‘meaning’, i.e. some external similarity exhibited by the nodes that has made it more likely for them to connect to other nodes in the network in a similar way, or vice versa. In practice, this ‘meaning’ is often attributed to additional information on the network nodes, which we call *metadata*. In the literature, node attributes available as part of network analyses have sometimes been assumed to be intrinsically linked to the network's structure, an assumption that—as has been demonstrated on multiple occasions^{8–10}—cannot be readily made.

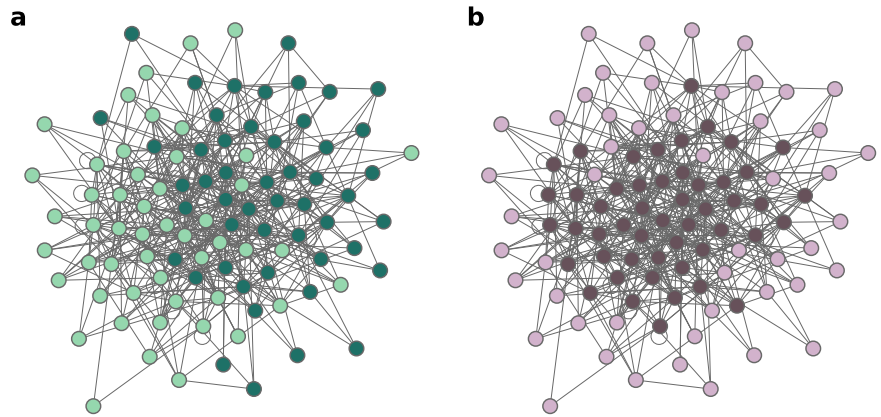
Imagine a set of users who interact with each other on some social media platform and for whom—in an idealised scenario—some metadata is known to us: for each user we know their preference for one of two political

parties. We can construct a network that represents the users' conversation around some political topic, by placing an edge between user nodes who interacted within the context of the topic. Assume we observe *homophily* in political leaning (called *assortativity* in network science): users with shared party preferences are more likely to endorse each other's content than that of users who support other parties, something that has been observed in the literature around online interactions repeatedly^{11–13}. The static ‘snapshot’ of an interaction network between these users, which might look something like the toy network in Fig. 1a, would show an alignment of a partition of users by party preference and one according to connectivity patterns. We could also say that the *generative process* of this network was, at least to some extent, likely governed by this particular set of metadata under an assortativity modelling assumption.

Such alignment between metadata partitions and block structure has been observed in social networks on many occasions^{14–16} which may partly explain the widespread assumption of an intrinsic connection between metadata and node structure. Motivated also by a lack of knowledge on the ‘real’ generative processes of empirical networks, node metadata has often been viewed as *the* ground truth for the block structure of a network, to evaluate^{4,17,18} or to improve^{19–23} the performance of community detection or other inference algorithms. However, a number of recent works in the complex networks community has challenged the notion of an intrinsic alignment between metadata and community structure^{8,9,24,25}; metadata might be entirely unrelated to structure or, similarly, multiple sets of metadata might be relevant to the structure of a network in different ways¹⁰. We can illustrate this on our example set of social media users for which we

¹Centre d'Analyse et de Mathématique Sociales (CNRS/EHESS), 54 Bd Raspail, 75006 Paris, France. ²Computational Social Science Team, Centre Marc Bloch (CNRS/MEAE), Friedrichstr. 191, 10117 Berlin, Germany. ✉ e-mail: lena.mangold@cnrs.fr

Fig. 1 | Two partitions of a toy network. **a** Example network described in main text, with nodes coloured according to their block membership in the planted bicomunity partition. **b** The same example network, with nodes coloured according to their block membership in the planted core-periphery partition. Both visualisations have been drawn using the graph-tool library³⁷, which is used for all network visualisations in this paper.



imagine a second set of node metadata: whether or not a user is an expert on the political topic that is being discussed in this particular network. Assume that in terms of structural features, the network exhibits some assortativity (of party preference) but the network also has a relatively well connected core of users—in which there are connections even between users of different party preference—and a loosely connected set of peripheral nodes. This structure then corresponds to the node attribute of expert (core) vs non-expert (periphery), an example of this can be seen in Fig. 1b. Overall, we thus have two sets of metadata that are both linked to the network structure while exhibiting different structural arrangements.

Besides emphasising the possible existence of multiple relevant metadata, this example can also be understood in the context of the diversity of likely node partitions exhibited by real networks: when focusing purely on the connectivity patterns in a network we can, in many cases, identify multiple, potentially qualitatively different partitions that divide the nodes in a ‘plausible’ way²⁶. This, in turn, accentuates the notion that multiple sets of metadata can be related to the network structure, even if they divide the network’s nodes in very different ways. It has been demonstrated, for example, that we can generate synthetic networks whose mesoscale—like the one in our toy example—is similarly well explained by (a) a division into two assortative communities and (b) a division into a well-connected core and a sparsely connected periphery²⁷. By fitting an SBM to such a network and sampling from the posterior distribution of partitions, we are likely to find two locally optimal partitions—a bicomunity and a core-periphery partition—which may be aligned with different sets of metadata. In other words: not only might multiple sets of metadata be relevant to the network structure in general, but they might be relevant in structurally very different ways.

To add to the complication, we can extend our social media thought experiment by considering the ‘reversed’ situation: instead of one network representing a political discussion with multiple sets of metadata, we imagine multiple networks with node attributes of one particular type. Firstly, we imagine that we are able to construct *multiple* networks for the same set of user nodes, such as networks of different interaction types (e.g. endorsement vs exchange of opinions) or snapshots recorded at different points in time. Secondly, we might have collected data on different topics from a variety of categories, leading to the construction of multiple topic-induced networks with different (but potentially overlapping) user node sets but shared type of metadata. To summarise and in all generality, we can divide our thought experiment into three scenarios worth thinking about: when studying social networks, we might be dealing with a situation with (I) a single network and multiple sets of metadata, (II) one set of metadata and multiple networks with the *same* user sets, and finally (III) one set of metadata and multiple networks in the same *context*, but with possibly different node sets.

Having outlined three types of scenarios, we come back to the questions we might be asking to investigate metadata-block structure relationships: Is there a way to measure the strength of the relationship between

metadata and network block structure, to enable comparisons between sets of metadata for one given network (I) and multiple metadata-network pairs (II and III)? If a set of node attributes is relevant to the structure of a particular network, can we quantify the particular structural arrangement at hand (assortativity vs some other structure)? Answering these questions requires rigorous, statistically grounded methods and we argue that generative models lend themselves particularly well in this context: they provide mathematical justification, help us deal with uncertainties connected to model selection and facilitate comparisons of different models and networks.

To our knowledge, we lack a framework that makes it possible, for a single given network and in a comparative setting for a collection of networks, to appraise the relative strength of the connection between various sets of metadata and various types of intermediate-scale structures. Related work that has gone furthest in this direction, and that serves as a motivation for our measure, is that by Peel et al.¹⁰. The authors demonstrated convincingly that multiple ground-truth partitions can be responsible for the generation of a given network. They proposed two separate approaches, one to measure the statistical significance of the metadata-structure connection and one to explore the relationship between specific sets of metadata and the partition landscape of a network. Their first measure (which we outline later on in this work) serves as a *p* value and thus answers a simple yes-no question with respect to the significance of metadata-structure relevance; their second measure is an inferential approach based on SBMs, which—upon visual inspection of the results—provides insights into the extent to which different sets of metadata are related to different parts of the partition landscape of a network. While these measures enable an in-depth analysis of individual networks and their metadata, they cannot be easily used for direct comparative purposes, since the strength of metadata-structure relationships cannot be measured and visual analysis is required for comparative studies of multiple sets of metadata; a direct comparison between networks is also not possible. Another relevant method is a label propagation approach to calculate a *p* value for the significance of metadata and structure²⁸ has similar problems and thus does not lend itself well for large-scale comparative meta analyses of multiple networks either.

In this work, we propose the *metadata block structure exploration* (*metablox*) tool, which utilises microcanonical SBMs and methods from information theory to quantify the connection between node metadata and the block structure of a network. Our measure exploits the feature of the minimum description length (MDL) principle²⁹, which penalises overly complex models, and enables comparison across multiple sets of network-metadata pairs with respect to the *strength* of the relevance of the metadata to the block structure and regarding the prominent *type* of block structure exhibited by the metadata. We design the measure to enable such comparative analyses for the scenarios I-III discussed above.

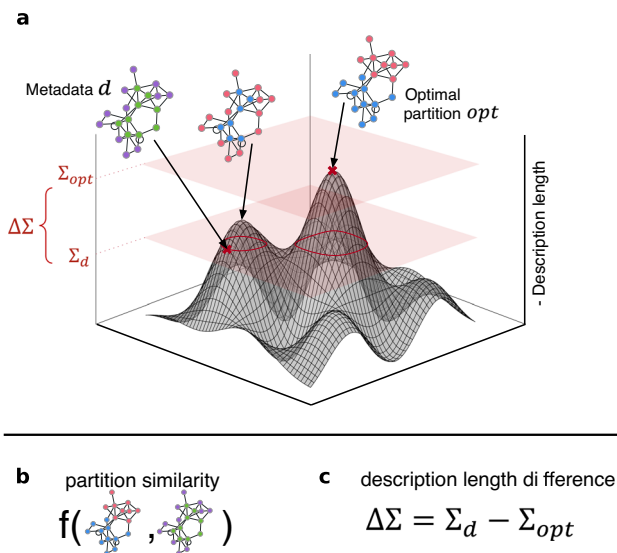


Fig. 2 | Schematic of a partition landscape. **a** Partition landscape^{10,26} for a toy network, with the negative description length on the vertical axis, for which we have highlighted the positions of the optimal inferred partition (blue/red), a second partition (blue/red), and a metadata partition d (purple/green). **b** Approach of measuring partition similarity directly. **c** Proposed metabox approach, of measuring the distance of the network's description length under the metadata partition from the network's description length under the optimal partition.

Results

General approach

To derive our measure, we first note that a set of categorical node attributes naturally divides the nodes of a network into a partition. We will refer to this division of the nodes as the *metadata partition* and denote it by the vector $\mathbf{d} = \{d_i\}$, where d_i denotes the metadata label of node i . The perhaps most obvious way to measure the relationship between a metadata partition and the block structure of a network would be to infer the optimal partition \mathbf{b}_{opt} of the network in some way and then measure the alignment between the metadata partition and the optimal partition through a similarity measure $s = f(\mathbf{b}_{opt}, \mathbf{d})$, such as the Rand index³⁰, variation of information³¹, or partition overlap²⁶. However, this approach may fail as soon as the detection method yields multiple similarly likely partitions that are not well-aligned in terms of their node groups: if the metadata partition is similar to one of the highly likely partitions that is not the optimal one, we would not capture the metadata relevance.

Instead of directly measuring partition similarities, we take a Bayesian approach to the problem. Specifically, we note that in the Bayesian SBM framework, one generally considers the entire distribution over the possible partitions that could have been responsible for generating the network under the given SBM. With this in mind, we hypothesise that a metadata partition could—in theory—appear among the partitions in the posterior distribution of the SBM with some non-zero probability. Unfortunately, calculating the posterior distribution of the SBM exactly turns out to be an intractable problem and one therefore needs to resort to approximations, e.g. by sampling from it using an MCMC method. Since we cannot rely on finding an exact copy of our metadata partition among the samples taken from the posterior, we instead take the reverse approach: we determine where in the posterior distribution the metadata partition *would* fall, were it to have been inferred by fitting an SBM to our observed network. The schematic in Fig. 2a shows an idealised partition landscape of a synthetic network (upon fitting an SBM) and illustrates the idea of the metadata being positioned more closely—in terms of description length (i.e. model fit)—to a non-optimal partition than to the optimal partition; in panels 2b and 2c, we show side-by-side the approach of directly measuring partition similarities described above and the approach we will be taking in this work, respectively.

Microcanonical SBMs and description length. To lay the groundwork for this approach, we first outline the microcanonical SBM framework and its connection to the concept of description length, which we will use to quantify the metadata relevance. We focus on the so-called degree-corrected SBM⁴, which configures an ensemble of networks with block structure, while making it possible to take into account heterogeneous degree distributions (in contrast to the traditional SBM). For a network with N nodes and B blocks, it is parameterised by an N -dimensional block membership vector $\mathbf{b} = \{b_i\}$, a degree sequence $\mathbf{k} = \{k_i\}$, and a $B \times B$ edge count matrix $\mathbf{e} = \{e_{rs}\}$. In its microcanonical form, the degree sequence and edge count matrix are specified *exactly*, making it possible to count the number of possible networks that can be generated with a set of parameters, so that the likelihood $P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})$, i.e. the probability of observing a network \mathbf{A} given the SBM parameters, can be calculated as $\Omega(\mathbf{k}, \mathbf{e}, \mathbf{b})^{-1}$, the inverse of the ensemble cardinality³². When using a parametric framework to detect the optimal partition from a network, one would find the partition \mathbf{b} that maximises the log-likelihood, which amounts to minimising the microcanonical entropy $S = \ln \Omega(\mathbf{k}, \mathbf{e}, \mathbf{b})$ ^{32,33}.

In the *nonparametric* approach—which allows for the number of model parameters to be inferred from the data—, one considers the joint distribution of the generative model, including the notion of priors on the model parameters in addition to the model likelihood. The hard constraints of the microcanonical approach imply that one does not need to sum over the remaining parameters to calculate the marginal likelihood. The joint distribution can therefore be written as $P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) = P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}, \mathbf{e})P(\mathbf{b})$. This is where an information theoretical lens lets us write $P(\mathbf{A}, \mathbf{k}, \mathbf{e}, \mathbf{b}) = e^{-\Sigma}$, where $\Sigma = -\ln P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) - \ln P(\mathbf{k}, \mathbf{e}, \mathbf{b})$ is called the description length of the data: the amount of information necessary to describe a network given a model plus the information required to describe the model itself, via its parameters⁷. Maximising the posterior distribution is therefore equivalent to minimising the SBM description length, which—in turn—is the same as finding the partition which provides the most compact *compression* of the network and therefore amounts to being the best model. Similarly, one can identify the more likely of two partitions \mathbf{b}_1 and \mathbf{b}_2 (under SBM variants m_1 and m_2 respectively) by comparing the network's description length for each of them, which amounts to calculating their posterior odds ratio. For $\Sigma_1 < \Sigma_2$, for example, partition \mathbf{b}_1 (under variant m_1) is the more likely of the two (see the “Methods” section for a detailed explanation). Note that we use description length as the basis for model selection (i.e. finding the best model by penalising overly complex modelling choices) since it is the natural choice when considering a posterior inference framework, but in certain circumstances other possible model selection criteria could be considered as well. In particular, we favour it over popular existing information criteria (such as AIC³⁴ and BIC³⁵), which are built on assumptions that do not hold for comparisons between the standard and degree-corrected SBM³⁶. Another commonly used model selection approach is cross-validation, which is widely used in machine learning and certain network science applications that allow for splitting the data into test and training sets (e.g. link prediction). While cross-validation can help reject models that are overly complex by favouring those that generalise better to unseen data, it is not inherently designed to penalise model complexity as directly as criteria like description length.

Measuring metadata relevance. To tie in these concepts with our metabox measure, we realise that one can simply count the nodes in each *metadata block* (i.e. those that share the same metadata category), and the edges within and between them, and arrive at quantities that are equivalent to the parameters of the microcanonical SBM. Given the metadata partition by \mathbf{d} , we denote the edge counts within and between metadata blocks by \mathbf{e}' . By plugging the quantities derived from the metadata and from the observed network into the SBM description length calculations—which, in the microcanonical case, can be done exactly by using combinatorics³—we can straightforwardly calculate Σ_d^m , the description length of the network under the metadata partition \mathbf{d} and a given SBM variant m . In this work, we use the description length

calculations from existing work⁵, which we provide in the “Methods” section. Note that a more consistent notation would be to denote a metadata partition by \mathbf{b}' , but we choose $\mathbf{d} = \mathbf{b}'$, to make a clearer distinction from the inferred partitions.

Calculating the metadata description length in this way does not yet tell us anything about the fit of the metadata *relative to the partition landscape* of the network. We therefore also separately fit the SBM variant m to our observed network (ignoring the metadata labels), using the graph-tool library³⁷, and compute the description length Σ_{opt}^m of the optimal partition. We can then simply calculate the description length difference of the two quantities, $\Delta\Sigma = \Sigma_d^m - \Sigma_{\text{opt}}^m$ to understand how well the metadata partition fits the given network, compared to the optimal partition: if $\Delta\Sigma$ is close to zero we conclude that the metadata partition is strongly relevant to the block structure; if $\Delta\Sigma \gg 0$, the metadata partition fits the network much less well than the optimal partition, under the SBM variant m . Note that for better readability, we dropped the superscript m for $\Delta\Sigma = \Delta\Sigma^m$.

Statistical significance and comparability. There are two major caveats of quantifying the metadata relevance by simply using $\Delta\Sigma$: description length increases with growing networks size, which means that simply taking the absolute description length difference does not enable inter-network comparisons and we therefore need to normalise the measure in some way. Furthermore, we need to include a notion of statistical significance, to ensure that the fit of the metadata partition is, in fact, better than a partition we would find at random. We address both caveats at once by turning our measure into a ratio and including a distinction from randomness as follows: additionally to Σ_d^m we calculate the description lengths of the same network under an SBM with multiple sets of *randomised* metadata. Each time we randomise the nodes’ metadata labels, we can plug in the resulting quantities \mathbf{d}_* (the randomised metadata partition) and \mathbf{e}_* (the edge counts between and within the nodes in each randomised metadata category) in the description length formulas again. This approach is inspired by the blockmodel entropy significance test (BESTest)¹⁰, which produces a metadata p value, i.e. the probability that a randomised version of the observed metadata describes the network better than (or equally well as) the observed metadata. The BESTest p value is calculated by generating a large set of randomised metadata partitions (whereby the number of elements in each metadata category is fixed), computing the SBM entropy under each of the randomised partitions, and finally comparing the SBM entropy under the observed metadata partition to that of the randomised partitions. However, SBM entropy decreases for an increasing number of groups which greatly complicates direct comparison of metadata models (i.e. different metadata partitions and/or different SBM variants). Description length enables model selection and can thus be used to compare different partitions, whether induced by metadata categories or inferred based on connectivity patterns.

Instead of calculating a p value from distributions of description lengths, we identify the *randomised metadata description length* Σ_*^m , which—for a significance level of $\alpha = 0.01$ —is equal to the first percentile of the description length distribution of the randomised metadata partitions under SBM variant m . The choice of α is, to some extent, arbitrary; here, we make the choice based on conventions in measuring statistical significance, according to which $\alpha = 0.01$ denotes strong evidence against the null hypothesis, i.e. that the network is described equally well by randomised metadata. We emphasise that other choices can be made when our measure is used, as the required level of significance may depend on the specific context. In our measure, we take Σ_*^m to be the maximum description length for which we would still consider the description length of the *observed* metadata to be relevant, and we use it to normalise $\Delta\Sigma$ by $\Delta\Sigma_* = \Sigma_*^m - \Sigma_{\text{opt}}^m$.

Probing structural arrangement. Before we put these components together, we note that we can use the above approach to formulate a measure that not only quantifies metadata relevance but that also helps us probe for specific *types of structural arrangements* of the metadata blocks:

Since we can follow the above procedure for any SBM variant for which there is a microcanonical formulation with a straightforward description length calculation, we can directly plug in the structure-specific SBM variant of interest. In this work, we focus on three SBM variants that are specific to three different types of block structures, which we briefly outline in the following. In the ‘standard’ SBM (i.e. non-degree-corrected, NDC), the placement of an edge between two nodes depends solely on the block membership of each node and on the probability of two nodes from the two blocks being connected. Similar to the random graph model, one major drawback of this model is that the generative process produces networks that have blocks within which node degrees are Poisson distributed. To overcome this caveat, that makes the model unlike many real networks whose node degrees often have power-law degree distributions, the degree-corrected variant (DC)—which we introduced above to motivate our approach—was proposed⁴. In this variant, edge placement depends on node degrees as well as block membership, thus accounting for heterogeneous degree distributions. A third variant we include in our analysis is the assortative ‘planted partition’ SBM (PP), which—as part of the generative process—assumes assortativity³⁸. The description length formulations for the three variants used in our measure are based on the work in refs. 5,32,38, as detailed in the “Methods” section.

Note that we are focusing here on undirected graphs, so that any network for which our measure is used has to be treated as such. An extension of our measure to directed networks is straightforward and should be considered as part of future research.

Measure. We finally put together these components to create the metadata block structure exploration (metablox) tool, which produces the vector γ_d for a metadata partition d , where each element represents the metadata relevance under a different SBM variant m . It can consist of as many elements as we include SBM variants to probe for different structural arrangements. Here, we will work with the three variants outlined above (DC, NDC, and PP), thus γ_d is a vector with three elements. For a metadata partition d , the metablox vector γ_d thus consists of elements

$$\gamma_d^m = \frac{\Delta\Sigma_d^m}{\Delta\Sigma_*^m} = \frac{\Sigma_d^m - \Sigma_{\text{opt}}^m}{\Sigma_*^m - \Sigma_{\text{opt}}^m} \quad (1)$$

for each SBM variant m . In summary, Σ_d^m is the description length of metadata partition d under model m , Σ_*^m is the randomised metadata description length, and Σ_{opt}^m is the description length of the optimal partition of the network. Recall that the difference in description length between two models is equal to the log of the posterior odds ratio of the two models. The numerator of γ_d^m thus measures how much more likely the optimal partition is compared to the metadata partition under model m . Similarly, the denominator represents how much more likely the optimal partition is to the randomised metadata partitions under all models. The total measure is the former normalised by the latter.

With this definition, each element can then be interpreted as follows: for $\gamma_d^m \geq 1$ we say that this set of metadata is not relevant to the structure under model m , since more than 1% of the randomised metadata partitions compress the network more efficiently than our observed metadata, under the given model. For $\gamma_d^m < 1$, the metadata is relevant to the structure under the given model, and the closer γ_d^m is to 0, the stronger the relevance of the set of metadata to the block structure and the closer the structural arrangement of the metadata to the typical type of structure generated under the given model m . Another way of interpreting γ_d^m is: relative to the best compression we can find, how efficient is the compression of the observed metadata d under m compared to partitioning the network by randomised versions of the metadata.

Limitations and network compressibility. One limitation of our approach is that the inequality $\gamma_d^{m_1} < \gamma_d^{m_2}$ for a network with some metadata partition d merely indicates that the metadata d is relatively closer to the best partition under model m_1 than to the best partition under model

m_2 . It does not necessarily mean that the block structure of our network as a whole is optimally explained by m_1 : since our measure is a ratio of description length differences, it is possible that in fact $\Sigma_{\text{opt}}^{m_1} > \Sigma_{\text{opt}}^{m_2}$ or that the model m_2 provides a better explanation for the network overall. This limitation arises from the comparative nature of the measure, which inherently focuses on relative rather than absolute fit to enable comparisons across metadata and networks.

Another limitation is related to the case of inter-network comparison, in particular in the context of different network topology. While metablox is robust with respect to the number of nodes in a network (see the “Methods” section), certain aspects of a network’s topology do affect the measured metadata relevance. Let us assume that we are comparing two node attributed networks A and A' , with the signal in the community structure of network A being significantly stronger than that of network A' (i.e. the communities of network A are more clearly separated). All other things being equal—including the level of alignment of the metadata partition with the optimal detected partition—we will have $\gamma < \gamma'$. Since Σ_{opt}^m decreases when the network becomes more compressible due to a stronger signal in the block structure, $\Delta\Sigma^*$ increases (since Σ_{opt}^m does not decrease with Σ_{opt}^m) and does so more quickly than $\Delta\Sigma$. This apparent conflation of block structure signal strength and metadata block-structure relevance points to a question around the definition of this ‘relevance’, which in turn relates back to our earlier point of going beyond the direct measurement of partition alignment. We argue that the concept of measuring the relevance of metadata to block structure in fact *should* include the notion of block structure strength: we do not want to retrieve the same γ for the described networks A and A' merely because there is alignment of metadata partition and optimal partition. Instead, the relevance of the metadata should be quantified as stronger (i.e., a lower γ) when the network exhibits more significant block structure.

The elements of the measure that lead to these limitations are part of the design that enables comparative work and what enables the measure to capture metadata-structure relationships in the intended way. However, there might be reasons for a researcher using this measure to require insights – alongside the metablox vector γ_d – into the absolute fit of the optimal partition, which acts as the main point of comparison in the partition landscape, or to disentangle block structure signal and metadata partition correlation. To account for this, the edge compression $c^m = \Sigma_{\text{opt}}^m / E$ can be referred to as a second dimension in our measure, which quantifies the absolute compression of the optimal partition per network edge. It can serve as a point of reference for the structure of the optimal partition and help us understand how much of γ is explained by correlation vs compressibility of the network under a given model. In the “Methods” section we include an analysis of the sensitivity of our measure to network topology and the role of the second dimension.

Simulations

We now summarise the metablox pipeline and state the choices we make for any calculation of γ_d in the remainder of this work: We calculate Σ_d^m by

plugging the required quantities, as obtained by the metadata partition \mathbf{d} , into the description length formulas for SBM variant m . For Σ_{opt}^m , we use the graph-tool library to fit variant m to the observed network and find the partition $\mathbf{b}_{\text{opt}}^m$ that minimises the description length, refining the result of the agglomerative algorithm by running 1000 sweeps of the merge-flip MCMC⁷. We calculate Σ_{opt}^m by performing 500 permutations of the metadata labels, calculating the description length of the network under each of them, and finding the description length for which only 1% of permutations have a lower description length.

Synthetic network. To illustrate the metablox measure, we generate a network with multiple, non-aligned, plausible explanations in terms of its block structure, create multiple sets of synthetic metadata and calculate metablox on two dimensions: under a degree-corrected (DC) and a planted partition (PP) SBM. More specifically, we generate the network with the Stochastic cross block model (SCBM)²⁷, which facilitates the generation of such networks with ‘ambiguous’ mesoscales structures, by ‘planting’ (similar to the planted partition model³⁹) multiple coexisting partitions. We use the SCBM to generate a network with $N = 100$ nodes, $B = 2$ blocks and expected degree $k = 10$. We plant two coexisting partitions, so that the network exhibits bicomunity (BC) structure (i.e. two assortative communities) as well as core-periphery (CP) structure (see the “Methods” section for details on how this network was constructed). Figure 1 shows this network in two separate visualisations, with nodes painted according to their block membership in the BC partition in Fig. 1a and according to their block membership in the CP partition in Fig. 1b. We also generate multiple sets of synthetic metadata, in a way such that the labels align with the block structure of the network to varying degrees, similar to related work¹⁰. We correlate half of the synthetic metadata with the BC partition and the other half with the CP partition, with varying level of correlation ρ . The metadata label of each node is equal to the node’s block assignment of the planted structure with probability $(1 + \rho)/2$. We increase ρ from 0 to 1 at steps of 0.01, yielding a total of 202 sets of metadata, half of which being ‘bicomunity-like’ and the other half being ‘core-periphery-like’.

In Fig. 3, each dot represents one set of metadata d for which we are plotting γ_d^{DC} and γ_d^{PP} for increasing correlation ρ , for the bicomunity-like metadata (Fig. 3a) and the core-periphery-like metadata (Fig. 3b). For Fig. 3 only, we additionally calculate the BESTest p value to illustrate the contribution of metablox compared to the most similar existing measure¹⁰.

For the bicomunity-like metadata, we can see that both dimensions of γ are a decreasing function of the correlation ρ : the more strongly correlated the metadata, the higher the relevance assigned to it by the metablox measure. As expected, we cannot make a distinction between the general (DC) and assortative (PP) SBM variant, since—by design—the metadata in this case is similar to the block structure that was created under an assortativity assumption. For core-periphery-like metadata, however, γ_d^{DC}

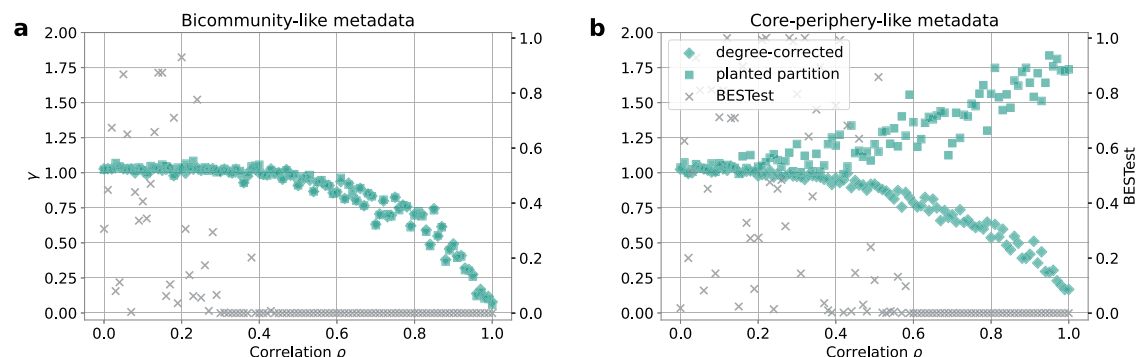


Fig. 3 | Metablox and BESTest values for a synthetic network. Metablox values for the degree-corrected and planted partition SBM, and BESTest values¹⁰ under the degree-corrected SBM, for a synthetic network with multiple sets of metadata.

a Values for the sets of bicomunity-like metadata. **b** Values for the sets of core-periphery-like metadata. Both panels show an increasing correlation ρ between metadata and block structure on the x-axis.

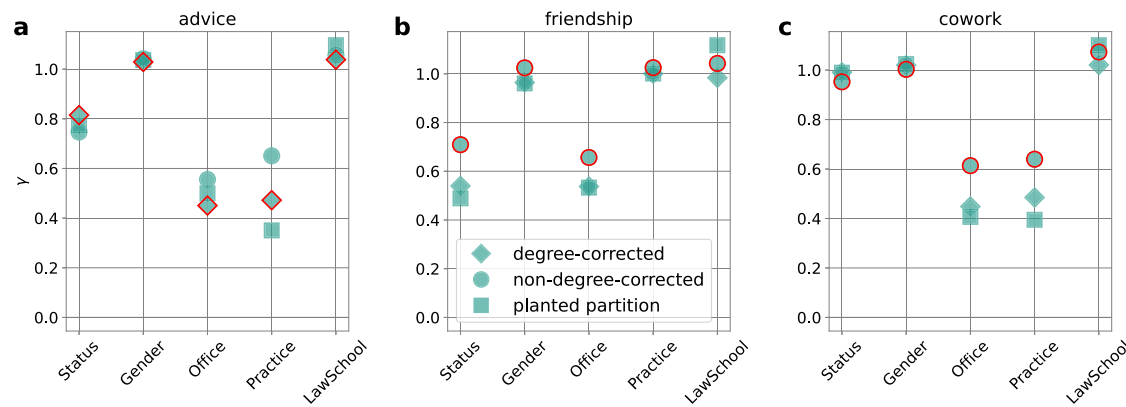


Fig. 4 | Metablob values for multiple metadata on law firm networks. Degree-corrected, non-degree-corrected and planted partition dimensions of the metablob vector for each of five sets of metadata (status, gender, office, type of law practised,

and law school attended) for three networks of employees of a law firm⁴⁰. **a** Advice network. **b** Friendship network. **c** Coworking network. On each figure, the SBM variant that gives the lowest edge compression for the network is highlighted in red.

decreases similarly to the previous case, while γ_a^{PP} is now an increasing function of ρ , suggesting that—as we know is true—if we were to assume that this metadata was responsible for generating this network, assortativity would be much less likely to explain its structure.

For comparison, we observe that the BESTest values become non-distinguishable once the correlation ρ has crossed a certain threshold. When the SBM entropy under a given metadata partition is lower than the SBM entropy under *all* sets of randomised metadata partitions then the p value is $1/n_p$ (where n_p is the number of randomised metadata partitions) and we cannot compare this metadata to another set of metadata for which this is also true. For example, BESTest does not allow us to compare a metadata partition that is nearly perfectly correlated with the planted block structure, to one for which $\rho = 0.7$, in both the bicomponent and core-periphery case. For comparison, the metablob values for DC cross the significance line $\gamma = 1$ at a similar point as BESTest becomes significant, for both the bicomponent and core-periphery-like metadata. However, it proceeds in the shape of a continuous decreasing function with a minimum of 0 at $\rho = 1$, which is what we would expect as the metadata partition is equal to the planted block structure at this point. The first contribution of metablob is therefore that, by comparing the metadata partition to the *optimal* partition (while still taking into account statistical significance), we are able to make direct comparisons between metadata sets. The second contribution is illustrated by the fact that BESTest does not tell us anything about likely structural arrangements of the metadata, while the difference between the trajectory of γ_a^{PP} on the left and right figure demonstrate the way in which metablob does allow for this.

We also apply the metablob measure to three real-world networks, in each case referring to one of the three scenarios (I–III) introduced above. As outlined previously, our measure can be extended to include any number of structure-specific SBMs. For our applications to real networks, we focus on the non-degree-corrected (NDC), the degree-corrected (DC), and the assortative ‘planted partition’ (PP) SBM. This means that we calculate a three-dimensional metablob vector for every network-metadata pair.

Application: law firm. We first demonstrate our method on a number of networks from the Lazega law firm networks collection⁴⁰. In this collection, there are three different networks among the same set of nodes, in which edges represents different types of connections between the employees of a corporate law firm: coworkers, friendship, and advice. For the employees, we have five sets of node attributes - their status in the firm, gender, the office in which they work, which type of law they practice, and which law school they attended. Figure 4 shows the metablob results for the three networks, each on a separate panel. On each panel, we show the results for the respective network and compare the five different metadata partitions. Each panel thus represents a case of

scenario I, while we can also compare network-metadata pairs which share a type of node attribute by comparing across panels (scenario II). Note that we have highlighted the SBM variant which gave the best per edge compression with a red marker. We can see some considerable variation in how the sets of metadata are related to the networks’ block structure under each of the SBM variants. Firstly, we observe that the strongest metadata-block structure relevance can be observed between the type of law practiced by an employee and the block structure in the advice network, under an assortative SBM assumption. This implies that employees are more likely to seek advice from colleagues who practice the same type of law as themselves. The same set of metadata is also strongly relevant in the cowork network. However, in the friendship network the law practiced by an employee does not seem relevant to the formation of blocks under any of the SBM variants. Employees’ status, on the other hand, is the most strongly relevant metadata out of all attributes we have for the friendship network, under PP and DC. This indicates that a division of the network’s nodes into the two available categories (partner and associate) in this set of metadata is more strongly aligned with a plausible partition according to shared connectivity patterns in the network, under the modelling assumptions of both PP and DC. In other words, the nodes that share a metadata attribute (partner or associate) are not only likely to share connectivity patterns independent of their node degrees (DC), they are also more likely to connect to employees with whom they share the status than with others (PP). Status matters less in the advice network (albeit still significantly relevant) and is essentially irrelevant in the cowork network. The office in which employees were located is the only set of metadata that is highly relevant for all three networks, indicating that the physical location of employees’ plays a large role in determining connections. Both employees’ gender and the law school they attended are not related to the block structure under any SBM variant and therefore appear to be essentially irrelevant in the tie generation process. In this example, we have been able to both compare different sets of metadata for a given network (scenario I) as well as different networks with the same node set and shared metadata (scenario II).

Peel et al.¹⁰ used the same networks to demonstrate both of their methods, which we outlined above. Using the BESTest significance test, they reached similar conclusions for the network-metadata pairs. Using a second method – which admittedly relies on visual interpretation and does not explicitly serve to compare the strength of metadata relevance—the authors concluded that for the friendship network, the law school attended by employees was more strongly structure-relevant than the office in which employees are located. This finding is the opposite of that given by metablob. While the methods proposed by the above authors can provide insights into the significance of metadata relevance and the quality of the relationship for a given network-metadata pair, they do not enable a *direct quantification* of




Fig. 5 | Metablox values for multiple Twitter/X networks. Degree-corrected, non-degree-corrected and planted partition dimensions of the metablox vector for various Twitter/X networks. **a** Static snapshots, representing a non-overlapping one year period each, of a Twitter/X retweet network among users discussing the topic of

impact investing, with user location (country) as shared metadata (scenario II).

b Three Twitter/X interaction networks^{43,44} with shared metadata representing the users' political stance (liberal vs conservative) (scenario III). The SBM variant that gives the lowest edge compression for each network is highlighted in red.

the strength of the relationship and of the likely prominent structure. As a result, they also do not enable a direct comparison of different networks (either made up of different or the same node sets), something our measure is designed to do as shown in the following paragraphs.

Application: impact investing. We continue with a second example for scenario II, in which we investigate the metadata-block structure relationship in Twitter/X retweet networks of users discussing the topic of ‘impact investing’⁴¹. In October 2021 we exhaustively collected tweets mentioning any term belonging to a small subset of hashtags directly relevant to this topic. This amounted to 1.89 m tweets published by around 300k accounts since 2007. We deemed users who posted less than a tweet a month on average or who never posted more than one year apart to be insufficiently active in that field and filtered them, leading to a core set of around 16k users. We subsequently inferred likely home countries from the geolocation voluntarily shared by users in their profile description. Most users appeared to come from an English-speaking country: countries with at least 1000 active accounts include only the US, Canada, Australia and the United Kingdom, covering 72% of the data. Focusing on these country users, we eventually study 13 static networks representing each a snapshot for each year between 2009 and 2021 (we excluded 2007 and 2008 which yielded trivial networks)⁴².

This time, the edges in the different networks represent the same type of interaction at different moments in time. The metadata here is the country which the user declares being located in and which we consider to be fixed over time, being inferred at the moment of data collection in 2021. We calculated the three metablox dimensions for each network snapshot and also highlighted, for each network, the SBM variant that provided the best compression per edge for its optimal partition. In Fig. 5a, we can see that the country in which a user is located starts being relevant to the network's block structure in 2010, under all SBM variants but most strongly under PP. In the next five years (until 2015) the metadata becomes more strongly relevant in general and assortativity remains the prominent arrangement of the metadata. In 2016, there seems to be a change in the structural organisation of the network in relation to the metadata: the country in which users are located remains relevant to the block structure but in 2017 and 2018, PP is replaced by DC as the best fit for the metadata, before it swaps again starting in 2019 up until the last year of data collection. It is worth noting that for each year, DC is the SBM variant for which the optimal inferred partition provided the best compression. Overall, the quantification of the relevance of the metadata afforded by metablox enables us to observe changes over time that should serve as a starting point to ‘zoom into’ the networks for those particular events to gain a deeper understanding of the relationship between the users' location and their tweeting behaviour: the change from 2009 to 2010 when the country user location becomes

significant, the years 2014 and 2019 when the metadata is most strongly relevant under PP; and the years 2016–2018, in which the networks seem to exhibit a change in the metadata-structure relationship.

Application: Twitter/X. In our final example, we proceed with another set of networks created from Twitter/X, this time demonstrating metablox for networks in the same context and not necessarily with the same set of nodes (scenario III). In particular, we use data from three debates on political topics in the US, for which we are studying the metadata-block structure relationship for the same type of node attributes. The tweets from which these networks were created were collected between 2015 and 2016, and are based on conversations on abortion, Obamacare, and gun control^{43,44}. In all three networks, the available node metadata reflects two categories of political orientation as either liberal or conservative. The liberal-conservative opinion categories are based on a continuous score between -1 and $+1$ that were calculated⁴⁴ based on a method that uses URLs shared by Twitter accounts and the categorisation of websites behind these URLs on <https://mediabiasfactcheck.com>⁴⁵. The two categories used as node metadata are based on users below and above a ‘neutral’ score of 0. Note that we do not know to what extent there is an overlap between the set of users participating in the debates represented by the three networks, neither do we care; here, the subject of comparison is the topic that is discussed in each network and to what extent the political stance of users is related to the block structure for each topic.

In Fig. 5b, we plot the metablox vectors for these interaction networks with binary political orientation metadata. We observe that for the topics of gun control and abortion, the metadata partition into liberals and conservatives is relevant under all variants but most strongly under PP, while DC provides—again—the best overall fit. For the Obamacare network, the metadata is barely relevant under NDC and not significantly relevant under the other variants. We can conclude that for gun control and abortion, the metadata partitions were likely to be at least somewhat related to the edge generating process and that assortativity is the prominent structural arrangement of the metadata in those cases. Our findings seem equivalent to previous findings on polarisation levels in these networks, which classified all three networks to be polarised to some extent with the Obamacare network being the least polarised⁴⁴. Similar to our work, their polarisation measure also considers both network structure as well as a metadata dimension. However, their measure considers continuous node attributes—specifically political ideology on a continuous left-right scale—and specialises on measuring polarisation rather than the more general approach we take here. So although our measure has a different purpose, it might be able to pick up specific network properties (such as polarisation or fragmentation) while also being general enough to enable broader comparisons.

Discussion

In this paper, we have introduced a measure for probing the relationship between network metadata and its structural organisation. Our metablox pipeline, which produces the vector \mathbf{y} , is designed to provide insights into the relevance of metadata to a network's block structure and the likely structural arrangement in various scenarios: for the comparison of multiple sets of metadata for one given network (scenario I) and to compare the metadata-block structure relationships for entire collections of networks that share the same type of metadata, for networks that have the same node set (scenario II) and for networks within the same context (scenario III).

We have demonstrated the metablox measure on a synthetic network and applied it to a number of real networks, including the Lazega law firm networks and networks representing different social media debates. The results reveal variations in the relevance of metadata partitions to block structure, and the likely structural arrangement, therefore providing insights into the underlying dynamics of these networks. By covering examples from the three scenarios (I–III), we have demonstrated that our measure allows for a comprehensive inter- and intra-network comparison, enabling researchers to quickly identify networks where specific metadata partitions are closely related to structure or where certain structural arrangements are likely or unlikely under the metadata.

We have discussed a number of limitations of the measure, related to the specific design of metablox as a comparative measure, and proposed ways to address these. Additional limitations are connected with the specific type of node attributes for which relevance can be measured. In its current framework, metablox can only quantify the relevance of non-overlapping categorical metadata, while overlapping, real-valued or multidimensional metadata cannot be evaluated. For the case of overlapping metadata, it is conceivable that an overlapping SBM variant may be used. For real-valued and multidimensional metadata, the options are less clear. We therefore see a possible research direction in an alternative version of the metablox tool for real-valued metadata. While this clearly requires a number of non-trivial decisions with respect to suitable generative models, providing an equivalent measure for non-discrete metadata seems a valuable extension.

Other potential research directions evolve around the particular type of structural arrangements of metadata categories. Currently, we focus on degree-corrected (DC), non-degree-corrected (NDC), and planted partition (PP) Stochastic block models. However, networks often exhibit more complex structural arrangements beyond these models and so might their metadata. Future work could explore the integration of other SBM variants that are tailored to specific structural motifs, such as core-periphery structures, bipartite structures or nested patterns. The current measure has been implemented for undirected graphs, but extensions to more complex network structures such as directed graphs are straightforward and should also be considered as part of future research.

Our measure has the potential to serve as a tool for conducting large-scale comparisons of collections of network-metadata pairs, for networks coming from a variety of research fields – as long as categorical node metadata is available. Obvious examples are: different types of social networks, for which metadata may include a range of demographics or affiliations; biological networks such as gene regulatory networks, protein-protein interaction networks, and ecological networks with metadata related to genes, proteins, or species; economic and financial networks, such as trade networks, supply chains, and stock market networks, for which metadata may relate to industries, sectors, or companies; or networks from science of science, where networks represent collaborations and knowledge flows, and metadata may include fields of scientific research.

Methods

Data compression & microcanonical SBMs

To justify the use of description length as part of the metablox formulation, we expand on the Results section in the main text and more thoroughly explain the minimum description length (MDL) principle and its relationship with microcanonical SBMs. MDL is a model selection criterion

according to which one should favour the model that achieves the smallest compression of the data. The idea behind this is that compression is possible when we find regularities in the data which, in turn, means that we ‘learn’ about patterns in the data²⁹. MDL is sometimes described as a formal interpretation of Occam’s razor—also known as the principle of parsimony – which is the idea that one should try to find the explanation with the smallest number of assumptions possible. In more formal terms, the best hypothesis H (e.g. a model with its parameters) for a data set D , is the one that minimises the sum $S(H) + S(D|H)$, where $S(D|H)$ is the amount of information required to describe the data D when it has been encoded with the hypothesis H and $S(H)$ is the amount of information necessary to describe the hypothesis itself. This demonstrates the ‘automatic’ overfitting-prevention property of MDL, which makes it an attractive model selection criterion: with a more complex hypothesis, we need less information to describe the data given the hypothesis, but we need more information to describe the hypothesis itself.

There is a strong relationship between MDL and Bayesian inference in general²⁹ and the Bayesian interpretation of the SBM more specifically, where it was first used to infer network partitions without knowing the number of blocks in advance⁴⁶. To calculate the description length of a network under a particular model, one needs to derive the entropy of the individual components of the SBM, which is an ensemble of networks that can be generated from a set of parameters (i.e. the partition). In general, *microcanonical* network ensembles—for which structural constraints need to be satisfied exactly rather than on average—can be described by their entropy $S = \ln \Omega(\theta)$, where $\Omega(\theta)$ is the total number of networks that can be generated under the given set of parameters θ ³³. The higher the entropy of a network ensemble, the more ‘disordered’ (or ‘random’) is the ensemble. More concretely, let us assume that we have a network A generated by a model with parameter set θ . $P(A|\theta)$ is the probability of observing the network A in an ensemble generated by the model with these parameters (i.e. the likelihood) and we assume that all networks occur with the same probability $P(A|\theta) = \frac{1}{\Omega(\theta)}$. From this assumption, one can straightforwardly make a connection between the microcanonical entropy S and the log-likelihood: $L = \ln P = -\ln \Omega(\theta) = -S$ ³². By minimising the entropy S one could therefore find the maximum likelihood parameters, such as the most likely partition, given an observed network. While maximum likelihood methods work well in many cases, in the particular case of model selection with SBMs, maximum likelihood estimation can lead to overfitting if the number of model parameters is not fixed. For example, if the number of blocks B is not known, minimising the entropy would lead to the trivial partition of every node being its own block, i.e. $B = N$, where N is the number of nodes. Peixoto proposed the MDL principle as part of a microcanonical *nonparametric* approach to SBM inference, to address the overfitting issue and enable the use of flexible priors and hyperpriors on the model parameters^{5,46}. Specifically, this means considering the full joint distribution of the network and the SBM model parameters as part of the inference process, rather than just the SBM likelihood, which can be interpreted as calculating the description length of the network under the given model, via its parameters.

We recall that for a network A of size N with E edges, we denote the degree of a node i by k_i , and we use the following notation for the parameters of a microcanonical SBM with B blocks. The block assignments of the nodes is denoted by the vector $\mathbf{b} = \{b_i\}$ of length N , the $B \times B$ matrix $\mathbf{e} = \{e_{rs}\}$ represents the edge counts within and between two blocks r and s (with twice the number of edges on the diagonal, as is convention), and the B -dimensional vector $\mathbf{n} = \{n_r\}$ describes the number of nodes in each block r . The full joint distribution of the degree-corrected SBM is, as introduced above, given by $P(A, \mathbf{k}, \mathbf{e}, \mathbf{b}) = P(A|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b})$. In terms of the generative process of this model, this means that one first samples a partition of the nodes into blocks, then samples the numbers of edges within and between the blocks, then samples the half-edges according to the node degrees, and finally connects half-edges accordingly to create the network (i.e. sampling from the networks that are possible given the partition and number of edges).

When using the nonparametric Bayesian framework for inference purposes, one can then use this formulation to maximise (or sample from) the *posterior* distribution of partitions

$$P(\mathbf{b}|\mathbf{A}) = \frac{P(\mathbf{A}, \mathbf{b})}{P(\mathbf{A})} = \frac{P(\mathbf{A}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{A})}. \quad (2)$$

This is where the microcanonical formulation helps simplify the inference problem, which allows us to write $P(\mathbf{A}|\mathbf{b}) = P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})$. The marginal likelihood is usually $P(\mathbf{A}|\mathbf{b}) = \sum_{\mathbf{e}} P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})$, but due to the ‘hard’ constraints of the microcanonical SBM, there is only one non-zero element in this sum⁴⁶.

It turns out that this is where we can re-introduce information theoretical interpretation. In particular, the posterior can be rewritten as

$$P(\mathbf{b}|\mathbf{A}) = \frac{P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b})}{P(\mathbf{A})} = \frac{e^{-\Sigma}}{P(\mathbf{A})} \quad (3)$$

where $\Sigma = -\ln P(\mathbf{A}|\mathbf{k}, \mathbf{e}, \mathbf{b}) - \ln P(\mathbf{k}, \mathbf{e}, \mathbf{b})$ is the description length. Equivalently to the more general introduction to description length above, the first component is the amount of information required to describe the network under the SBM and the given parameters and the second component is the amount of information needed to describe the parameters themselves (in nats). Due to this definition, finding the partition \mathbf{b} that minimises the description length is equivalent to finding the partition that maximises the posterior, i.e. to finding the most likely partition of the network.

The description length can be used to identify the most likely model given the observed data, upon comparing multiple competing models. One way of interpreting a ‘model’ in this context is as a particular partition \mathbf{b}_1 under an SBM variant m_1 . We can identify for which of two partitions \mathbf{b}_1 and \mathbf{b}_2 (under models m_1 and m_2 respectively) there is more evidence in the data, by calculating their posterior odds ratio

$$\begin{aligned} \Lambda &= \frac{P(\mathbf{b}_1, m_1|\mathbf{A})}{P(\mathbf{b}_2, m_2|\mathbf{A})} \\ &= \frac{P(\mathbf{A}|\mathbf{k}, \mathbf{e}_1, \mathbf{b}_1, m_1)P(\mathbf{k}|\mathbf{e}_1, \mathbf{b}_1, m_1)P(\mathbf{e}_1|\mathbf{b}_1, m_1)P(\mathbf{b}_1|m_1)P(m_1)}{P(\mathbf{A}|\mathbf{k}, \mathbf{e}_2, \mathbf{b}_2, m_2)P(\mathbf{k}|\mathbf{e}_2, \mathbf{b}_2, m_2)P(\mathbf{e}_2|\mathbf{b}_2, m_2)P(\mathbf{b}_2|m_2)P(m_2)} \quad (4) \\ &= e^{-\Delta\Sigma}, \end{aligned}$$

where $\Delta\Sigma = \Sigma_1 - \Sigma_2$ and $\Sigma_i = -\ln P(\mathbf{A}|\mathbf{k}, \mathbf{e}_i, \mathbf{b}_i, m_i) - \ln P(\mathbf{k}, \mathbf{e}_i, \mathbf{b}_i, m_i)$ is the description length of model i (e.g. of the network under SBM m_i with partition \mathbf{b}_i). Here we assume that both variants are equally likely i.e., $P(m_1) = P(m_2)$ ⁵. The first component of Σ_i is the amount of information required to describe the network under model i and the given parameters; the second component is the amount of information needed to describe the parameters themselves. One can therefore identify the more likely model (in terms of the specific parameters) by calculating the description lengths of the network under each model and partition. For $\Lambda = 1$ or, equivalently, $\Sigma_1 = \Sigma_2$, the models are equally likely and for $\Lambda > 1$ ($\Sigma_1 < \Sigma_2$) model m_1 is more likely than model m_2 .

Description length calculation

As described above, the description length Σ of a network \mathbf{A} under an SBM with a set of parameters θ can be directly derived from its full joint distribution, since $\Sigma = -\ln P(\mathbf{A}, \theta) = -\ln P(\mathbf{A}|\theta) - \ln P(\theta)$. Here, we detail the description length calculation for each of the SBM variants discussed in the main text, by providing their joint distributions, made up of model likelihood and priors. We will see that the microcanonical framework makes it possible to derive the likelihood and priors through combinatorics. Specifically, to be as parsimonious as possible, the priors tend to be uniform distributions over the number of possible realisations of a particular parameter under the given modelling assumptions.

Note that in the formulas of the description length calculations, we use the notation introduced for the SBM parameters. However, in the

metablock measure, the description length is not only used as part of the inference of the optimal partition but it is also calculated for the metadata partition. To calculate the metadata description length, we simply replace the parameters $\mathbf{b}, \mathbf{e}, \mathbf{n}, B$ by the respective metadata quantities that can be directly induced by \mathbf{d} . For example, e' is the number of edges within and between the node sets with shared node attributes, n' is the number of nodes that share each type of node attributes, and B' is the number of unique node attributes in \mathbf{d} .

The formulations we provide here for the likelihood and priors for the non-degree-corrected (NDC) and degree-corrected (DC) SBM are based on the work in refs. 5,32, those for the assortative planted partition (PP) SBM are from ref. 38.

For NDC, the only model parameters are the edge counts e_{rs} between blocks r and s and the block assignment vector \mathbf{b} , so the model is fully described by $P(\mathbf{A}, \mathbf{e}, \mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b})$, where $P(\mathbf{A}|\mathbf{e}, \mathbf{b})$ is the model likelihood of NDC, $P(\mathbf{e}|\mathbf{b})$ is the prior on edge counts and $P(\mathbf{b})$ is the prior on the partition. For DC, we need to consider the additional prior on the degree sequence \mathbf{k} , and we thus have $P(\mathbf{A}, \mathbf{e}, \mathbf{k}, \mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b})$ ⁵, where $P(\mathbf{k}|\mathbf{e}, \mathbf{b})$ is the probability of the degree sequence. For PP, the prior on the edge counts serves as a constraint on the network to favour assortative structure. The full joint distribution can be written as $P(\mathbf{A}, \mathbf{e}, \mathbf{k}, \mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b})P(\mathbf{k}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{e}_{\text{in}}, \mathbf{e}_{\text{out}}, \mathbf{b})P(\mathbf{e}_{\text{in}}, \mathbf{e}_{\text{out}}|E, \mathbf{b})P(E)P(\mathbf{b})$, where E is the total number of edges in the network, e_{in} and e_{out} are the number of edges within and between blocks respectively, and where we use the model likelihood of DC³⁸.

We start with the elements of the joint distribution of NDC,

$$P(\mathbf{A}, \mathbf{e}, \mathbf{b}) = P(\mathbf{A}|\mathbf{e}, \mathbf{b})P(\mathbf{e}|\mathbf{b})P(\mathbf{b}) \quad (5)$$

The model likelihood of NDC is given by

$$P(\mathbf{A}|\mathbf{e}, \mathbf{b}) = \frac{\prod_{r < s} e_{rs}! \prod_r e_{rr}!}{\prod_r n_r! \prod_{i < j} A_{ij}! \prod_i A_{ii}!} \quad (6)$$

The prior for the block matrix e_{rs} is a uniform distribution over the total possible number of symmetric block matrices given B , with the constraint that the sum of all elements must equal $2E$:

$$P(\mathbf{e}|\mathbf{b}) = \left(\binom{B(B+1)/2}{E} \right)^{-1} \quad (7)$$

The prior on the partition is defined as

$$\begin{aligned} P(\mathbf{b}) &= P(\mathbf{b}|\mathbf{n})P(\mathbf{n}|B)P(B) \\ &= \frac{\sum_r n_r!}{N!} \binom{N-1}{B-1}^{-1} \frac{1}{N} \end{aligned} \quad (8)$$

Here, $P(\mathbf{B})$ and $P(\mathbf{n}|B)$ are hyperpriors on the number of blocks B and on the block sizes n_r respectively, to be as parsimonious as possible about these parameters.

In the case of DC, the model likelihood includes terms for the degree sequence \mathbf{k} , so that:

$$P(\mathbf{A}|\mathbf{e}, \mathbf{k}, \mathbf{b}) = \frac{\prod_{r < s} e_{rs}! \prod_r e_{rr}! \prod_i k_i!}{\prod_r e_r! \prod_{i < j} A_{ij}! \prod_i A_{ii}!} \quad (9)$$

Additionally, the DC case also includes a prior on the degree sequence \mathbf{k} , namely

$$P(\mathbf{k}|\mathbf{e}, \mathbf{b}) = \prod_r \frac{\prod_k \eta_k!}{n_r!} \prod_r q(e_r, n_r)^{-1} \quad (10)$$

where η_k denotes the number of degree- k nodes in group r and $q(x, y)$ is the number of ways an integer x can be partitioned into a maximum of y parts⁵.

For PP, the prior on the block matrix needs to be defined differently, to encode the constraint that is put on the structural arrangement⁴⁷. In fact, Zhang and Peixoto³⁸ proposed two different versions of this probability: one which assumes that *uniform* expected number of edges within each community and one that allows the number of expected edges to vary across communities (*non-uniform*). Here, we give the formulation for both versions, since in our analysis, we use the uniform version in the case of synthetic networks (since they are generated with equal size blocks) and the non-uniform version in the analysis of the metablox vector on real networks. The uniform version of the prior on the edge counts in PP is described by

$$P(\mathbf{e}|e_{\text{in}}, e_{\text{out}}, \mathbf{b})P(e_{\text{in}}, e_{\text{out}}|E, \mathbf{b}) \quad (11)$$

with

$$P(\mathbf{e}|e_{\text{in}}, e_{\text{out}}, \mathbf{b}) = \frac{e_{\text{in}}!e_{\text{out}}!}{B^{e_{\text{in}}} \prod_r (e_{rr}/2)! \binom{B}{2}^{e_{\text{out}}} \prod_{r<s} e_{rs}!} \quad (12)$$

This is equivalent to the product of two uniform multinomial distributions: one for the elements of the block matrix that correspond to the within-block edge counts and one for those that correspond to the between-block edge counts, given e_{in} and e_{out} . The second part is then the hyperprior on e_{in} and e_{out} :

$$P(e_{\text{in}}, e_{\text{out}}|E, \mathbf{b}) = \left(\frac{1}{E+1} \right)^{1-\delta_{B,1}} \quad (13)$$

For the non-uniform version, the prior is also made up of two probabilities

$$P(\mathbf{e}|\{e_{rr}\}, e_{\text{out}}, \mathbf{b})P(\{e_{rr}\}, e_{\text{out}}|\mathbf{b}, E) \quad (14)$$

where

$$P(\mathbf{e}|\{e_{rr}\}, e_{\text{out}}, \mathbf{b}) = \frac{e_{\text{out}}!}{\binom{B}{2}^{e_{\text{out}}} \prod_{r<s} e_{rs}!} \quad (15)$$

corresponds to a uniform multinomial distribution for the off-diagonal elements of the block matrix given e_{out} . The second component is made out of a uniform distribution over all possible values e_{in} from E , and a uniform distribution over all ways of choosing the set of diagonal block matrix values $\{e_{rr}\}$, given e_{in} :

$$P(\{e_{rr}\}, e_{\text{out}}|\mathbf{b}, E) = P(\{e_{rr}\}|e_{\text{in}}, \mathbf{b})P(e_{\text{in}}|E, \mathbf{b}) \\ = \binom{B+e_{\text{in}}-1}{e_{\text{in}}}^{-1} \left(\frac{1}{E+1} \right)^{1-\delta_{B,1}} \quad (16)$$

The role of network compressibility

In the main text, we discussed the limitations of our measure and the way in which they can be mitigated by using the compression of the network under the *optimal* partition per network edge, $c^m = \sum_{\text{opt}}^m / E$, as a second dimension.

Recall that we considered two limitations: one related to identifying the best SBM variant as part of a comparison of different dimensions of γ and another regarding the conflation of block structure signals and metadata correlation. In Figs. 4, 5, we used edge compression to highlight the particular SBM variant whose optimal partition provided the best partition of the network and therefore illustrated how this second partition can be used to address the first limitation. Including this additional information enabled a distinction between the cases in which the metadata is closest—in terms of how well it compresses the network—to the optimal partition among the variants included in the analysis, or merely to the particular variant being

measured. In other words, when the strongest relevance is measured under DC and DC is also the model whose optimal partition provides the best fit overall, we know that the metadata partition is strongly relevant not just under the particular model but in comparison to the most optimal partition we have for this network.

Here, we provide a concrete example for the ability of the second dimension to disentangle the signal of the block structure and the metadata correlation. For the purpose of this demonstration, we generate a total of 300 synthetic networks with community structure, for each of which we fix size $N = 200$, expected degree $k = 10$ and number of blocks B . We create the networks in this collection such that we end up with three sets of 100 networks with differently ‘strong’ community structure. The stronger the signal of the community structure, the more pronounced the difference between the within- and between-block edge probability. Specifically, we generate the networks using an SBM with block matrix $\theta = 2E \begin{pmatrix} 1-\mu & \mu \\ \mu & 1-\mu \end{pmatrix}$, with $\mu = 0.1$ for the strongest signal, $\mu = 0.2$ for a medium signal and $\mu = 0.3$ for weak community structure. For each network, we generate three sets of categorical node metadata, correlating with the planted block structure with correlation $\rho = 0.7; 0.8; 0.9$. In Fig. 6, we visualise an analysis of the NDC element of metablox for these networks; we refrain from including it for other dimensions as those yielded very similar results. In the panel on the left-hand side of Fig. 6a, we see that increasing metadata correlation leads to lower values of γ . For Fig. 6b, we disaggregated the data by signal strength. We observe that for each correlation value ρ , the γ values for the networks with strong and medium community structure are essentially indistinguishable, while the networks with weak community structure have a larger value of γ for the same ρ . While γ on the whole does a good job at differentiating between the correlation strength of the different network-metadata pairs, the differences we observe upon separating by community strength call for further analysis. Especially in light of working with real networks, whose generative process is unknown to us, these insights demand a tool to disentangle the different impacting factors on the metadata relevance. In Fig. 6c–e, we demonstrate the way in which the edge compression mitigates this limitation: As expected, the stronger the community structure, the ‘better’ (i.e. smaller c) the compression of the optimal partition per network edge.

Robustness analysis

To understand the suitability of metablox to be used as part of comparative studies, we analyse the sensitivity of the measure to networks of different sizes and topologies. We first discuss the parameters to which we do and do not expect the measure to be robust. For an inter-network comparison, it is important that our measure is robust with respect to varying network size N , which is not a trivial question since the description length of a network increases with growing N . For structural differences between two networks that are related to the compressibility under an SBM, such as the number of blocks, the edge density, or the signal strength of the block structure, the question becomes more nuanced. In fact, we have already discussed in ‘Limitations and network compressibility’ that—due to our definition of metadata-block structure relevance—we *expect* the measure to be impacted by the compressibility of a network: metadata is more relevant to a network with stronger block structure signal. We therefore expect metablox to yield stronger relevance for network-metadata pairs for which the SBM variant provides better compression e.g. for increasing numbers of blocks.

To test these assumptions, we generate a number of synthetic networks. Specifically, we increase N from $N = 100$ to $N = 1000$ at steps of 100, and generate 50 networks for each N , all with expected degree $k = 10$ and with planted bicomunity structure ($B = 2$). For the within- and between block edge counts, we use the same block matrix θ as above, with $\mu = 0.1$. In Fig. 7, we plot the mean description lengths and metablox (plus 95% confidence intervals) for each value of N and three SBM variants, for correlation values of $\rho = 0.7, 0.8, 0.9$. We clearly see the intended normalising effect, as γ remains stable for growing N for all three variants.

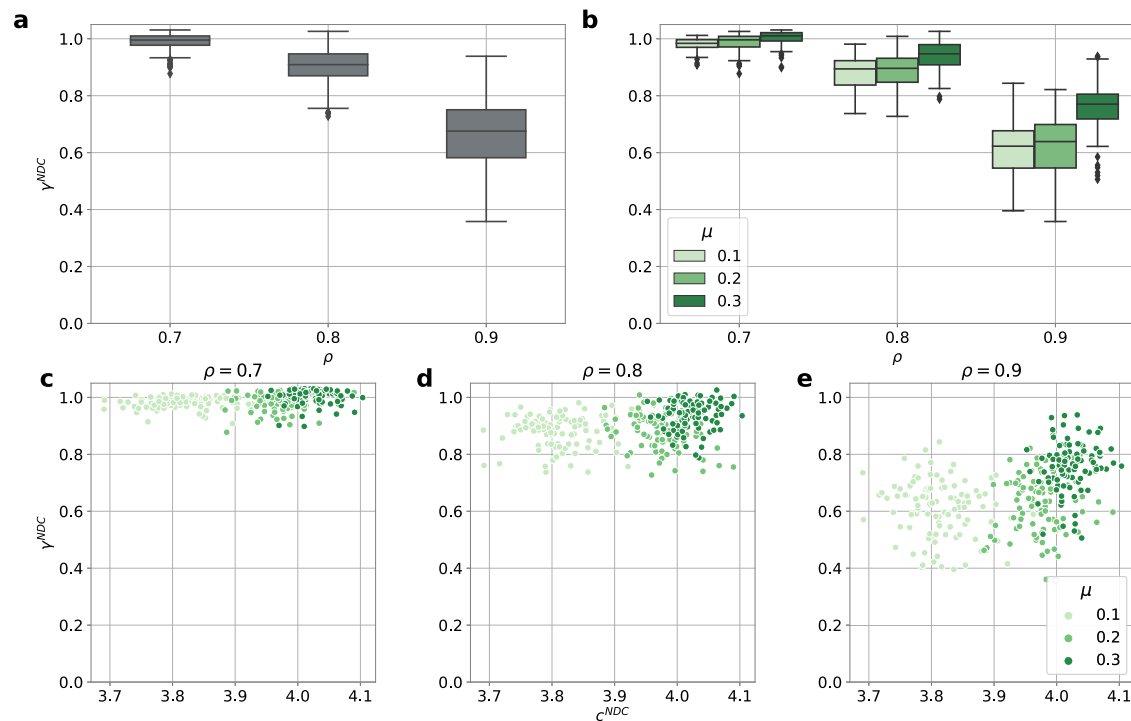


Fig. 6 | Using edge compression to disentangle block structure signal strength and metadata correlation. Distributions of non-degree-corrected metablox values for a total of 300 networks, each with three sets of metadata. For box plots, the centre line refers to the median, box limits to upper and lower quartiles and whiskers to 1.5 times the interquartile range; remaining points are outliers. **a** Box plots of the

metablox values for three values of correlation ρ . **b** Same box plots, disaggregated by block structure signal strength μ . Metablox values for the same data, as a function of edge compression c under the non-degree-corrected variant, for correlation $\rho = 0.7$ (**c**), $\rho = 0.8$ (**d**), $\rho = 0.9$ (**e**).




Fig. 7 | Metablox values for networks of varying size. Metablox values for networks of varying size N , for three different correlation values ρ and three variants (degree-corrected, non-degree-corrected, and planted partition). The lines show the mean metablox values across 50 networks, the shaded areas show 95% confidence intervals.

To test the measure's behaviour for varying numbers of blocks, we fix the network size at $N = 400$ and leave all other parameters as above, increasing B from 2 to 10, again generating 50 networks at each step. In Fig. 8a we show the results for γ_a^{NDC} (the results for the other two variants are similar). We observe that, for each value of ρ , γ decreases with the number of blocks, confirming our expectation of the measure capturing stronger metadata relevance for more compressible networks. In Fig. 8b–d, we show how the edge compression of the network can be used as a second dimension, if a distinction between metadata correlation and block structure signal strength is required. The three figures show the edge compression c on the x-axis and γ on the y-axis. A lower value of c implies better compression, and we can see that the second dimension separates the compressibility from the metadata correlation by placing the most compressible

networks ($B=10$) low on the x-axis and low on the y-axis. In a network comparison, in which disentangling these two factors is important, researchers can plot the two dimensions in this way to draw the correct conclusions.

Heterogeneous partition landscape

For the synthetic network used to illustrate the measure in the Results section, we demonstrate here that this network does, in fact, exhibit a heterogeneous partition landscape—as designed by the SCBM²⁷. Recall that we created a network with $N = 100$ nodes, $B = 2$ blocks and expected degree $k = 10$, generated in a way such that upon fitting an SBM to the network and sampling from the posterior distribution, one recovers at least two ‘clusters’ of partitions—a bicomunity partition (BC) and a second

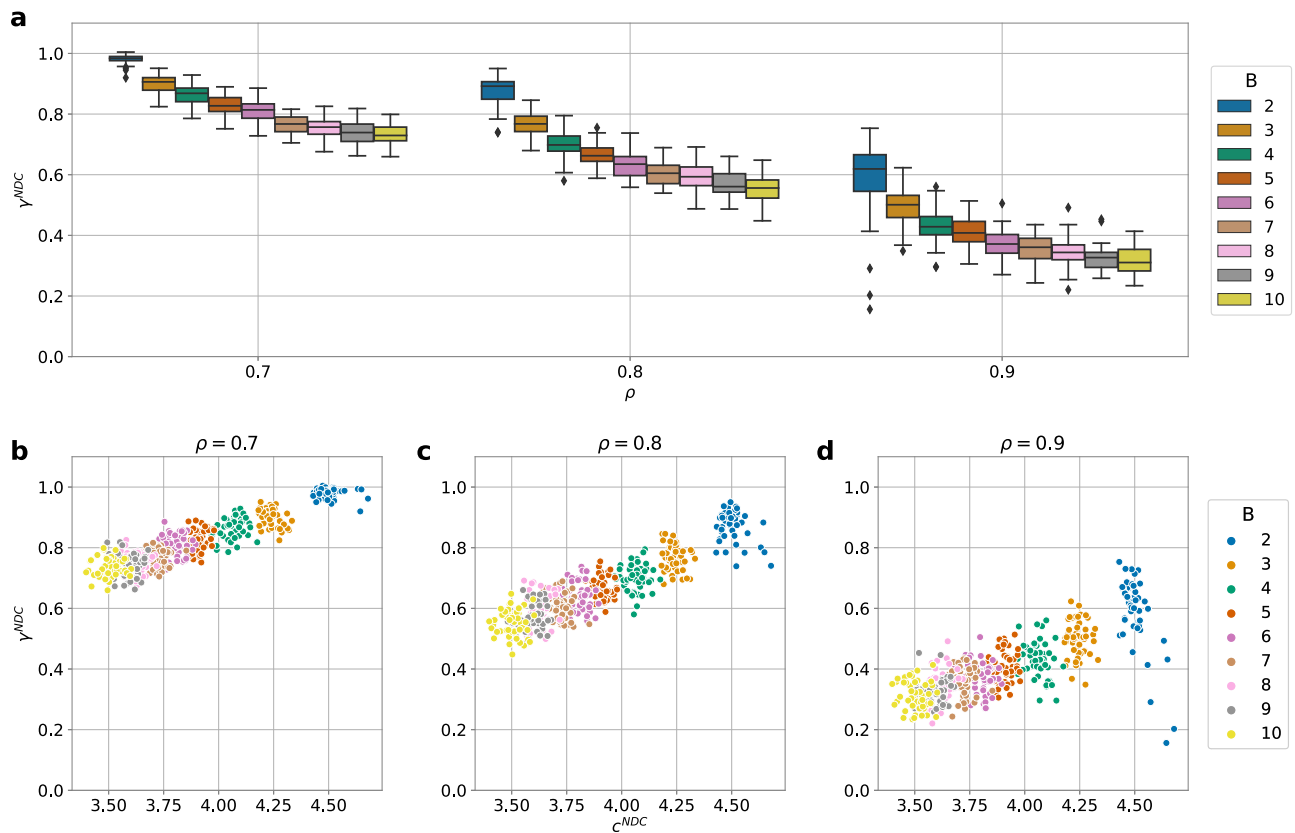


Fig. 8 | Metablob values for varying number of blocks. Distributions of non-degree-corrected metablob values for networks of size $N = 400$ with varying number of blocks B . 50 networks were created for each B and metadata was created with varying levels for correlation ρ with the planted block structure. For box plots, the centre line refers to the median, box limits to upper and lower quartiles and whiskers

to 1.5 times the interquartile range; remaining points are outliers. **a** Box plots of the metablob values for three values of correlation ρ , disaggregated by number of blocks B . The same metablob values as a function of edge compression c under the non-degree-corrected variant, for correlation $\rho = 0.7$ (**b**), $\rho = 0.8$ (**c**), $\rho = 0.9$ (**d**).

one dividing the network into a core and a periphery (CP). We choose $\theta_{BC} = 2E \begin{pmatrix} 1-\mu & \mu \\ \mu & 1-\mu \end{pmatrix}$ and $\theta_{CP} = 2E \begin{pmatrix} 1-\lambda & \frac{1}{2} \\ \frac{1}{2} & \lambda \end{pmatrix}$ as block matrices, with $\mu = 0.25$ and $\lambda = 0.05$ and where E denotes the total number of edges. The choice of μ and λ is motivated by the findings in ref. 27, where it was demonstrated that the partition landscape inferred by a degree-corrected SBM on a network with two partitions planted with these parameters does, in fact, uncover both the planted partitions. The probability of the existence of edges in this network depends solely on the block membership of nodes and on the edge probability given by these two matrices.

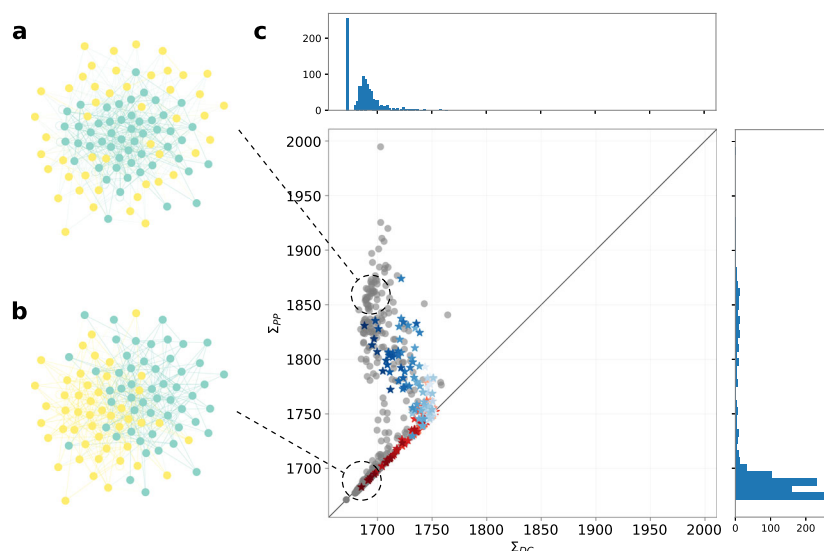
For the purpose of demonstrating that this network's partition landscape does have multiple competing explanations, we fit an SBM and sample from the posterior distribution by using the methods from the graph-tool library³⁷. In particular, we use the degree-corrected SBM variant (DC), which is meant to account for heterogeneous degree distributions within blocks⁴⁸. The description lengths Σ^{DC} of the network according to these partitions are plotted in grey dots on the x-axis in Fig. 9c. We also calculate the description lengths Σ^{PP} of the network according to each partition under the planted partition SBM (PP), the variant of the model that assumes assortativity³⁸; these are shown on the y-axis of the same figure. Note that we have $\Sigma^{PP} \geq \Sigma^{DC}$ for all partitions, since partitions were found by DC and corresponding description lengths were then calculated for the same partitions under PP. If we take a closer look at the partitions that we sampled, we find groups of similar partitions, here marked by the dashed-line circles, that look like the two partitions we planted. We show representative partitions for each partition cluster in the two network visualisations in Fig. 9a and b, in which nodes are painted according to block assignments. We clearly see a strong similarity between these inferred partitions and each of the two

planted partitions in Fig. 1 and we therefore conclude that we have successfully created a network with the desired diverse partition landscape. We also observe that for those inferred partitions that are similar to the planted bicomponent partition, the PP offers a similarly good encoding of the network as the DC, since $\Sigma^{PP} \approx \Sigma^{DC}$ for those partitions. In contrast, for the inferred partitions that are similar to the planted core-periphery partition, we have $\Sigma^{PP} > \Sigma^{DC}$. This illustrates that calculating the description length of a network under different models can help us probe its partition landscape.

In line with the objective of our measure, we use this example network to support our hypothesis that using description length is a suitable tool to understand the way in which metadata is relevant for different parts of the partition landscape. For this purpose, we calculate Σ_d^{DC} and Σ_d^{PP} for each of the 202 sets of metadata that we generated for this network, as introduced in the “Results” section and plot the resulting description length values alongside the grey dots (i.e. alongside the description length values of the *inferred* partitions) on Fig. 9c. The red stars represent the metadata partitions that are correlated with the planted bicomponent structure, the blue stars represent those that are similar to the core-periphery structure, with darker colours depicting higher values of ρ . As expected, we observe that under the more general of the two models, the partitions with the highest values of ρ have the lowest description length: the stronger the correlation of the metadata with the planted structure, the lower the description length, since the metadata partition is similar to the two ground truth partitions that were responsible for generating the network. The description lengths under the PP on the y-axis illustrate that additional to measuring the extent to which metadata is related to structure, we can also probe the type of structural arrangement: the bicomponent-like metadata are encoded as well under PP as under DC, indicating that assortativity was a prominent feature

Fig. 9 | Heterogeneity of the partition landscape.

a Visualisation of the synthetic network described in the main text, with nodes painted according to a partition that resembles the planted core-periphery structure. **b** The same network, with nodes painted according to a partition that resembles the planted bicomunity structure. **c** Description lengths under the degree-corrected and planted partition Stochastic block model of the inferred partitions (grey dots), sampled by fitting a degree-corrected Stochastic block model, and description lengths of the same network under a degree-corrected and planted partition variant, with the partition parameter set equal to different sets of metadata partitions. Red (blue) stars represent description lengths under the bicomunity-like (core-periphery-like) metadata. The darker the colour of the stars representing the metadata partitions, the higher the correlation of the metadata labels with the respective planted structure. The insets show histograms of the description lengths of the inferred partitions.



of the network generation process³⁸. The core-periphery-like metadata, however, yield much higher description lengths under PP compared to DC, suggesting that—as we know is true—if we were to assume that this metadata was responsible for generating this network, assortativity was much less likely the prominent structure compared to some other more general structure.

Data availability

The data on law firm networks⁴⁰ is available in “The Netzschleuder network catalogue and repository” (<https://doi.org/10.5281/zenodo.7839981>)⁴⁹. The impact investing data is available in the NAKALA repository (<https://doi.org/10.34847/nkl.dbd8q853>)⁴². The scripts to generate all other data are available in <https://github.com/lenafm/quantifying-metadata-relevance>.

Code availability

A Python library to use the *metablox* measure can be downloaded and installed from <https://github.com/lenafm/metablox>. The scripts to generate the results in the paper are available in <https://github.com/lenafm/quantifying-metadata-relevance>.

Received: 19 March 2024; Accepted: 30 September 2024;

Published online: 11 October 2024

References

- Lorrain, F. & White, H. C. Structural equivalence of individuals in social networks. *J. Math. Sociol.* **1**, 49–80 (1971).
- White, H. C., Boorman, S. A. & Breiger, R. L. Social structure from multiple networks. i. blockmodels of roles and positions. *Am. J. Sociol.* **81**, 730–780 (1976).
- Holland, P. W., Laskey, K. B. & Leinhardt, S. Stochastic blockmodels: First steps. *Soc. Netw.* **5**, 109–137 (1983).
- Karrer, B. & Newman, M. E. J. Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83**, 016107 (2011).
- Peixoto, T. P. Nonparametric Bayesian inference of the microcanonical stochastic block model. *Phys. Rev. E* **95**, 012317 (2017).
- Peixoto, T. P. Efficient monte carlo and greedy heuristic for the inference of stochastic block models. *Phys. Rev. E* **89**, 012804 (2014).
- Peixoto, T. P. Merge-split markov chain monte carlo for community detection. *Phys. Rev. E* **102**, 012305 (2020).
- Newman, M. E. & Clauset, A. Structure and inference in annotated networks. *Nat. Commun.* **7**, 1–11 (2016).
- Hric, D., Peixoto, T. P. & Fortunato, S. Network Structure, Metadata, and the Prediction of Missing Nodes and Annotations. *Phys. Rev. X* **6**, 031038 (2016).
- Peel, L., Larremore, D. B. & Clauset, A. The ground truth about metadata and community detection in networks. *Sci. Adv.* **3**, e1602548 (2017).
- Adamic, L. A. & Glance, N. The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, 36–43 (2005).
- Conover, M. et al. Political polarization on twitter. In *Proceedings of the International Aaai Conference on Web and Social Media*, **5**, 89–96 (2011).
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A. & Bonneau, R. Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Sci.* **26**, 1531–1542 (2015).
- Small, H. & Griffith, B. C. The structure of scientific literatures i: Identifying and graphing specialties. *Sci. Stud.* **4**, 17–40 (1974).
- Zachary, W. W. An Information Flow Model for Conflict and Fission in Small Groups. *J. Anthropological Res.* **33**, 452–473 (1977).
- Traud, A. L., Mucha, P. J. & Porter, M. A. Social structure of facebook networks. *Phys. A: Stat. Mech. its Appl.* **391**, 4165–4180 (2012).
- Yang, J. & Leskovec, J. Defining and evaluating network communities based on ground-truth. In *Proceedings of the ACM SIGKDD Workshop on Mining Data Semantics*, 1–8 (2012).
- Chakrabort, T., Sikdar, S., Tammana, V., Ganguly, N. & Mukherjee, A. Computer science fields as ground-truth communities: Their impact, rise and fall. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 426–433 (2013).
- Yang, J., McAuley, J. & Leskovec, J. Community Detection in Networks with Node Attributes. In *2013 IEEE 13th International Conference on Data Mining*, 1151–1156 (2013).
- Bothorel, C., Cruz, J. D., Magnani, M. & Micenkova, B. Clustering attributed graphs: Models, measures and methods. *Netw. Sci.* **3**, 408–444 (2015).
- Binkiewicz, N., Vogelstein, J. T. & Rohe, K. Covariate-assisted spectral clustering. *Biometrika* **104**, 361–377 (2017).
- Contisciani, M., Power, E. A. & De Bacco, C. Community detection with node attributes in multilayer networks. *Sci. Rep.* **10**, 15736 (2020).
- Fajardo-Fontiveros, O., Guimerà, R. & Sales-Pardo, M. Node metadata can produce predictability crossovers in network inference problems. *Phys. Rev. X* **12**, 011010 (2022).

24. Hric, D., Darst, R. K. & Fortunato, S. Community detection in networks: Structural communities versus ground truth. *Phys. Rev. E* **90**, 062805 (2014).
25. Sweet, T. M. & Zheng, Q. Estimating the effects of network covariates on subgroup insularity with a hierarchical mixed membership stochastic blockmodel. *Soc. Netw.* **52**, 100–114 (2018).
26. Peixoto, T. P. Revealing Consensus and Dissensus between Network Partitions. *Phys. Rev. X* **11**, 021003 (2021).
27. Mangold, L. & Roth, C. Generative models for two-ground-truth partitions in networks. *Phys. Rev. E* **108**, 054308 (2023).
28. Stanley, N., Niethammer, M. & Mucha, P. J. Testing alignment of node attributes with network structure through label propagation. <https://doi.org/10.48550/arXiv.1805.07375> (2018).
29. Grünwald, P. D. *The minimum description length principle* (MIT press, 2007).
30. Rand, W. M. Objective criteria for the evaluation of clustering methods. *J. Am. Stat. Assoc.* **66**, 846–850 (1971).
31. Meilä, M. Comparing clusterings by the variation of information. In *Learning Theory and Kernel Machines: 16th Annual Conference on Learning Theory and 7th Kernel Workshop, COLT/Kernel 2003, Washington, DC, USA, August 24–27, 2003. Proceedings*, 173–187 (Springer, 2003).
32. Peixoto, T. P. Entropy of stochastic blockmodel ensembles. *Phys. Rev. E* **85**, 056122 (2012).
33. Bianconi, G. Entropy of network ensembles. *Phys. Rev. E* **79**, 036114 (2009).
34. Akaike, H. A new look at the statistical model identification. *IEEE Trans. Autom. control* **19**, 716–723 (1974).
35. Schwarz, G. Estimating the dimension of a model. *The annals of statistics* 461–464 (1978).
36. Yan, X. et al. Model selection for degree-corrected block models. *J. Stat. Mech.: Theory Exp.* **2014**, P05007 (2014).
37. Peixoto, T. P. The graph-tool python library. *figshare* http://figshare.com/articles/graph_tool/1164194 (2014).
38. Zhang, L. & Peixoto, T. P. Statistical inference of assortative community structures. *Phys. Rev. Res.* **2**, 043271 (2020).
39. Condon, A. & Karp, R. M. Algorithms for graph partitioning on the planted partition model. *Random Struct. Algorithms* **18**, 116–140 (2001).
40. Lazega, E. *The collegial phenomenon: The social mechanisms of cooperation among peers in a corporate law partnership* (Oxford University Press, USA, 2001).
41. Chiapello, E. & Knoll, L. Social finance and impact investing: governing welfare in the era of financialization. *Historical Soc. Res./Historische Sozialforschung* **45**, 7–30 (2020).
42. Roth, C. ERC socsemics - dataset description impact investing - global dataset - retweet graphs 2008–22. [Dataset] NAKALA (2022).
43. Garimella, K., De Francisci Morales, G., Gionis, A. & Mathioudakis, M. Political discourse on social media: Echo chambers, gatekeepers, and the price of bipartisanship. In *Proceedings of the 2018 world wide web conference*, 913–922 (2018).
44. Hohmann, M., Devriendt, K. & Coscia, M. Quantifying ideological polarization on a network using generalized Euclidean distance. *Sci. Adv.* **9**, eabq2044 (2023).
45. Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W. & Starnini, M. The echo chamber effect on social media. *Proc. Natl Acad. Sci.* **118**, e2023301118 (2021).
46. Peixoto, T. P. Parsimonious module inference in large networks. *Phys. Rev. Lett.* **110**, 148701 (2013).
47. Young, J.-G., St-Onge, G., Desrosiers, P. & Dubé, L. J. Universality of the stochastic block model. *Phys. Rev. E* **98**, 032309 (2018).
48. Karrer, B., Levina, E. & Newman, M. E. Robustness of community structure in networks. *Phys. Rev. E* **77**, 046119 (2008).
49. Peixoto, T. P. The netzscheider network catalogue and repository <https://networks.skewed.de/> (2020).

Acknowledgements

We are thankful to Manuel Tonneau for his contribution to the Impact Investing data collection endeavour, as well as Ève Chiapello for the initial impulse in investigating this field on Twitter. This work was partially supported by the “Socsemics” Consolidator grant from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 772743).

Author contributions

L.M. and C.R. designed the analysis. L.M. wrote the code, performed the analysis, and prepared the figures. L.M. and C.R. wrote and approved the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Lena Mangold.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024