BISoN: A Bayesian Framework for Inference of Social Networks

Jordan D. A. Hart¹, Michael N. Weiss^{1,2}, Daniel W. Franks^{*3}, and Lauren J. N. Brent^{*1}

¹Centre for Research in Behaviour, University of Exeter, UK

²Center for Whale Research, Friday Harbour, WA, USA

³Departments of Biology and Computer Science, University of York, UK

Abstract

Social networks are often constructed from point estimates of edge weights. In many contexts, edge weights are inferred from observational data, and the uncertainty around point estimates can be affected by various factors. Though this has been acknowledged in previous work, methods that explicitly quantify uncertainty in edge weights have not yet been widely adopted, and remain undeveloped for common types of data. Furthermore, existing methods are unable to cope with some of the complexities often found in observational data, and do not propagate uncertainty in edge weights to subsequent analyses. We introduce a unified Bayesian framework for modelling social networks based on observational data. This framework, which we call *BISoN*, can accommodate many common types of observational social data, can capture confounds and model effects at the level of observations, and is fully compatible with popular methods of social network analysis. We show how the framework can be applied to common types of data and how various types of downstream analyses can be performed, including non-random association tests and regressions on network properties. Our framework opens up the opportunity to test new types of hypotheses, make full use of observational datasets, and increase the reliability of scientific inferences. We have made example R code available to enable adoption of the framework.

* Joint senior authors.

Introduction

Social network analysis is arguably one of the most popular frameworks in the study of sociality (Croft et al., 2016). In many scientific contexts, network connections (edges) are inferred from observational data, from human observers recording social interactions between macaques, to biologgers capturing proximity events between people. The type of data used to infer networks and the manner in which the data are collected affects both the interpretation of social network analyses and their accuracy (Whitehead, 2008).

Networks are usually constructed by taking a normalised measure of sociality, such as the proportion of sampling periods each pair spends engaged in a social behaviour (e.g. the Simple Ratio Index; SRI). These normalised measures ensure that pairs that are observed for longer are not erroneously determined to be more social (Whitehead, 2008). This is important because observing social events can be challenging and uniform sampling over all pairs is not always possible (Croft et al., 2010). Though normalised measures of sociality will not be biased by sampling time, they will be accompanied by varying levels of certainty. For example, edge weights will be treated as a certain 0.5 for both a case where individuals have been seen together once and apart once, and equally where individuals have been together 100 times and apart 100 times. Despite there being considerably more certainty around the values of the latter example than the former, methods to estimate and propagate uncertainty through social network analyses remain largely underdeveloped (but see Farine and Strandburg-Peshkin (2015) for a Bayesian approach to estimating uncertainty over edge weights).

In this paper we introduce BISO_N: a general unified Bayesian framework for modelling social network data. BISO_N captures uncertainty in edge weights, making full use of available data; propagates uncertainty through downstream analyses; and can control for social and non-social effects at any level (individual, dyad, group, observation, etc). Any type of social network analysis can be conducted within our framework, including dyadic and nodal regressions, non-random edge weight tests, and estimation of structural network properties. We will focus on how our framework can be applied to animal systems, but the underlying principle can be applied to any type of network analysis where network edges are inferred. The BISO_N framework presents an opportunity to generate reliable, flexible, and rich scientific inference in the study of social systems.

BISO_N - Bayesian Inference of Social Networks

In the following section we outline the three base models we have developed for binary, count, and duration data (see Table 1 for definitions of these types of data). The models are presented without priors, as the choice of prior should depend on the context and the question in hand (van de Schoot et al., 2021). These base models can be used to

Table 1: Definitions of the three types of social observational data we have developed models for. Examples of standard edge weight interpretations are included in the third column.

Data type	Description	Edge interpretation
Binary	The presence or absence of a social event for each dyad is recorded in each sampling period.	Probability of engaging in a social event in a fixed period of time.
Count	The number of social events for each dyad is recorded in each sampling period.	Rate of occurrence of social events per unit time.
Duration	The amount of time each dyad spends engaged in a social event is recorded in each sampling period.	Proportion of time spent engaging in a social event.

generate estimates for network edge weights ω_{ij} . Once estimates for edge weights ω_{ij} are obtained, they can be used in downstream analyses. We will briefly outline three types of analyses that propagate uncertainty through the analysis. These analyses are: 1) non-random edge weights, 2) dyadic regression, and 3) nodal regression. When presenting the downstream analyses, for brevity we will refer to them in the context of the count data model, but these downstream analyses are freely applicable to all types of data. The models we present here are simple examples that can be extended and improved in many ways. The Bayesian nature of this framework means that these methods will be useful and reliable, providing they are used appropriately in the context of the data, the scientific question, and with appropriate diagnostic tools (Kruschke, 2015; McElreath, 2020). We have provided example code at: www.github.com/JHart96/BISON_examples.

Quantifying uncertainty in edge weights

Binary data model

To demonstrate the notion behind BISO_N, we will first describe the base model for the case of binary data, where presence or absence of a social event per fixed sampling period is recorded. First, note that for binary data, edge weights are commonly defined as:

$$W_{ij} = \frac{X_{ij}}{D_{ij}}$$

where X_{ij} is the number of social events that occurred between i and j , and D_{ij} is the number of sampling periods where a social event could have been observed between i and j .

The edge weight is therefore equivalent to the probability of i and j engaging in a social event in any given sampling period. This can also be interpreted as an estimate of the proportion of time i and j spend engaging in a social event. Let us refer to this probability/proportion as p_{ij} . Without making any additional assumptions, an equivalent formulation is:

$$X_{ij} \sim \text{Binomial}(D_{ij}, p_{ij})$$

Note that the maximum likelihood estimate for p_{ij} in this process is $X_{ij}/D_{ij} = W_{ij}$. The binomial process is the natural process for multiple Bernoulli trials, which makes this formulation a natural extension of point estimates for edge weights. In a Bayesian context, the parameter p_{ij} is treated as a random variable, and inherently captures any uncertainty around its value.

Social data can sometimes be influenced by non-social effects at the observation level, but modelling the aggregated social events X_{ij} at the dyad level does not allow such effects to be modelled. We propose to model these effects by decomposing the binomial process over the aggregated observations to distinct Bernoulli processes over unaggregated observations, as follows:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Bernoulli}(p_{ij}^{(n)}) \\ \text{logit}(p_{ij}^{(n)}) &= \omega_{ij} + \dots \end{aligned}$$

where $p_{ij}^{(n)}$ is the probability of a social event occurring between i and j in the n -th observation of i and j , and ω_{ij} is the edge weight between i and j . The ellipsis represents additional terms that can be included (if desired) to model various effects. For example, if individuals move through different locations with different visibilities, it may be that some events are missed in some locations more than others, so in some cases including an observation-level location effect will account for this. We use a linear combination of effects in these examples for simplicity, but the effects could equally be modelled as non-linear functions, depending on the question. Alternatively dyad- or individual-level effects can also be included to separate out different types of social effect.

Modelling non-social effects as predictors in a regression-like model is conceptually similar to that of generalised affiliation indices (GAIs), introduced by Whitehead and James (2015). GAIs use the residuals of regressions of standard affiliation indices on non-social effects to remove the effects of non-social drivers on observations. Our approach differs in that it directly captures non-social drivers in the core model, not in residuals, and ensures that uncertainty is preserved throughout.

The binary data model can also be used for group-based (gambit) data (Franks et al., 2010). If being used for group data, an additional term in the model can be used to account for non-independence within group observations. See the supplementary material for further discussions on these points.

Count data model

The same concept as above can also be applied to the case of count data, where counts of social events between individuals per sampling period is recorded. First, note that for count data, edge weights are often defined as:

$$W_{ij} = \frac{X_{ij}}{D_{ij}}$$

where X_{ij} is the total number of social events that occurred between i and j , and D_{ij} is the total amount of time a social event could have been observed between i and j .

The edge weight is therefore equivalent to the rate at which i and j engaged in a social event per unit time. Let us refer to this rate as λ_{ij} . When events occur with a fixed rate, the number of events expected per unit time is described by the Poisson process (Wasserman, 2004). Since we assume that the rate of events characterises edge weights well, then without making any additional assumptions, we can model social events as:

$$X_{ij} \sim \text{Poisson}(\lambda_{ij}D_{ij})$$

where λ_{ij} is the underlying rate parameter. When modelled in a Bayesian framework, this provides a natural way to model the uncertainty around edge weights because the maximum likelihood estimate for λ_{ij} is $X_{ij}/D_{ij} = W_{ij}$ (Held and Sabanés Bové, 2014).

Again, this model doesn't allow modelling of observation-level effects, so we must decompose the aggregated data into the original sampling periods. Fortunately, the sum of Poisson-distributed random variables is also described by a Poisson distribution. This means that the full model can take a similar form to the previous one:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)}D_{ij}^{(n)}) \\ \log \lambda_{ij}^{(n)} &= \omega_{ij} + \dots \end{aligned}$$

where $\lambda_{ij}^{(n)}$ is the rate at which social events occur between i and j in the n -th observation of i and j , $D_{ij}^{(n)}$ is the length of time of the n -th sampling period for the dyad, and ω_{ij} is again the edge weight between i and j . As above, the ellipsis represents potential additional effects that could be included in the model if necessary. This is discussed further in a later section and in the supplementary material.

Duration data

The duration data model is slightly more complex than the binary and count data models, because it models two different types of data simultaneously: the duration of social

events, and the frequency with which they occur. With duration data, the edge weight is calculated as:

$$W_{ij} = \frac{X_{ij}}{D_{ij}}$$

where X_{ij} is the amount of time i and j were seen engaged in a social event, and D_{ij} is the maximum amount of time i and j could have been seen engaged in a social event. The total amount of time X_{ij} is the sum of times from K_{ij} social events. The edge weight is then modelled as the proportion of time t_{ij} each dyad could engage in a social event that is actually spent engaging in the social event:

$$t_{ij} = \mu_{ij} \lambda_{ij}$$

where μ_{ij} is the mean event duration for the dyad and λ_{ij} is the mean rate of events per unit time for the dyad. Inter-arrival times in a Poisson process are distributed according to the exponential distribution, so we use this to model the duration of social events once they have started. The count of events can again be considered to be a Poisson process. Bringing these two assumptions together gives the following model (see supplementary material for the full derivation):

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Exponential}(\lambda_{ij}/t_{ij}^{(n)}) \\ K_{ij} &\sim \text{Poisson}(\lambda_{ij} D_{ij}) \\ \text{logit}(t_{ij}^{(n)}) &= \omega_{ij} + \dots \end{aligned}$$

The mean event time μ_{ij} is modelled implicitly in $t_{ij}^{(n)}$ and can be recovered from the fitted model afterwards using $\mu_{ij} = \omega_{ij} / \lambda_{ij}$.

Accounting for social and non-social effects

By modelling social events at the observation-level rather than the dyad-level, it is possible to account for observation-level effects that may influence the estimated edge weights. An example of this is where the study population might move through various locations with different visibilities, leading to location-dependent missed social events. Assuming that the location doesn't actually influence sociality, and that differences in estimated sociality are just due to differences in visibility, we can include a location effect in the model. An example of how this can be achieved in the count model is:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)} d_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + L_{ij}^{(n)} \\ L_{ij}^{(n)} &\sim \text{Normal}(0, \sigma_L^2) \end{aligned}$$

where $L_{ij}^{(n)}$ is a location effect corresponding to the location of the social events observed for the dyad ij in the n -th observation period. The location effect has an adaptive normal prior centered at zero (often known as a random effect).

Depending on the assumed causal process that generates observations, it won't always be possible to control for non-social effects in this way, as the social effect of interest may partially be borne out through apparent non-social effects, such as space use. For example individuals may genuinely be more social in some locations than others, and controlling for location would remove the social effect of interest. A discussion on the the considerations around modelling non-social effects is included in the supplementary material.

Non-random edge weights

BISoN also enables a Bayesian test analogous to the classic non-random association test developed by Bejder et al. (1998). At its core, this test asks whether the structure of an observed network is an artefact of sampling or if there is genuine variation in edge weights between individuals. If the structure of an observed network is simply due to sampling, then there is no variation in edge weight, and each dyad is equally likely to engage in social events as all other dyads. In BISoN, this is equivalent to

$$X_{ij}^{(n)} \sim \text{Poisson}(\lambda_{ij}^{(n)} D_{ij}^{(n)})$$

$$\log(\lambda_{ij}^{(n)}) = \omega + L_{ij}^{(n)}$$

where ω is now a single edge weight parameter shared between all dyads, indicating no variation in underlying edge weights between dyads. This model constitutes the null hypothesis H_0 of random edge weights. This can be compared to the alternative hypothesis H_1 of non-random edge weights, where the ω parameter is replaced with a dyad-level parameter ω_{ij} , allowing each dyad to exhibit a different edge weights. In a Bayesian context, these two models ($H_0 : \omega_{ij} = \omega$ vs $H_1 : \omega_{ij} \neq \omega$) can be compared using the *Bayes factor*. Instead of attempting to reject the null model, as in null hypothesis significance testing, the Bayes factor quantifies the support for or against the hypotheses. The Bayes factor in favour of H_1 over H_0 is defined as:

$$\text{BF}_{10} = \frac{P(H_1|X)}{P(H_0|X)} = \frac{P(X|H_1) P(H_1)}{P(X|H_0) P(H_0)}$$

where $P(H_1|X)$ and $P(H_0|X)$ denote the probability of the alternative hypothesis and the null hypothesis given the data X respectively. $P(X|H_1)$ and $P(X|H_0)$ are known as the marginal likelihoods, and can be estimated numerically from a fitted model (Gronau et al., 2017). The Bayes factor can be directly interpreted as the relative likelihood of the competing hypotheses. For example, a Bayes factor of 5 in favour of H_1 over H_0 would indicate that H_1 is 5 times more likely than H_0 given the data and model. A useful

property of Bayes factors is that they can quantify support for or against hypotheses, and can even be used to compare multiple hypotheses, unlike frequentist p -values.

The original test from Beijder et al. (1998) is extendable to test for non-random associations within categories (such as sex or age). However, it requires continuous variables such as age to be discretised into categories. Discretising variables is generally not ideal, and removes information, weakens statistical power, and can lose useful insights (Altman and Royston, 2006). Our proposed test can account for variables such as age without discretisation. This can be achieved by adding variables to the linear predictor in the following way:

$$\log(\lambda_{ij}^{(n)}) = \omega + S_{\text{sex}[ij]} + L_{ij}^{(n)}$$

where S is a vector of length 3 describing the effects for each combination of the two sexes (*Female-Female*, *Female-Male*, *Male-Male*) on the log rate of social events, and sex_{ij} indexes the sex variable for each dyad ij .

Dyadic regression

A common type of analysis in studies of social networks is to consider factors that affect edge weights. This type of analysis is often called dyadic regression, where explanatory factors such as sex or age difference are regressed against edge weight. The edge weights in dyadic regression are non-independent, as variables such as age difference or sex difference are inherently linked to individual-level attributes. This means that effects due to age or sex in individual i affect all dyads that connect i . This non-independence can be controlled for by including node-level effects in the regression (Tranmer et al., 2014).

Using the count data example discussed earlier, we propose to conduct dyadic regression using a standard regression model of the form:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + L_{ij}^{(n)} \\ \hline \omega_{ij} &\sim \text{Normal}(\beta_0 + \beta_1 x_{ij} + u_i + u_j, \sigma^2) \end{aligned}$$

where β_0 is the intercept parameter, β_1 is the slope parameter, u_i are parameters accounting for the effect of node membership on edge weight, and σ is the standard deviation of the residuals. The dashed line indicates that the model can be fit in two separate parts: the first part being the core edge weight model, and the second part being the dyadic regression model. It is not strictly necessary to separate out the model like this, but doing so makes it possible to fit the edge weight model once and conduct multiple types of analysis on it afterwards without fitting the entire model again.

The β_0 and β_1 parameters can be interpreted in the same way as usual regression coefficients, and will be accompanied with credible intervals describing the probability distribution of coefficient values. There are several ways to gain inference from these values. The overlap of the credible interval with zero is often used to indicate the presence of a ‘significant effect’, though this approach has been criticised for several reasons (Gelman et al., 2012; Ogle et al., 2019). An alternative option is to use Bayes factors to quantify support for or against the hypothesis $H_1 : \beta_1 \neq 0$ (Kass and Raftery, 1995; Jeffreys, 1998). In regression settings, the Bayes factor for a point null hypothesis can be calculated quickly using the *Savage-Dickey* method (Wagenmakers et al., 2010).

The model shown here is intended as a minimal example of a simple linear regression with node-level effects to account for the non-independence of edges. The regression can be extended to support different family distributions, hierarchical effects, non-linearities, and more. We have included examples of diagnostics and some extensions in the example code.

Nodal regression

The final common type of network analysis we will cover here is nodal regression, where a regression is performed to analyse the relationship between a nodal network metric (such as centrality) and nodal traits (such as age and sex). These analyses are usually used to assess how network position depends on various biological factors, but can also be used where network position is a predictor. Since node metrics are derivative measures of the network, uncertainty in edge weights should ideally propagate through to node metrics, and on to coefficients in regression analyses, giving an accurate estimate of the total uncertainty in inferred parameters. The core of the nodal regression is similar to dyadic regression, taking the form:

$$X_{ij}^{(n)} \sim \text{Poisson}(\lambda_{ij}^{(n)} D_{ij}^{(n)})$$

$$\log(\lambda_{ij}^{(n)}) = \omega_{ij} + L_{ij}^{(n)}$$

$$M_i(\omega) \sim \text{Normal}(\beta_0 + \beta_1 x_i, \sigma^2)$$

where β_0 is the intercept parameter, β_1 is the slope parameter, and σ is the standard deviation of the residuals. $M_i(\omega)$ denotes the node metric estimates when applied to the edge weights estimated in the top part of the model.

Calculating node metrics within the model may present a practical challenge when using standard model fitting software, as many node metrics can not be described purely in terms of closed-form equations (Freeman, 1978). In this case, splitting up the model along the dashed line becomes an important step, as it allows the edge weights to be

estimated using a standard piece of software, and leaves the regression part of the model to be fit using a sampler that supports numeric estimates of network centralities. As part of the code we have made available, we have written a custom Metropolis-Hastings sampler that allows numeric estimates of node metrics to be used when fitting the model. Importantly, our sampler samples from the joint distribution of edge weights, rather than sampling edge weights independently. This ensures that the effect of any structure in the observation data (such as location effects) is maintained and propagated through to the node metric estimates, and subsequently to regression coefficient estimates.

Example: Synthetic dataset

To demonstrate our framework, we present an example based on a synthetic dataset of social events simulating count data. The synthetic dataset was generated by assigning underlying undirected edge weights to the 28 dyads corresponding to all pairs of 8 individuals living in a closed social group. Individuals 1 - 4 were assigned to the treatment condition and individuals 4 - 8 were assigned to the control. Dyads between two individuals with the treatment condition were assigned a higher edge weight than other dyads, to induce an effect of treatment condition on centrality. Social events were simulated from a Poisson point process based on both the underlying edge weight and an observation-level, non-social effect of location. There were 6 possible locations for observations of social events that were used to simulate the effect of varying visibilities in different locations. Between 10 and 50 observation periods of fixed length (arbitrary time units) were simulated for each dyad. Full details of the simulation can be found in the example code.

The synthetic dataset describes the count of events for each observation period, therefore the appropriate model for this data is the count model. The full model was defined as:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + L_{ij}^{(n)} \\ \omega_{ij} &\sim \text{Normal}(0, 1) \\ L_{ij}^{(n)} &\sim \text{Normal}(0, \sigma_L^2) \\ \sigma_L &\sim \text{Half-Normal}(0, 1) \end{aligned}$$

where $X_{ij}^{(n)}$ is the count of social events between i and j in their n -th observation period, $\lambda_{ij}^{(n)}$ is the social event rate from the corresponding period, ω_{ij} is the dyad-level social effect, and $L_{ij}^{(n)}$ is an observation-level social effect corresponding to one of the 6 locations. The prior distributions were chosen using a prior predictive check to ensure the parameters only take on biologically plausible values. The model was implemented in

Figure 1: A) Visualisation of the network generated from uncertain edge weights. The width of edges denotes their weights, and the bands around the edges is proportional to the width of their 95% credible intervals. Low edge weights are arbitrarily removed from the plot (though not the analysis) for visualisation purposes. B) Posterior medians and credible intervals of log-scale edge strengths ω_{ij} against the underlying edge weights intended by the simulation for each of the 28 dyads.

Stan, and fitted in R, using the package *Rstan* (R Core Team, 2020; Stan Development Team, 2020).

The model fit was checked visually by examining the chains and computing the Gelman-Rubin $\hat{R}$ convergence statistic (Gelman and Rubin, 1992). Once the model fit was verified, the parameter estimates of edge weights ω_{ij} between dyads were extracted. Figure 1A shows a visualisation of the network based on edge weight estimates ω_{ij} , where the bands around edges are proportional to the widths of the 95% credible intervals. Figure 1B shows the posterior medians and 95% credible intervals of edge weights ω_{ij} against underlying edge weights intended by the simulation. At this point the estimated network can be treated much the same as a standard network of point estimates, as long as care is taken to preserve uncertainty through subsequent analyses. To demonstrate the power of BISoN for conducting downstream analyses, we conducted a regression analysis where the response was eigenvector centrality and the predictor was a categorical variable with two categories, representing the treatment condition. To improve model fit, the response was standardised by subtracting the mean and dividing by the standard deviation. The

model specification for the regression analysis was:

$$\begin{aligned} EV_i &\sim \text{Normal}(\beta_{\text{Treat}[i]}, \sigma_2) \\ \beta_{\text{Treat}[i]} &\sim \text{Normal}(0, 1) \\ \sigma &\sim \text{Half-Normal}(0, 1) \end{aligned}$$

where EV_i denotes the standardised eigenvector centrality for the i -th node, $\beta_{\text{Treat}[i]}$ is the effect parameter corresponding to the treatment condition (control or treatment) for the i -th sample, and σ is the standard deviation of the residuals. The regression model was fit using our custom Metropolis-Hastings sampler and chain convergence was checked visually and with the $\hat{R}$ statistic. Posterior predictive checks were used to ensure the model fitted the data well.

The posterior distributions of unstandardised eigenvector centralities for the 8 nodes are shown in Figure 2A. The treatment condition parameters are interpretable as the mean standardised eigenvector centrality of nodes for each of the two conditions. The 95% credible interval for the control parameter β_{Control} was $[-1.27, -0.36]$, with a median of -0.83 , and for the treatment parameter $\beta_{\text{Treatment}}$ was $[0.35, 1.29]$, with a median of 0.85 . One of the many benefits of Bayesian inference is that posterior distributions of different parameters can themselves be used gain statistical insights. In this case, the question of real interest is the difference between the two treatment conditions: $\beta_{\text{Treatment}} - \beta_{\text{Control}}$. This quantity can simply be calculated by subtracting the posteriors corresponding to the treatment parameter from those corresponding to the control posterior. We found that the difference in standardised centrality between treatment conditions, $\beta_{\text{Treatment}} - \beta_{\text{Control}}$, had a 95% credible interval of $[0.99, 2.33]$ and a median of 1.69 . The posterior distribution of the difference is also shown in Figure 2B. This indicates that the treatment condition appears to correlate strongly with a higher centrality.

This brief example outlines just some of the types of analysis that can be carried out in our framework. We have included several additional examples in the code included with this paper.

Discussion

In this paper we presented BISO_N, a general, highly flexible framework for modelling and analysing social networks. BISO_N uses models of data collection processes to quantify uncertainty in edge weights. These models are versatile, and can be extended to capture additional social and non-social effects. We demonstrated how uncertainty in the network can then be propagated through to subsequent analyses, such as regressions involving network centrality. We have provided example code for how to implement the BISO_N framework in R with Stan at: www.github.com/JHart96/BISON_examples.

Figure 2: A) Posterior distributions of eigenvector centralities for the 8 nodes. B) Posterior distribution of difference in centrality between treatment conditions, extracted from the regression model.

BISoN acknowledges the inherent uncertainty in inferred edge weights, and by propagating this through downstream analyses can offer nuanced, balanced, and powerful statistical analyses, focused on quantifying biological effects and the uncertainty around them. As well as propagating uncertainty into downstream analyses, the core edge weight model itself can also be adapted to ask questions about individual variation in gregariousness, reciprocation of directed behaviours, and underlying differentiation of edge weights, among many others (see the supplementary material for further information).

One of the main benefits of our framework is its inherent flexibility, but this also introduces a number of considerations when using it. Setting priors is an important part of Bayesian modelling, and it is important that they reflect realistic prior expectations of parameter values (van de Schoot et al., 2021). The priors included in the example code are not intended as good defaults, and instead the priors should be determined by careful consideration of the meaning of the parameter, prior expectations of likely values for the parameter, and using prior predictive checks. It is also important to note that using additional terms and altering the linear predictor part of the model can change the interpretation of model parameters, which may require the priors to be altered as well.

Another consideration is that network centrality measures based on binarised edges, such as degree and transitivity, will not perform as expected when using estimated edge

weights from this framework. This is because, even though a particular dyad may never have been seen interacting in 100 samples, it is still possible that they would interact in the 101st sample. The edge weight models BISO_N uses naturally acknowledge this, so there is uncertainty around all edge weights. If binarised centralities are needed, there are two possible ways to use them in BISO_N: 1) the edge weights can be thresholded, to generate a binary distribution on 0 and 1 for each edge; or 2) the edge weights can be assigned to categories, and membership of a specific category (e.g. weak or strong) becomes the new binary measure, again maintaining uncertainty over the membership. Option 1 can simply be applied to the sampling posterior by choosing a biologically relevant cut-off point and applying the threshold. This approach may be challenging because it requires mapping domain knowledge and the question, in context, to a single threshold point for the analysis. Option 2 requires partial pooling over edge weights, assuming that they are drawn from a mixture of distributions, and assigns each edge weight to each type of connection with a certain probability. This approach is conceptually similar to the social bond categorisation method introduced by Ellis et al. (2021). If it is biologically sensible to believe there are different categories of connections in a system, this can provide a natural and principled way to compute binary network measures (see supplementary material for an example of how to do this).

A final consideration of our modelling framework is that, though non-social effects can be modelled in the linear predictor, this alone cannot remove some non-social effects, and could instead absorb the social effect of interest, depending on the underlying causal system. We recommend that these effects be treated with caution, and used only with explicit assumptions about the underlying causal system. These considerations are also discussed in further depth in the supplementary material using causal models.

The versatility of our proposed framework makes the approach powerful, but comes with a number of limitations. Though highly flexible, one important constraint on BISO_N is that knowledge of sampling effort will be required to properly model the processes that generate observations. This constraint may prevent these models from being applied to some published datasets, where only point estimates of network edges are made available. On a practical level, BISO_N models will usually need to be fit using an MCMC sampler. MCMC can be slow and computationally expensive, especially with large datasets. If additional effects are not included, this issue can be partially overcome by aggregating data at the dyad-level and fitting collapsed versions of the models (see supplementary materials). Furthermore, MCMC methods do not guarantee convergence, so additional checks will be required to make sure the sampler has performed correctly (Cowles and Carlin, 1996; Draper, 2008). Together these limitations make the process of model building, fitting, and checking a more complex process than simply building networks from point estimates. However, in return for these investments, researchers will have access to a powerful suite of tools not otherwise available.

The BISO_N framework is open-ended and can be extended in many ways. One particularly useful aspect of the Bayesian approach is the ability to treat data with uncertainty. This allows missing data to be modelled within the statistical model without the need for repeated model fitting routines. This allows missing data to be estimated from other variables, and could be especially useful in populations where traits such as age are often estimated or unknown (Little and Rubin, 2002). Another potential direction for the framework is to move beyond static sociality measures and instead explicitly model the dynamics of edge weights, as discussed by Pinter-Wollman et al. (2014). Dynamic edge weights could be modelled by incorporating autoregressive time series models into the framework (Wei, 2013). This would make it possible to model the general trends in edge weights, how changes in edge weights propagate through the network, and even model flow processes such as disease or information transmission.

Conclusion

Uncertainty affects all social network analyses, yet until now has been difficult to quantify. The BISO_N framework we have introduced is designed to help quantify uncertainty and ensure it is preserved through subsequent analyses to give robust and reliable statistical inference. Our framework is general and widely applicable to almost any type of social data. Estimated edge weights from our models can be used with many different types of downstream analysis while propagating uncertainty. We believe the ideas presented in this paper are only the first step in developing powerful, flexible, and robust statistical models for social networks. BISO_N has the potential to enable us to ask a wide range of nuanced questions, and obtain a deeper insight into the nature of sociality.

Author contributions

JDAH conceived of the initial idea behind BISO_N, which was further developed with the input from MNW, DWF, and LJNB. The example code was written by JDAH. The manuscript was written by JDAH with input from MNW, DWF, and LJNB.

Acknowledgements

JDAH acknowledges funding from the Engineering and Physical Sciences Research Council [grant number EP/R513210/1]. LJNB acknowledges funding from a European Research Council Consolidator Grant (FriendOrigins - 864461) and from the National Institute of Health (R01AG060931, R01MH118203). DWF and MNW acknowledge funding

from the Natural Environment Research Council [grant number NE/S010327/1]. DWF also acknowledges funding from the Natural Environment Research Council [grant number NE/S009914/1]. The authors declare no conflict of interest.

References

- Douglas G Altman and Patrick Royston. The cost of dichotomising continuous variables. *BMJ : British Medical Journal*, 332(7549):1080, May 2006. ISSN 0959-8138.
- Lars Bejder, David Fletcher, and Stefan Bräger. A method for testing association patterns of social animals. *Animal Behaviour*, 56(3):719–725, September 1998. ISSN 0003-3472. doi: 10.1006/anbe.1998.0802.
- Mary Kathryn Cowles and Bradley P Carlin. Markov chain monte carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association*, 91(434): 883–904, 1996.
- Darren P Croft, Richard James, and Jens Krause. *Exploring animal social networks*. Princeton University Press, Princeton, N.J., 2010. ISBN 978-1-4008-3776-2. OCLC: 939935939.
- Darren P Croft, Safi K Darden, and Tina W Wey. Current directions in animal social networks. *Current Opinion in Behavioral Sciences*, 12:52–58, December 2016. ISSN 2352-1546. doi: 10.1016/j.cobeha.2016.09.001.
- David Draper. *Bayesian Multilevel Analysis and MCMC*, pages 77–139. Springer New York, New York, NY, 2008. ISBN 978-0-387-73186-5. doi: 10.1007/978-0-387-73186-5_2.
- Samuel Ellis, Daniel W Franks, Michael N Weiss, Michael A Cant, Paolo Domenici, Kenneth C Balcomb, David K Ellifrit, and Darren P Croft. Mixture models as a method for comparative sociality: social networks and demographic change in resident killer whales. *Behavioral Ecology and Sociobiology*, 75(4):1–15, 2021.
- Damien R. Farine and Gerald G. Carter. Permutation tests for hypothesis testing with animal social network data: Problems and potential solutions. *Methods in Ecology and Evolution*, 2021.
- Damien R. Farine and Ariana Strandburg-Peshkin. Estimating uncertainty and reliability of social network data using Bayesian inference. *Royal Society Open Science*, 2(9): 150367, September 2015. ISSN 2054-5703. doi: 10.1098/rsos.150367.
- Daniel W. Franks, Graeme D. Ruxton, and Richard James. Sampling animal association networks with the gambit of the group. *Behavioral Ecology and Sociobiology*, 64(3):493–503, January 2010. ISSN 0340-5443, 1432-0762. doi: 10.1007/s00265-009-0865-8.
- Linton C Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1978.

- Andrew Gelman and Donald B. Rubin. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457–472, November 1992. ISSN 0883-4237, 2168-8745. doi: 10.1214/ss/1177011136. Publisher: Institute of Mathematical Statistics.
- Andrew Gelman, Jennifer Hill, and Masanao Yajima. Why We (Usually) Don’t Have to Worry About Multiple Comparisons. *Journal of Research on Educational Effectiveness*, 5(2):189–211, April 2012. ISSN 1934-5747, 1934-5739. doi: 10.1080/19345747.2011.618213.
- Quentin F Gronau, Alexandra Sarafoglou, Dora Matzke, Alexander Ly, Udo Boehm, Maarten Marsman, David S Leslie, Jonathan J Forster, Eric-Jan Wagenmakers, and Helen Steingroever. A tutorial on bridge sampling. *Journal of mathematical psychology*, 81:80–97, 2017.
- Leonhard Held and Daniel Sabanés Bové. *Applied Statistical Inference*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014. ISBN 978-3-642-37886-7 978-3-642-37887-4. doi: 10.1007/978-3-642-37887-4.
- Harold Jeffreys. *The Theory of Probability*. OUP Oxford, August 1998. ISBN 978-0-19-158967-6. Google-Books-ID: vh9Act9rtzQC.
- Robert E. Kass and Adrian E. Raftery. Bayes Factors. *Journal of the American Statistical Association*, 90(430):773–795, June 1995. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.1995.10476572.
- David A Kenny and Lawrence La Voie. The social relations model. In *Advances in experimental social psychology*, volume 18, pages 141–182. Elsevier, 1984.
- John K. Kruschke. *Doing Bayesian data analysis: a tutorial with R, JAGS, and Stan*. Academic Press, Boston, edition 2 edition, 2015. ISBN 978-0-12-405888-0.
- Roderick JA Little and Donald B Rubin. Bayes and multiple imputation. *Statistical analysis with missing data*, pages 200–220, 2002.
- Richard McElreath. *Statistical rethinking: a Bayesian course with examples in R and Stan*. CRC texts in statistical science. Taylor and Francis, CRC Press, Boca Raton, 2 edition, 2020. ISBN 978-0-367-13991-9.
- Kiona Ogle, Drew Peltier, Michael Fell, Jessica Guo, Heather Kropp, and Jarrett Barber. Should we be concerned about multiple comparisons in hierarchical Bayesian models? *Methods in Ecology and Evolution*, 10(4):553–564, April 2019. ISSN 2041-210X. doi: 10.1111/2041-210X.13139. Publisher: John Wiley & Sons, Ltd.

Judea Pearl, Madelyn Glymour, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.

Noa Pinter-Wollman, Elizabeth A Hobson, Jennifer E Smith, Andrew J Edelman, Daizaburo Shizuka, Shermin De Silva, James S Waters, Steven D Prager, Takao Sasaki, George Wittemyer, et al. The dynamics of animal social networks: analytical, conceptual, and theoretical advances. *Behavioral Ecology*, 25(2):242–255, 2014.

R Core Team. R: A language and environment for statistical computing. 2020.

Sebastian Sosa, Cristian Pasquaretta, Ivan Puga-Gonzalez, F. Stephen Dobson, Vincent A. Viblanc, and William Hoppitt. Reliable hypotheses testing in animal social network analyses: global index, index of interactions and residual regression. Technical report, December 2021. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

Stan Development Team. RStan: the R interface to Stan, 2020. R package version 2.21.2.

Mark Tranmer, David Steel, and William J. Browne. Multiple-membership multiple-classification models for social network and group dependences. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 177(2):439–455, 2014. doi: <https://doi.org/10.1111/rssa.12021>.

Rens van de Schoot, Sarah Depaoli, Ruth King, Bianca Kramer, Kaspar Märtens, Mahlet G Tadesse, Marina Vannucci, Andrew Gelman, Duco Veen, Joukje Willemsen, et al. Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1):1–26, 2021.

Eric-Jan Wagenmakers, Tom Lodewyckx, Himanshu Kuriyal, and Raoul Grasman. Bayesian hypothesis testing for psychologists: A tutorial on the savage–dickey method. *Cognitive psychology*, 60(3):158–189, 2010.

Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. Springer New York, New York, NY, 2004. ISBN 978-1-4419-2322-6 978-0-387-21736-9. doi: 10.1007/978-0-387-21736-9.

William W. S. Wei. Time Series Analysis, March 2013. ISBN: 9780199934898.

Hal Whitehead. *Analyzing animal societies*. University of Chicago Press, 2008.

Hal Whitehead and Richard James. Generalized affiliation indices extract affiliations from social network data. *Methods in Ecology and Evolution*, 6(7):836–844, 2015. ISSN 2041-210X. doi: 10.1111/2041-210X.12383. [_eprint: https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.12383](https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.12383).

Supplementary Material: Derivation of duration model

With duration data, edge weights are often calculated as:

$$W_{ij} = \frac{X_{ij}}{D_{ij}}$$

where D_{ij} is the maximum amount of time i and j could have been seen engaged in a social event. The total amount of time X_{ij} is the sum of times from K_{ij} social events:

$$X_{ij} = \sum_{n=1}^{K_{ij}} X_{ij}^{(n)}$$

where $X_{ij}^{(n)}$ is the n -th observed social event between i and j , and K_{ij} is the number of social events between i and j over the course of the study period.

This means that edge weights could alternatively be modelled in terms of the mean rate of events per unit time, λ_{ij} and the mean event duration μ_{ij} , as follows:

$$W_{ij} = \mu_{ij}\lambda_{ij} = t_{ij}$$

where t_{ij} is the proportion of time i and j spend engaged in a social event.

Unlike the binary and count models, there is no natural process model for duration data. This leaves the modelling decisions for λ_{ij} and μ_{ij} less clear. Therefore we introduce a minimal process model that makes few assumptions but still attempts to capture as much detail about the generating process for the data as possible. Time-until-event processes are often modelled as exponential processes. We propose to use the exponential distribution, centered on the mean event duration μ_{ij} , to model the duration of events:

$$X_{ij}^{(n)} \sim \text{Exponential}(1/\mu_{ij}).$$

Additionally, the count of events can naturally be modelled as a Poisson process, with rate λ_{ij} . Combining these assumptions gives the following model:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Exponential}(1/\mu_{ij}) \\ K_{ij} &\sim \text{Poisson}(\lambda_{ij}d_{ij}) \\ \text{logit}(\lambda_{ij}\mu_{ij}^{(n)}) &= \text{logit}(t_{ij}^{(n)}) = \omega_{ij} + \dots \end{aligned}$$

Unfortunately this model is challenging to fit because it requires introducing a latent variable to describe $\lambda_{ij}\mu_{ij}^{(n)}$. This latent variable can be marginalised out by using the equality $\lambda_{ij}\mu_{ij}^{(n)} = t_{ij}^{(n)}$, giving an alternative parameterisation of the model:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Exponential}(\lambda_{ij}/t_{ij}^{(n)}) \\ K_{ij} &\sim \text{Poisson}(\lambda_{ij}d_{ij}) \\ \text{logit}(t_{ij}^{(n)}) &= \omega_{ij} + \dots \end{aligned}$$

In this specification, it is theoretically possible for the duration time $X_{ij}^{(n)}$ to exceed the maximum possible duration time. A truncated exponential distribution could be used to prevent this from occurring, but for simplicity we have only presented the base model here. There are many alternative possible specifications for the duration model, so predictive checks will be especially important when using these kinds of model.

Supplementary Material: Non-random edge weight test for duration data

The non-random edge weight test compares two alternative edge weight models, one where social events between dyads depend on a dyad-level edge weight, and one where they depend on a single, population-level edge weight. This tests the hypothesis that apparent differences in edge weight between dyads is due only to sampling, or if there is an underlying difference in edge weights. For duration data, the two models are described by:

H1:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Exponential}(\lambda_{ij}/t_{ij}^{(n)}) \\ K_{ij} &\sim \text{Poisson}(\lambda_{ij}d_{ij}) \\ \log(t_{ij}^{(n)}) &= \omega_{ij} + \dots \end{aligned}$$

H0:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Exponential}(\lambda/t^{(n)}) \\ K_{ij} &\sim \text{Poisson}(\lambda d_{ij}) \\ \log(t^{(n)}) &= \omega + \dots \end{aligned}$$

Supplementary Material: Modelling individual gregariousness

The base edge weight model uses a single term to model edge weight, ω_{ij} , where observed social events depend on it in some way, for example in the binary model:

$$\text{logit}(p_{ij}) = \omega_{ij} + \dots$$

where $\dots$ represent additional control terms for nuisance effects/confounds. This is designed to capture a dyad-level measure of edge weight, but depending on the question and context, it may be necessary to consider additional effects of edge weight, for example node-level gregariousness. This could be modelled quite simply by including a node-level term in the above equation:

$$\text{logit}(p_{ij}) = \gamma_i + \gamma_j + \omega'_{ij} + \dots$$

where γ_i is a latent measure of node-level gregariousness, and ω'_{ij} is a measure of the edge weight of a dyad beyond that explained by individual gregariousness. The original edge weight from the base model can be recovered as $\omega_{ij} = \gamma_i + \gamma_j + \omega'_{ij}$.

This alternative formulation may be useful if individual gregariousness is of particular interest.

Supplementary Material: Including an intercept term in the linear predictor

In some cases it may be useful to include an ‘intercept’ term in the linear predictor part of the edge weight model. For example, in the base model of the edge weight model for binary data, the linear predictor:

$$\text{logit}(p_{ij}) = \omega_{ij} + \dots$$

could be modified to:

$$\text{logit}(p_{ij}) = \bar{\omega} + \omega'_{ij} + \dots$$

where $\bar{\omega}$ is the mean log-odds edge weight of all the dyads, and ω'_{ij} is a zero-centered relative measure of edge weight. This measure is equivalent to taking the standard edge weight ω'_{ij} and mean-centering it.

Supplementary Material: Sharing information between dyads with partial pooling

In some cases it can be useful to assume that edge weights are drawn from the same distribution. This assumption allows information about edge weights between dyads to be shared, and decrease the uncertainty around estimates. This may not always be the case, but the assumption can be checked using predictive checks. An important consideration is for dyads of individuals are rarely seen in the same sampling period. There is little information for these dyads and they will be heavily influenced by any type of pooling. Depending on the accuracy of the pooling assumption, this could greatly improve the inference. Predictive checks should be used when determining the appropriate model for any given dataset.

Partial pooling of edge weights can be implemented by adding an adaptive prior to the edge weight terms. Using the count model as an example, a model using unimodal partial pooling would take the form:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + \dots \\ \omega_{ij} &\sim \text{Normal}(\mu_{\omega}, \sigma_{\omega}^2) \end{aligned}$$

where μ_{ω} and σ_{ω} are the mean and standard deviation of the adaptive prior over edge weights.

edge weights are not necessarily unimodally distributed, and may comprise multiple different types of edge weight. Mixture models have previously been used to study social complexity, but they can also be incorporated into our framework to capture the multimodal nature of edge weights. A model using multimodal partial pooling would take the form:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + \dots \\ \omega_{ij} &\sim \text{Normal}(\mu_{z[ij]}, \sigma_{z[ij]}^2) \\ z_{ij} &\sim \text{Categorical}(\alpha_1, \dots, \alpha_K) \end{aligned}$$

where z_{ij} is an index variable describing the mixture component to which each dyad belongs. The α_k parameters describe the weightings of each mixture component, centered on the component with mean μ_k and standard deviation σ_k .

Supplementary Material: Directed networks and reciprocal effects

The social relations model often used in human social network studies includes a provision for reciprocal effects using correlated effects for sending and receiving effects (Kenny and La Voie, 1984). In undirected networks there is no need for modelling reciprocal effects, but it may be useful to do so in directed networks for two reasons: 1) reciprocal effects may be of interest for the question in hand, and 2) including a reciprocal effect may improve inference and reduce uncertainty around parameter estimates if there is in fact a reciprocal effect at play.

To include reciprocal effects in the base edge weight model, the prior for the directed edge weights can be replaced with a multivariate normal distribution:

$$\begin{pmatrix} \omega_{ij} \\ \omega_{ji} \end{pmatrix} \sim \text{MultivariateNormal} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{\omega}^2 & \sigma_{\omega} r \\ \sigma_{\omega} r & \sigma_{\omega}^2 \end{pmatrix} \right)$$

where σ_{ω} is the standard deviation of edge weights, and r is the correlation between sending and receiving effects. The correlation parameter r can be interpreted as the correlation between sending effects and receiving effects. If the correlation is positive this would indicate that directed social events tend to be reciprocated, whereas a negative correlation implies that social events are imbalanced.

Supplementary Material: Discussion on controlling for non-social effects

There has been considerable attention in the social network literature focusing on methods for controlling for non-social effects in statistical analyses. The key concern is that edge weights may be affected by non-social factors, for example preferences for certain spatial locations, rather than preferences for specific individuals. This is of obvious importance when conducting social network analyses, as it could cause the wrong conclusions to be drawn from statistical models.

A number of methods have been proposed to attempt to control for non-social effects, also termed nuisance effects or confounds. Recent works have focused on computing correction terms to remove the effect of non-social factors from analyses (Farine and Carter, 2021; Sosa et al., 2021). Our framework also introduces a mechanism for controlling for non-social effects. However, there are implicit limitations to all of these methods (including ours) because of the nature of the underlying causal systems. Here we present a few examples in a causal framework using directed acyclic graphs (DAGs) to illustrate the considerations when attempting to control for non-social effects. DAGs and causal inference are a powerful tool for reasoning over causal problems, for those unfamiliar with causal inference, Pearl et al. (2016) provides a thorough primer.

The DAG shown in Figure 3 describes the standard causal system used for inferring sociality in social network analysis. Consider the hypothesis that the trait T of a dyad (e.g. age difference) affects the dyad's sociality, S . Unfortunately, S is unknown and unknowable, but we assume it is borne out in observable social events X , which can be measured in any number of ways. X is therefore a good measure of sociality, and can be used to ask questions about the relationship between T and S .

Example 1

Problems can start to arise when there is an additional effect on X , a non-social effect N . The first example we will consider is shown in Figure 4. In this diagram, T influences S , which in turn influences X . X is also influenced by N , causing a non-social effect on S . N and S are independent, so using X as a measure of S will give an unconfounded, but imprecise measure of sociality. Alternatively, if N can be controlled for, X will be a more precise measure of sociality. In this example, controlling for the non-social effect is not essential, but will yield more precise estimates and improved inference.

Figure 3: T influences S , which influences X . X is only affected by S , so it can serve us well as a proxy for sociality.

Example 2

The second example we will consider is shown in Figure 5. T influences S , which in turn influences X , but T also influences N , which in turn influences X . N and S are now independent, conditional on T . Using X as a proxy of S will no longer work. The relationship between T and X is now confounded by N , so N must be controlled for before X can be used as a measure of sociality. Controlling for N essentially blocks the causal influence of T on X through N , so X becomes a good measure of sociality only when N is controlled for.

Example 3

The final example we consider is shown in Figure 6. In this example T influences S , which in turn influences X , and again T also influences N , which in turn influences X . This time, however, N and S are no longer independent, conditionally or otherwise, there is now a direct influence of S on N . This means that T influences X on two distinct paths: $T \rightarrow S \rightarrow X$, and importantly, $T \rightarrow S \rightarrow N \rightarrow X$. The overall influence of T on X can no longer be parsed out from the effect of N on X . There is no way to separate out the social and non-social effects without additional information.

This example might arise where X is derived from a spatial measure, N is a spatial variable, and T affects both sociality and spatial preference. Since sociality requires individuals to spent time in the same location, but the trait also affects location preference, the non-social spatial preference cannot be parsed out from the edge weight. Including the non-social effect as a control here would remove one of the paths for T to influence

Figure 4: DAG 1: T influences S , which in turn influences X . X is also influenced by N , causing a non-social effect on S . N and S are independent, so controlling for the non-social effect is not essential, but will yield more precise estimates and improved inference.

X , and could remove the effect of interest, but excluding the non-social effect as a control opens up the possibility for a non-social effect to be mistaken as a social effect.

Figure 6: DAG 3: T influences S , which in turn influences X , and again T also influences N , which in turn influences X . There is a direct influence of S on N , preventing social and non-social effects on X from being separated.

Figure 5: DAG 2: T influences S , which in turn influences X , but T also influences N , which in turn influences X . N and S are still independent, but now using X as a proxy of S will no longer work. The relationship between T and X is now confounded by N , so N must be controlled for before X can be used as a measure of S .

Modelling non-social effects

These examples show the potential pitfalls when controlling for non-social effects. In examples 1 and 2, controlling for the non-social effect is always useful and improves the precision of estimates while removing the effect of unimportant non-social effects. However in example 3, including the control could remove the effect of interest when it is there, but excluding the control could deceive us and detect an effect of interest when there is none.

Supplementary Material: Models for aggregated data (collapsed models)

Aggregated binary model

$$X_{ij} \sim \text{Binomial}(D_{ij}, p_{ij})$$

$$\text{logit}(p_{ij}) = \omega_{ij} + \dots$$

Aggregated count model

$$X_{ij} \sim \text{Poisson}(\lambda_{ij} D_{ij})$$

$$\log(\lambda_{ij}) = \omega_{ij} + \dots$$

Aggregated duration model

$$X_{ij} \sim \text{Gamma}(K_{ij}, \lambda_{ij}/t_{ij})$$

$$K_{ij} \sim \text{Poisson}(\lambda_{ij} d_{ij})$$

$$\log(t_{ij}) = \omega_{ij} + \dots$$

Supplementary Material: Modelling group-based data with the binary model

A popular method of data collection in fission-fusion societies is “gambit of the group”. This method of data collection assumes individuals gather in groups according to their social network, so a sighting of three individuals A, B, and C in a group or gathering counts as a social event between A and B, between A and C, and between B and C. Since the observations are either presence or absence of a social event, the binary model is an appropriate model for these data. However, since social events derived from the same grouping are unlikely to be independent, we propose to include a varying effect to control for group membership, as follows:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Bernoulli}(p_{ij}^{(n)}) \\ \text{logit}(p_{ij}^{(n)}) &= \omega_{ij} + G_{ij}^{(n)} + \dots \\ G_{ij}^{(n)} &\sim \text{Normal}(\mu_G, \sigma_G^2) \end{aligned}$$

where $G_{ij}^{(n)}$ is the group-level effect for the n -th grouping that ij were seen in. This parameter will be shared by any observations for dyads in the same grouping. An example of this type of model is included in the example code. Additionally, the `convert_gbi.Rmd` file gives an example of how group-by-individual matrices can be converted into the long format required by Stan.

The long format used to fit these types of models can generate large numbers of observations depending on the number of individuals and groupings. Be aware it may be computationally intensive to process the data and fit the model.