

RESEARCH ARTICLE

BISoN: A Bayesian framework for inference of social networks

Jordan Hart¹  | Michael Nash Weiss^{2,3}  | Daniel Franks^{4,5}  | Lauren Brent⁶ 

¹Department of Psychology, University of Exeter, Exeter, UK; ²Centre for Research in Animal Behaviour, Exeter, UK; ³Center for Whale Research, Friday Harbor, UK; ⁴Department of Biology, University of York, York, UK; ⁵The University of York Heslington, York, UK and ⁶Center for Research in Animal Behaviour, University of Exeter, Exeter, UK

Correspondence

Jordan Hart

Email: jordan.da.hart@gmail.com

Funding information

Engineering and Physical Sciences Research Council, Grant/Award Number: EP/R513210/1; European Research Council Consolidator Grant, Grant/Award Number: 864461; National Institute of Health, Grant/Award Number: R01MH118203 and R01AG060931; Natural Environment Research Council, Grant/Award Number: NE/S009914/1 and NE/S010327/1

Abstract

1. Animal 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 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 many 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 statistical analyses.
2. 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 used in social network analysis.
3. We show how the framework can be applied to common types of data and how various types of downstream statistical analyses can be performed, including non-random association tests and regressions on network properties.
4. 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 both an R package and example R scripts available to enable adoption of the framework.

KEYWORDS

animal social network analysis, Bayesian inference, dyadic regression, network metrics, nodal regression

1 | 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, such as human observers recording social interactions between monkeys or biologists capturing proximity events between people. The type of data used to infer networks and the manner in

Daniel Franks and Lauren Brent are Joint senior authors.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Methods in Ecology and Evolution* published by John Wiley & Sons Ltd on behalf of British Ecological Society.

| 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 (or proportion of time) 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 |

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.

which the data are collected affects both the interpretation of social network analyses and their accuracy (Whitehead, 2008a).

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]; Cairns & Schwager, 1987). These normalised measures ensure that pairs that are observed for longer are not erroneously determined to be more social (Whitehead, 2008a). This is important because observing social events can be challenging and uniform sampling over all pairs is not always possible (Croft et al., 2008). Though normalised measures of social structure 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. Existing solutions to this problem aim to either a priori quantify accuracy and power of social networks (Hart et al., 2021; Whitehead, 2008b), filter out highly uncertain data points (James et al., 2009) or include sampling effort as terms in statistical models (Franks et al., 2021). These existing solutions can be useful, but also necessarily lose information and may introduce biases into analyses. Though uncertainty is likely present in most measurements (e.g. age, size or perhaps even sex), the above-mentioned studies have highlighted that uncertainty due to sampling effort can be orders of magnitude greater than other sources of measurement uncertainty. Any major sources of uncertainty can drastically impact the performance of statistical models, so it is important that uncertainty due to sampling effort is accounted for properly (Wasserman, 2004). Despite this, methods to fully estimate and propagate uncertainty through social network analyses remain largely underdeveloped.

In this paper, we introduce BISoN: a general unified Bayesian framework for modelling social network data. BISoN 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. The BISoN framework comprises (a) an *edge weight model*, dependent on data type, that builds edge weights and networks with uncertainty from empirical data, and (b) downstream analyses that use estimated edge weights to take into account uncertainty in the network. This manuscript is intended as a technical overview of BISoN and assumes some familiarity with the main concepts of Bayesian statistics and social network analysis. Readers unfamiliar with the Bayesian statistical workflow and social network analysis are directed to introductory texts such as Kruschke (2015), McElreath (2020), Carrington et al. (2005), and Croft et al. (2008). 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 BISoN framework presents an opportunity to generate reliable, flexible and rich scientific inference in the study of social systems.

2 | BISoN—BAYESIAN INFERENCE OF SOCIAL NETWORKS

In the following section, we outline the three edge weight models we have developed for binary, count, and duration data (see Table 1 for definitions of these types of data). For brevity, the models are presented without detailing any specific priors but when fitting these models, priors should always be specified, and the choice of prior should depend on the context and the question in hand (van de Schoot et al., 2021). These edge weight models can be used to generate estimates for network edge weights ω_{ij} . We describe how to obtain edge weight estimates for binary, count and duration data below. 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) testing for 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. An overview of the BISoN modelling workflow is

shown in Figure 1. The models we present here are simple examples that can be extended and refined in many ways. 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 released an R package to enable most common types of model and analysis to be used by researchers, at: <https://doi.org/10.5281/zenodo.7611719> (Hart et al., 2022). We have also provided example code for users wishing to develop advanced bespoke models, at: <https://doi.org/10.5281/zenodo.6603327>.

2.1 | Quantifying uncertainty in edge weights

2.1.1 | Edge weight model: Binary data

To demonstrate the notion behind BISoN, we will first describe the edge weight model for the case of binary data. We define binary data as the case where the 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 presence/absence social events that occurred between i and j , and D_{ij} is the number of sampling periods where it is considered possible for a social event to have been observed between i and j . The exact definition will depend on context,

but this may often be the number of periods in which at least one of the individuals of the dyad was seen.

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 (or proportion) as p_{ij} . Without making any additional assumptions beyond that of the standard SRI calculation, 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. A similar formulation was used by (Farine & Strandburg-Peshkin, 2015) to estimate uncertainty over edge weights using a beta conjugate prior over the p_{ij} parameters (as detailed in Fink, 1997).

Social data can sometimes be influenced by non-social effects at the observation level. For example, varying levels of habitat visibility depending on location, variance in observer reliability, or time of day can all affect the recorded observations of social events. These effects cannot be modelled by aggregating social events X_{ij} at the dyad level. To solve this, we propose to model these effects by decomposing the binomial process over the aggregated observations to distinct Bernoulli processes over unaggregated observations, as follows:

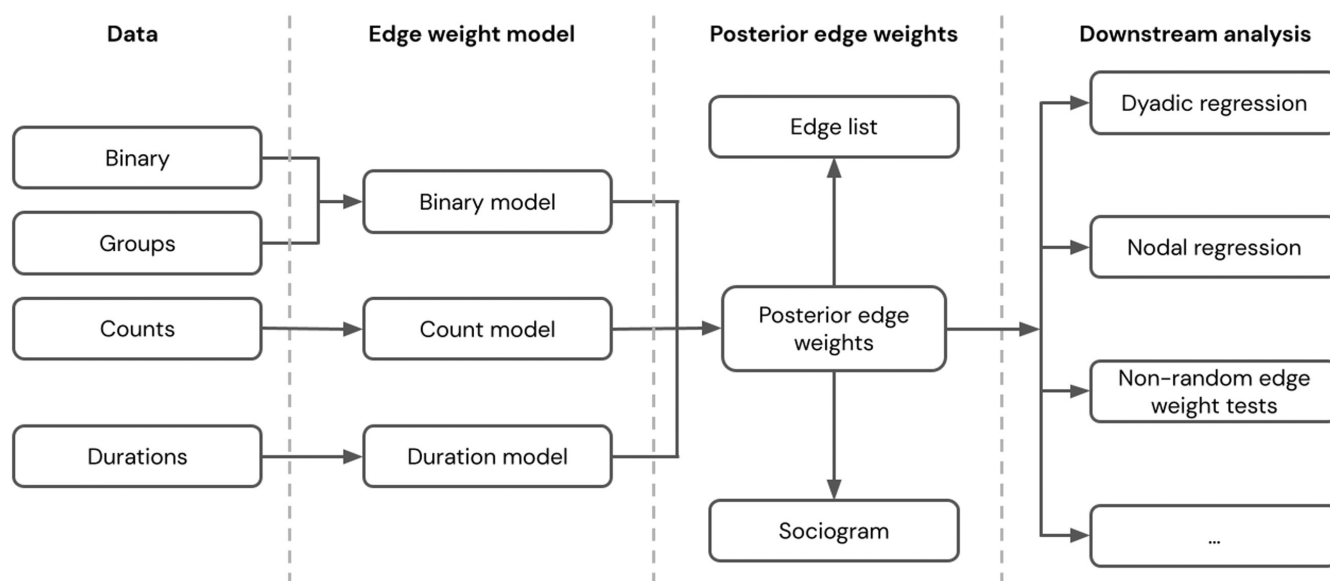


FIGURE 1 Schematic overview of the BISoN framework. Different types of observation data in various levels of aggregation are fed into the appropriate edge weight model for the type of data. The edge weight model generates posterior distributions over edge weights, quantifying uncertainty over them. Posterior edge weights can be visualised in a sociogram with uncertainty, or extracted as an edge list with credible intervals. Posterior edge weights can also be extracted for downstream analyses such as regression analyses, non-random edge weight tests and more.

$$X_{ij}^{(n)} \sim \text{Bernoulli}(p_{ij}^{(n)})$$

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

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. The term $\text{logit}(p_{ij}^{(n)})$ can be seen as analogous to the predictor terms in regression models. In the location-dependent visibility example, including an observation-level location effect will account for this visibility effect. When additional effects are included, the edge weight will become an estimate of relative sociality in the presence of non-social effects. We use a linear combination of effects in these examples for simplicity, but the effects could equally be modelled as non-linear functions if deemed necessary in context. Alternatively dyad- or individual-level effects can also be included to separate out different types of social effect.

The binary data model can also be used for group-based (gam-bit of the group) data (Whitehead & Dufault, 1999). 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 [Supporting Information \(S1\)](#) for further discussions on these points.

2.1.2 | Edge weight model: Count data

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 are recorded. First, note that for count data, point estimates of edge weights are often also defined as $W_{ij} = X_{ij} / D_{ij}$ but 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 & Bové, 2014).

Again, this model does not allow observation-level effects to be modelled, 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 model:

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

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

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 in a later section and in the [Supporting Information \(S2\)](#).

2.1.3 | Edge weight model: Duration data

The edge weight model for duration data is slightly more complex than those of 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 again calculated using $W_{ij} = X_{ij} / D_{ij}$, but 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 [Supporting Information S2](#) for the full derivation):

$$X_{ij}^{(n)} \sim \text{Exponential}(\lambda_{ij} / t_{ij}^{(n)})$$

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

$$\text{logit}(t_{ij}^{(n)}) = \omega_{ij} + \dots$$

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

2.1.4 | 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 variation in sociality between location is not of direct interest to the researcher, 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). The location effect is designed to capture the differences in visibility at different locations. Note that if more direct information such as vegetation density is available that can directly model, for example, visibility, this will always be preferable to weaker proxy variables. In theory, the effect of visibility on edge weights will always be negative but we choose to model it without bounds. This is because the true visibilities in different locations are unknown, so the edge weights cannot be estimated exactly, only proportionally. It is therefore unnecessary to place an upper bound on the visibility effect, making the model easier to fit.

Depending on the assumed causal process that generates observations, it will not 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 considerations around modelling non-social effects is included in the [Supporting Information \(S7\)](#).

3 | DOWNSTREAM ANALYSIS WITH UNCERTAINTY

The previous section has shown how uncertainty can be estimated over edge weights. Once these estimates have been obtained, it is then possible to conduct various types of statistical analysis on the network while propagating uncertainty through the entire analysis. This approach will lead to lower 'false positive' error rates because network quantities will have higher variances. Despite this, because the uncertainty captured by BISoN reflects true uncertainty in the data, it will also have a 'false negative' error rate no higher than comparable analyses that do not quantify uncertainty. In this section we will discuss how several common types of social network analysis can be applied to social networks constructed with BISoN edge weight models. In particular, we outline how three specific types of network analysis can be conducted: (1) testing for non-random edge weights, (2) dyadic regression and (3) nodal regression.

3.1 | Testing for non-random edge weights

BISoN enables a Bayesian test analogous to the classic non-random association test developed by (Bejder et al., 1998) from a species co-occurrence test developed by (Manly, 1995), and originally proposed as the 'chance sociogram' by (Moreno & Jennings, 1938). At its core, this test asks whether the structure of an observed network is an artefact of sampling alone 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

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

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 weight. 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

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

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.

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

$$\log(\lambda_{ij}^{(n)}) = \omega + \beta_{\text{Age}} \cdot A_{ij} + L_{ij}^{(n)},$$

where β_{Age} is a parameter describing the effects of age on the log rate of social events, and A_{ij} is the age difference for the dyad ij . In addition to modelling continuous variables, this approach to the non-random edge weight testing also makes it possible to compare multiple competing hypotheses of edge structure, model the effects of different dyad-level covariates, and compare hypotheses relating to global network metrics such as network density or clustering coefficients.

3.2 | 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:

$$\begin{aligned} X_{ij}^{(n)} &\sim \text{Poisson}(\lambda_{ij}^{(n)} D_{ij}^{(n)}) \\ \log(\lambda_{ij}^{(n)}) &= \omega_{ij} + L_{ij}^{(n)} \\ \hline \omega_{ij} &\sim \text{Normal}(\beta_0 + \beta_1 x_{ij} + u_i + u_j, \sigma^2) \\ u &\sim \text{Normal}(0, \sigma_u^2) \\ \sigma_u &\sim \dots \end{aligned}$$

where β_0 is the intercept parameter, β_1 is the slope parameter, u_i are parameters accounting for the effect of some edges sharing the same node, σ is the standard deviation of the residuals, and σ_u is a hyperprior for the standard deviation of the edge-sharing effect (undefined here to avoid suggesting a default prior). The u parameters are used to capture non-independence of dyads, and are sometimes known as multi-membership effects (Tranmer et al., 2014). Further discussion of multi-membership models in social network analysis can be found in (Hart et al., 2021). The dashed line indicates that the model can be fit in two separate parts: the first part being the 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 posterior distributions describing plausible values of the coefficients.

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.

3.3 | 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

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

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 cannot 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 maintains 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.

4 | 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 28 dyads corresponding to all possible pairs of 8 individuals living in a closed social group. Individuals 1–4 were assigned to a treatment condition and individuals 4–8 were assigned to a control. Dyads between two individuals within 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 six 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

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

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, $D_{ij}^{(n)}$ is the length of the corresponding sampling period (fixed to 1 here for simplicity), ω_{ij} is the dyad-level edge weight, and $L_{ij}^{(n)}$ is an observation-level effect corresponding to one of the six 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 Stan, and fitted in R, using the package *Rstan* (R Core Team, 2020; Stan Development Team, 2020). Other approaches such as integrated nested Laplace approximation may also be used, see the Supporting Information for a brief discussion of available fitting algorithms (S6) (Rue et al., 2009).

The model fit was checked visually by examining the chains and computing the Gelman-Rubin \hat{R} convergence statistic (Gelman & Rubin, 1992). Once the model fit was verified, the

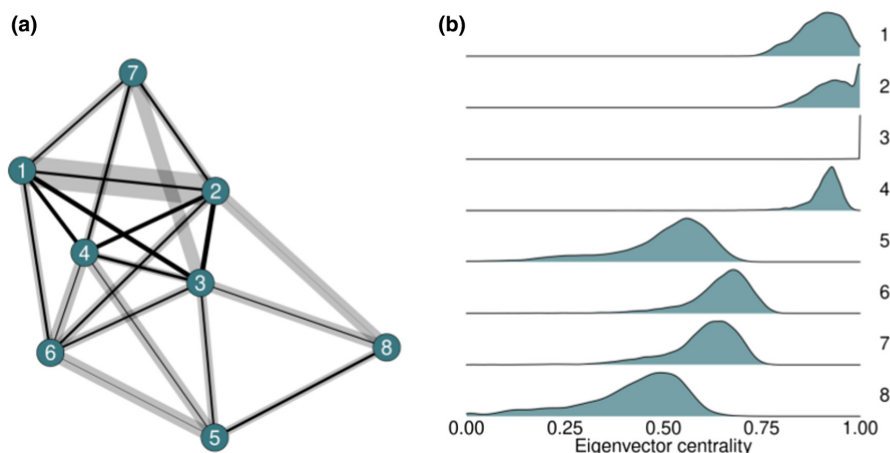
parameter estimates of edge weights ω_{ij} between dyads were extracted. Figure 2a shows a sociogram visualisation of the network based on edge weight estimates ω_{ij} , where the bands around edges are proportional to the widths of the 95% credible intervals. At this point the estimated network can 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 BISO 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 eight nodes are shown in Figure 2b. 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

FIGURE 2 (a) Sociogram visualisation of a network generated from uncertain edge weights. The width of edges denotes their weights, and the bands around the edges are proportional to the width of their 95% credible intervals. Low edge weights are arbitrarily removed from this plot (though not the analysis) for visualisation purposes. (b) Posterior distributions of node-level eigenvector centralities for each node in the network. Nodes 1–4 are those assigned to the treatment and nodes 5–8 to the control.



the control parameter. 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.

5 | EXAMPLE: SOUTHERN RESIDENT KILLER WHALES

To demonstrate BISoN on empirical data, we used a publicly-available dataset of near-surface physical contact interactions between 22 southern resident killer whales (*Orcinus orca*) (Weiss et al., 2021). We fitted a BISoN count edge model to the data and ran a regression analysis to study the associations between node strength and both age and sex. See Supporting Information (S4) for replicable code of this analysis. We opted for a standard linear model with a Gaussian family distribution, implemented using the `bison_brm` function in `bisonR`. From previous work, we expected a negative association between node centrality and age, and we also expected males to be less central than females (Weiss et al., 2021). Our findings were in line with this, and we found a negative association of age with centrality, with an estimated age coefficient of -0.57 (95% CI: $-0.80, -0.35$). We also found that males had a lower centrality than females, with an estimated sex coefficient of -9.99 (95% CI: $-15.32, -4.67$). We hope this findings serve as a useful empirical example of how BISoN can be used to run regression analyses.

6 | DISCUSSION

In this paper, we presented BISoN, a general, highly flexible framework for modelling and analysing social networks. BISoN 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 BISoN framework in R with Stan at: <https://doi.org/10.5281/zenodo.6603327>.

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 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 Supporting Information S3 for further information).

We see BISoN as a generalised synthesis and extension of a wide range of existing methods. In particular, for binary data, BISoN implements a variant of the beta binomial *measurement error model*, which has been previously used in animal social network analysis (Farine & Strandburg-Peshkin, 2015; Fink, 1997; Fuller, 1987). BISoN's mechanism for modelling multiple effects was inspired by generalised linear models, and can be seen as an extension of social relations models (Kenny & La Voie, 1984; McCulloch & Searle, 2004). This concept was introduced to the animal social network literature by (Whitehead & James, 2015) as generalised affiliation indices. The BISoN framework shares some conceptual similarities with the recently developed STRAND framework, but implements measurement models specifically designed for common types of empirical data, and allows propagation of uncertainty to downstream analyses.

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 (Conn et al., 2018). It is also important to note that using additional terms and altering the 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, have a slightly different interpretation in this framework. This is because, even though a particular dyad may never have been seen interacting in, say 100 samples, it is still possible that they would interact in the 101st sample. The edge weight models BISoN 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 BISoN: (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. In some cases this approach may prove 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 [Supporting Information S5](#) 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 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 [Supporting Information \(S7\)](#) 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. On a practical level, BISO-N models will usually need to be fit using a Markov chain Monte Carlo (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 [Supporting Information S8](#)). Furthermore, MCMC methods do not guarantee convergence, so additional checks will be required to make sure the sampler has performed correctly (Cowles & Carlin, 1996; Draper, 2008). Together these limitations make the process of model building, fitting and checking a more complex undertaking than simply building networks from point estimates. However, in return for these investments, researchers will have access to a powerful suite of tools that can make maximum use of hard-won data.

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. Furthermore, missing data can be estimated from other variables, which could be especially useful in populations where traits such as age are often estimated or unknown (Little & 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.

7 | 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 relational 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

Jordan Hart conceived of the initial idea behind BISO-N, which was further developed with input from Michael Nash Weiss, Daniel Franks and Lauren Brent. The example code was written by Jordan Hart. The manuscript was written by Jordan Hart with input from Michael Nash Weiss, Daniel Franks and Lauren Brent.

ACKNOWLEDGEMENTS

The authors thank Rebecca Padgett, Josefine Bohr Brask, Delphine De Moor, Samuel Ellis, Kenneth Keuk, Mario Angst and members of the CRAB social network club for useful discussions on early versions of the manuscript.

FUNDING INFORMATION

JH acknowledges funding from the Engineering and Physical Sciences Research Council (grant number EP/R513210/1). LB acknowledges funding from a European Research Council Consolidator Grant (FriendOrigins—864461) and from the National Institute of Health (R01AG060931, R01MH118203). DF and MNW acknowledge funding from the Natural Environment Research Council (grant number NE/S010327/1). DF also acknowledges funding from the Natural Environment Research Council (grant number NE/S009914/1).

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interest.

DATA AVAILABILITY STATEMENT

The BISO-N example code is indexed at <https://doi.org/10.5281/zenodo.6603327>. The bisonR package is indexed at <https://doi.org/10.5281/zenodo.7611719>.

ORCID

Jordan Hart  <https://orcid.org/0000-0002-4636-0760>

Michael Nash Weiss  <https://orcid.org/0000-0002-7422-0538>

Daniel Franks  <https://orcid.org/0000-0002-4832-7470>

Lauren Brent  <https://orcid.org/0000-0002-1202-1939>

REFERENCES

- Altman, D. G., & Royston, P. (2006). The cost of dichotomising continuous variables. *British Medical Journal*, 332(7549), 1080. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1458573/>
- Bejder, L., Fletcher, D., & Bräger, S. (1998). A method for testing association patterns of social animals. *Animal Behaviour*, 56(3), 719–725. <https://doi.org/10.1006/anbe.1998.0802>
- Cairns, S. J., & Schwager, S. J. (1987). A comparison of association indices. *Animal Behaviour*, 35(5), 1454–1469.

- Carrington, P. J., Scott, J., & Wasserman, S. (2005). *Models and methods in social network analysis* (Vol. 28). Cambridge University Press.
- Conn, P. B., Johnson, D. S., Williams, P. J., Melin, S. R., & Hooten, M. B. (2018). A guide to Bayesian model checking for ecologists. *Ecological Monographs*, 88(4), 526–542.
- Cowles, M. K., & Carlin, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91(434), 883–904.
- Croft, D. P., Darden, S. K., & Wey, T. W. (2016). Current directions in animal social networks. *Current Opinion in Behavioral Sciences*, *Behavioral Ecology*, 12(December), 52–58. <https://doi.org/10.1016/j.cobeha.2016.09.001>
- Croft, D. P., James, R., & Krause, J. (2008). *Exploring animal social networks*. Princeton University Press.
- Draper, D. (2008). Bayesian multilevel analysis and MCMC. In J. de Leeuw & E. Meijer (Eds.), *Handbook of multilevel analysis* (pp. 77–139). Springer New York. https://doi.org/10.1007/978-0-387-73186-5_2
- Ellis, S., Franks, D. W., Weiss, M. N., Cant, M. A., Domenici, P., Balcomb, K. C., Ellifrit, D. K., & Croft, D. P. (2021). 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.
- Farine, D. R., & Strandburg-Peshkin, A. (2015). Estimating uncertainty and reliability of social network data using Bayesian inference. *Royal Society Open Science*, 2(9), 150367. <https://doi.org/10.1098/rsos.150367>
- Fink, D. (1997). A compendium of conjugate priors. <http://www.people.cornell.edu/pages/df36/conjintnrnew%20tex.pdf46>
- Franks, D. W., Weiss, M. N., Silk, M. J., Perryman, R. J. Y., & Croft, D. P. (2021). Calculating effect sizes in animal social network analysis. *Methods in Ecology and Evolution*, 12(1), 33–41. <https://doi.org/10.1111/2041-210X.13429>
- Freeman, L. C. (1978). Centrality in social networks conceptual clarification. *Social Networks*, 1(3), 215–239.
- Fuller, W. A. (1987). *Measurement error models*. John Wiley & Sons.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457–472. <https://doi.org/10.1214/ss/1177011136>
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., Leslie, D. S., Forster, J. J., Wagenmakers, E.-J., & Steingrover, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, 81, 80–97.
- Hart, J. D. A., Franks, D. W., Brent, L. J. N., & Weiss, M. N. (2021). Accuracy and power analysis of social interaction networks. *bioRxiv*. <https://doi.org/10.1101/2021.05.07.443094>
- Hart, J. D. A., Franks, D. W., Brent, L., & Weiss, M. N. (2022). *bisonR—Bayesian inference of social networks with r*. <https://osf.io/ywu7j>
- Held, L., & Bové, D. S. (2014). *Applied statistical inference*. Springer Berlin Heidelberg. <https://doi.org/10.1007/978-3-642-37887-4>
- James, R., Croft, D. P., & Krause, J. (2009). Potential Banana skins in animal social network analysis. *Behavioral Ecology and Sociobiology*, 63(7), 989–997. <https://doi.org/10.1007/s00265-009-0742-5>
- Kenny, D. A., & La Voie, L. (1984). The social relations model. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 18, pp. 141–182). Elsevier.
- Kruschke, J. K. (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). Academic Press.
- Little, R. J. A., & Rubin, D. B. (Eds.). (2002). Bayes and multiple imputation. In *Statistical analysis with missing data* (pp. 200–220). John Wiley & Sons. <https://doi.org/10.1002/9781119013563>
- Manly, B. F. J. (1995). A note on the analysis of species co-occurrences. *Ecology*, 76(4), 1109–1115.
- McCulloch, C. E., & Searle, S. R. (2004). *Generalized, linear, and mixed models*. John Wiley & Sons.
- McElreath, R. (Ed.). (2020). Statistical rethinking: A Bayesian course with examples in R and Stan. In *CRC texts in statistical science* (2nd ed.). CRC Press. <https://doi.org/10.1201/9780429029608>
- Moreno, J. L., & Jennings, H. H. (1938). *Statistics of social configurations* (Vol. 1, pp. 342–374). Sociometry.
- Pinter-Wollman, N., Hobson, E. A., Smith, J. E., Edelman, A. J., Shizuka, D., De Silva, S., Waters, J. S., Prager, S. D., Sasaki, T., Wittemyer, G., & Fewell, J. (2014). The dynamics of animal social networks: Analytical, conceptual, and theoretical advances. *Behavioral Ecology*, 25(2), 242–255.
- R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Rue, H., Martino, S., & Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(2), 319–392. <https://doi.org/10.1111/j.1467-9868.2008.00700.x>
- Stan Development Team. (2020). *RStan: The R interface to Stan*. <http://mc-stan.org/>
- Tranmer, M., Steel, D., & Browne, W. J. (2014). Multiple-membership multiple-classification models for social network and group dependencies. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 177(2), 439–455. <https://doi.org/10.1111/rssa.12021>
- van de Schoot, R., Depaoli, S., King, R., Kramer, B., Märtens, K., Tadesse, M. G., Vannucci, M., Gelman, A., Veen, D., Willemssen, J., & Yau, C. (2021). Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1), 1–26.
- Wasserman, L. (2004). *All of statistics: A concise course in statistical inference*. Springer texts in statistics. Springer. <https://doi.org/10.1007/978-0-387-21736-9>
- Wei, W. W. S. (2013). Time series analysis. In *The Oxford Handbook of Quantitative Methods in Psychology* (Vol. 2). Statistical Analysis, Oxford Library of Psychology. <https://doi.org/10.1093/oxfordhb/9780199934898.013.0022>
- Weiss, M. N., Franks, D. W., Giles, D. A., Youngstrom, S., Wasser, S. K., Balcomb, K. C., Ellifrit, D. K., Domenici, P., Cant, M. A., Ellis, S., Nielsen, M. L. K., Grimes, C., & Croft, D. P. (2021). Age and sex influence social interactions, but not associations, within a killer whale pod. *Proceedings of the Royal Society B*, 288(1953), 20210617.
- Whitehead, H. (2008a). *Analyzing animal societies*. University of Chicago Press.
- Whitehead, H. (2008b). Precision and power in the analysis of social structure using associations. *Animal Behaviour*, 75(3), 1093–1099. <https://doi.org/10.1016/j.anbehav.2007.08.022>
- Whitehead, H., & Dufault, S. (1999). Techniques for analyzing vertebrate social structure using identified individuals. *Advances in the Study of Behavior*, 28, 33–74.
- Whitehead, H., & James, R. (2015). Generalized affiliation indices extract affiliations from social network data. *Methods in Ecology and Evolution*, 6(7), 836–844. <https://doi.org/10.1111/2041-210X.12383>

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

Appendix S1.

How to cite this article: Hart, J., Weiss, M. N., Franks, D., & Brent, L. (2023). BISON: A Bayesian framework for inference of social networks. *Methods in Ecology and Evolution*, 00, 1–10. <https://doi.org/10.1111/2041-210X.14171>