BGGM: A R Package for Bayesian Gaussian Graphical Models

Donald R. Williams

University of California, Davis

Joris Mulder

Tilburg University

Abstract

Gaussian graphical models (GGM) allow for learning conditional independence structures that are encoded by partial correlations. Whereas there are several R packages for classical (i.e., frequentist) methods, there are only two that implement a Bayesian approach. These are exclusively focused on identifying the graphical structure. The R package **BGGM** not only contains novel Bayesian methods for this purpose, but it also includes Bayesian methodology for extending inference beyond identifying non-zero relations. **BGGM** is built around two Bayesian approaches for inference—estimation and hypothesis testing. The former focuses on the posterior distribution and includes extensions to assess predictability, as well as methodology to compare partial correlations. The latter includes methods for Bayesian hypothesis testing, in both exploratory and confirmatory contexts, with the novel matrix- F prior distribution. This allows for testing order and equality constrained hypotheses, as well as a combination of both with the Bayes factor. Further, there are two approaches for comparing GGMs across any number of groups with either the posterior predictive distribution or with Bayesian hypothesis testing. This work describes the software implementation of these methods. We end by discussing future directions for **BGGM**.

Keywords: Bayesian, Gaussian graphical model, estimation, hypothesis testing, Bayes factor, prediction, partial correlation, R.

1. Introduction

Gaussian graphical models (GGM) capture conditional dependencies between random variables (Dempster 1972). This is accomplished by identifying off-diagonal elements in the inverse of the covariance matrix (i.e., precision matrix) that are not zero (Lauritzen 1996). When these elements are standardized and the sign is reversed, they correspond to partial correlations. Assuming multivariate normality, non-zero partial correlations imply pairwise, conditional dependent relations, in which all other variables in the model have been controlled for (Baba, Shibata, and Sibuya 2004; Baba and Sibuya 2005). This powerful framework for learning the structure of multivariate data has been used extensively across the sciences, for example, spanning from financial (Anufriev and

Panchenko 2015) to climate “networks” (Zerenner, Friederichs, Lehnertz, and Hense 2014). To date, the most common applications have been gene co-expression (Schäfer and Strimmer 2005) and functional connectivity applications (Das, Sampson, Lainscsek, Muller, Lin, Doyle, Cash, Halgren, and Sejnowski 2017). Here it is common to have more variables p than observations n (i.e., high dimensional data), which requires some form of regularization to make estimation possible (Ledoit and Wolf 2004; Kuismin and Sillanpää 2017). Perhaps the most popular approach employs ℓ_1 -regularization (Tibshirani 1996). This is implemented in various R packages, including **glasso** (Friedman, Hastie, and Tibshirani 2008), **huge** (Zhao, Liu, Roeder, Lafferty, and Wasserman 2012), and **flare** (Li and Redden 2015). Importantly, these packages provide point estimates which do not readily allow for statistical inference. The topic of computing p -values and confidence intervals from ℓ_1 -regularized models is an active area of research (Janková and van de Geer 2017; Javanmard and Montanari 2013; Zhang and Zhang 2014), with methods specifically for GGMs provided in **SILGGM** (Zhang, Ren, and Chen 2018).

More recently, GGMs have become increasingly popular in the social-behavioral sciences (Epskamp, Waldorp, Mottus, and Borsboom 2018b). In these applications, it is uncommon to have high dimensional data and thus regularization is not necessary (Williams, Rhemtulla, Wysocki, and Rast 2019c; Williams and Rast 2018). **BGGM** could accommodate high dimensional data, but the current version is explicitly built for low dimensional situations ($p < n$). Further, the focus is not exclusively on point estimates, but extending inference to include, for example, partial correlation differences (Epskamp, Borsboom, and Fried 2018a) and predictive accuracy (Haslbeck and Waldorp 2018; Haslbeck and Fried 2017). **bootnet** provides bootstrapping capabilities for the former (Epskamp *et al.* 2018a). The approach is computationally intensive, in that a ℓ_1 -regularized model is selected for each iteration. Conversely, **BGGM** provides a computationally efficient approach that allows for comparing partial correlations. The **mgm** package also employs ℓ_1 -regularization and allows for assessing predictability of “networks” (Haslbeck and Waldorp 2015). Here point estimates for variance explained (R^2) and mean squared error are provided. **BGGM** not only provides both measures of prediction error, but each includes posterior uncertainty—e.g., Bayesian R^2 (Gelman, Goodrich, Gabry, and Ali 2017) and Bayesian leave-one-out cross-validation (Vehtari, Gelman, and Gabry 2017).

Another area of interest, particularly in psychopathology applications (Fried, Eidhof, Palic, Costantini, Dijk, Bockting, Engelhard, Armour, Nielsen, and Karstoft 2018), is comparing GGMs estimated from different groups. To our knowledge, **NetworkComparisonTest** is the only R package with this capability (van Borkulo, Boschloo, Kossakowski, Tio, Schoevers, Borsboom, and Waldorp 2016). It is limited to detecting differences between two groups. That is, it cannot provide evidence for equality or the null hypothesis across any number of groups. Conversely, **BGGM** includes methods for comparing GGMs estimated from any number groups with either the predictive distribution or Bayesian hypothesis testing. In (Williams, Rast, Pericchi, and Mulder 2019a), the predictive approach was shown to have more power than the methods implemented in **NetworkComparisonTest**. Further, our package allows for testing invariances in graphical structures with the Bayes factor. The

methodology included in **BGGM** is currently the only approach available for comparing GGMs across any number of groups.

To our knowledge, **BDgraph** (Mohammadi and Wit 2015a,b) and **beam** (Leday and Richardson 2018) are the only R packages currently available for estimating Bayesian GGMs. They exclusively focus on determining the graphical structure—i.e., detecting conditional relations. For this purpose, **BGGM** uses either an analytic solution or the encompassing prior approach (Klugkist, Kato, and Hoijtink 2005). For the latter, estimating the graph is formulated as hypothesized effect sizes expressed with the novel matrix- F prior distribution (Mulder and Raúl Pericchi 2018). **BGGM** also includes extensions that go beyond detecting the conditional independence structure. These include Bayesian approaches for assessing in- and out-of-sample predictive accuracy, testing for differences between partial correlations, and importantly, *confirmatory* hypothesis testing with the Bayes factor. This allows for rigorously testing scientific theories or expectations in GGMs, as well as explicitly gaining evidence for the null hypothesis of conditional independence (Williams and Mulder 2019). Further, **BGGM** includes two approaches for comparing any number of GGMs (Williams *et al.* 2019a)—e.g., those estimated from different groups (male vs. female). Together, **BGGM** is a comprehensive R package that includes newly introduced methodology for Bayesian Gaussian graphical models.

This work is organized as follows. Section 2 introduces notation specific to GGMs. The following sections are then structured around two general approaches for Bayesian inference—estimation and hypothesis testing. Section 3 includes estimation based methods that make exclusive use of the posterior distributions. Section 4 includes method based on Bayesian model selection with the Bayes factor. Section 5 includes the methods for comparing GGMs with the posterior predictive distribution and Bayesian hypothesis testing. We end by discussing future directions for **BGGM**.

2. The Gaussian graphical model

The Gaussian graphical model captures undirected, conditional relationships (Lauritzen 1996), that are typically visualized to infer the underlying structure (i.e., the “network”; Højsgaard, Edwards, and Lauritzen 2012). The undirected graph is $G = (V, E)$, and includes a vertex set $V = \{1, \dots, p\}$ as well as an edge set $E \subset V \times V$. Let $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_p)^\top$ be a $n \times p$ matrix, where each $\mathbf{y}$ is a n dimensional vector that is indexed by the graphs vertices. Assume they follow a multivariate normal distribution $\mathbf{Y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with the mean vector $\boldsymbol{\mu} = (0_1, \dots, 0_p)^\top$ and $p \times p$ positive definite covariance matrix $\boldsymbol{\Sigma}$. Denote the precision matrix $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$. The graph is obtained from the off-diagonal elements $\theta_{ij} \in \boldsymbol{\Theta}_{ij}$. This is used to construct an adjacency matrix A that follows

$$A_{ij} = \begin{cases} 1, & \text{if } \theta_{ij} \neq 0, \ 1 \leq i < j \leq p \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

with $1 \leq i < j \leq p$ denoting the elements in the upper-triangular of the $p \times p$ matrix. Further, $(i, j) \in E$ when the variables i and j are *not* conditionally independent and set to

zero otherwise. Note that the selected edges, when standardized, are partial correlations (ρ) determined to be non-zero. These are computed directly from the precision matrix—i.e.,

$$\rho_{ij} = \frac{-\theta_{ij}}{\sqrt{\theta_{ii}\theta_{jj}}}, \quad 1 \leq i < j \leq p. \quad (2)$$

In **BGGM**, these partial correlations are explicitly used for the Bayes factor based approaches (Section 4), whereas the precision matrix is primarily used for the estimation based methods (Section 3).

3. Estimation

The following methods are described in Williams (2018). They build upon two basic ideas. First, the Wishart distribution is a conjugate prior distribution for Θ . This provides an analytic solution for determining E , and allows for conveniently drawing posterior samples. Second, there is an exact relationship between estimating Θ directly and with multiple regression. Here the individual elements, from joint posterior distribution of Θ , can be converted to their respective regression counterparts. This allows for assessing nodewise (for each variable in the model) predictability.

3.1. Wishart distribution

The joint posterior density for the precision matrix follows

$$p(\Theta|\mathbf{Y}) \propto p(\mathbf{Y}|\Theta)p(\Theta), \quad (3)$$

where $\mathbf{Y}$ is a $n \times p$ matrix drawn from a multivariate normal distribution—i.e.,

$$\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \Theta^{-1}). \quad (4)$$

When using the conjugate Wishart prior, $\mathcal{W}(k, \epsilon \mathbf{I}_p)$, with k degrees of freedom and identity matrix $\mathbf{I}_p$, the posterior distribution also has a Wishart distribution—i.e.,

$$\Theta|\mathbf{Y} \sim \mathcal{W}(k+n, (\mathbf{S} + \epsilon \mathbf{I}_p)^{-1}), \quad (5)$$

where $\mathbf{S}$ is the sums of squares matrix $\mathbf{Y}'\mathbf{Y}$ and ϵ is a constant. The posterior mode then has a closed form

$$\text{argmax}_{\Theta} p(\Theta|\mathbf{Y}) = (k+n-p-1)(\mathbf{S} + \epsilon \mathbf{I}_p)^{-1}, \quad (6)$$

as well as the posterior posterior variance

$$\text{Var}(\Theta|\mathbf{Y}) = (k+n)((\mathbf{S} + \epsilon \mathbf{I}_p)^{-1^2} + \mathbf{d}\mathbf{d}'), \quad (7)$$

where $\mathbf{d} = \text{diag}(\mathbf{S} + \epsilon \mathbf{I}_p)$. By setting $k = p + 1$, and ϵ to a small value, say, 1.0×10^{-10} , results (approximately) in the maximum likelihood estimate $n(\mathbf{S}^{-1})$. Consequently,

by assuming $p(\theta_{ij}|\mathbf{Y})$ is normally distributed (see pg. 88: “Asymptotic normality and consistency”; Gelman, Carlin, Stern, Dunson, Vehtari, and Rubin 2014), allows for constructing (equal tailed) credible intervals and computing posterior probabilities. For example, the former can be obtained as

$$\int_l^u p(\theta_{ij}|\mathbf{Y})d\theta_{ij} = 1 - \alpha, \quad (8)$$

where l and u denote the lower and upper bounds of the interval. What remains is the choice of α , which as shown in Williams (2018), can be (approximately) calibrated to the desired level of specificity¹. The latter follows

$$\Pr(\theta_{ij} > 0|\mathbf{Y}), \quad (9)$$

which corresponds to the posterior probability of a positive effect. Both can be computed analytically with the point estimate in (6) and variance in (7). In **BGGM**, this simple solution can be used for determining E and analytically deriving “network” predictability (Section 3.4).

There is one drawback of this analytic form. Because there are no posterior samples, this limits its applicability to some of the functions in **BGGM**. Indeed, capturing uncertainty is a key aspect of **BGGM** compared to, say, **glasso** and **BDgraph** (both provide point estimates for Θ). We thus include an additional approach for conveniently drawing samples from the posterior distribution. We implemented the noninformative Jeffreys’ prior $|\Theta|^{(p+1)/2}$. This allows for posterior sampling directly from a Wishart distribution—i.e.,

$$\Theta|\mathbf{Y} \sim \mathcal{W}(n - 1, \mathbf{S}^{-1}). \quad (10)$$

These samples, $s = 1, \dots, S$, are used to compute the posterior distribution for the $p \times p$ partial correlation matrix, with $\rho_{ij} \in \mathcal{P}$ —i.e.,

$$\mathcal{P}^{(s)} = -([\text{diag}(\boldsymbol{\sigma})^{(s)}]^{-1}\Theta^{(s)}[\text{diag}(\boldsymbol{\sigma})^{(s)}]^{-1}), \quad (11)$$

where $\boldsymbol{\sigma}$ are the square roots of $\text{diag}(\Theta)$ (Barnard, McCulloch, and Meng 2000) and multiplying by -1 reverses the direction ($\pm$) as in (2). Because ρ is a standardized effect, and thus each ρ_{ij} is on the same scale, **BGGM** allows for defining a *region of practical equivalence* (Kruschke 2011). Rather than determining E in reference to 0, as in (8) and (9), a neighborhood around zero is defined (a null area; Betancourt 2018). Consequently, given some region, $|\rho_{ij}| \leq \vartheta_0$, there is support for null values when the posterior probability is above a pre-determined threshold,

¹Specificity is a diagnostic measure commonly used to assess the performance of methods used to estimate GGMs. It is equal to $1 - \alpha$.

$$\int_{-\vartheta_0}^{\vartheta_0} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} > 1 - \alpha, \quad (12)$$

where $p(\rho_{ij}|\mathbf{Y})$ is the posterior distribution. Conversely,

$$\int_{-\infty}^{-\vartheta_0} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} + \int_{\vartheta_0}^{\infty} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} > 1 - \alpha, \quad (13)$$

denotes the posterior probability *outside* of the region of practical equivalence exceeds $1 - \alpha$. This is implemented in **BGGM** for both determining conditional dependent and practically independent relations, as well as comparing two edges (e.g., $\rho_{1,2} - \rho_{1,3}$).

3.2. Structure learning

By structure learning we are referring to E , which consists of those edges determined to be non-zero. For demonstrative purposes, we consider a relatively small number of variables ($p = 5$). This keeps the returned information concise.

The analytic form is implemented with:

```
R> Y <- BGGM::bfi[,1:5]
R> fit_analytic <- estimate(Y, analytic = T)
R> summary(fit_analytic)
```

```
BGGM: Bayesian Gaussian Graphical Models
```

```
---
```

```
Type: Estimation (Analytic Solution)
```

```
Posterior Samples:
```

```
Observations (n): 2709
```

```
Variables (p): 5
```

```
Edges: 10
```

```
---
```

```
Call:
```

```
estimate.default(x = Y, analytic = T)
```

```
---
```

```
Date: Mon May 06 09:49:06 2019
```

Note `summary(.)` provides information about the fitted model, including that the analytic solution was used, the number of observations (n) and variables (p), and the number of edges.

The edge set E is then selected with:

```
R> E <- select(fit_analytic, ci_width = 0.95)
R> summary(E)
```

```

BGGM: Bayesian Gaussian Graphical Models
---
Type: Selected Graph (Analytic Solution)
Credible Interval: 95 %
Connectivity: 80 %
---
Call:
select.estimate(x = fit_analytic, ci_width = 0.95)
---
Selected:

Partial correlations

      1      2      3      4      5
1  0.00 -0.24 -0.11  0.00  0.00
2 -0.24  0.00  0.29  0.16  0.16
3 -0.11  0.29  0.00  0.18  0.36
4  0.00  0.16  0.18  0.00  0.12
5  0.00  0.16  0.36  0.12  0.00
---

Adjacency

  1 2 3 4 5
1 0 1 1 0 0
2 1 0 1 1 1
3 1 1 0 1 1
4 0 1 1 0 1
5 0 1 1 1 0
---

```

The analytic solution estimates Θ . The signs are in the opposite direction of the partial correlations (Equation 2), which in our experience, can lead to confusion. Thus further information (e.g., credible intervals) is not provided about the analytic solution. The sampling based method can be used to summarize the partial correlations. This is implemented with:

```

R> fit_sampling <- estimate(Y, analytic = F)
R> E <- select(fit_analytic, ci_width = 0.95, rope = NULL, prob = NULL)
R> summary(E, summarize = T, digits = 2)

```

```

BGGM: Bayesian Gaussian Graphical Models
---

```

```

Type: Selected Graph (Sampling)
Credible Interval: 95 %
Connectivity: 80 %
---
Call:
select.estimate(x = fit_analytic, ci_width = 0.95)
---
Estimates:

     egde post_mean post_sd  2.5% 97.5%
1--2   -0.2404    0.018 -0.275 -0.204
1--3   -0.1075    0.019 -0.146 -0.071
2--3    0.2861    0.018  0.251  0.320
1--4   -0.0073    0.019 -0.044  0.030
2--4    0.1650    0.019  0.128  0.200
3--4    0.1778    0.019  0.140  0.213
1--5   -0.0086    0.019 -0.046  0.030
2--5    0.1560    0.019  0.118  0.194
3--5    0.3590    0.017  0.326  0.391
4--5    0.1216    0.019  0.083  0.159
---
```

Note that we set `analytic = F` and `summarize = T` (which are the default settings). When the region of practical equivalence is desired, e.g., `rope = 0.1`, the output also includes the posterior probability in and out of the rope. Further, the `prob` argument is the decision rule for determining not only E but also partial correlations that are practically zero.

In GGMs, it is commonplace to estimate many parameters. For example, with 20 variables, this results in 190 partial correlations in total. This limits the usefulness of tables, and instead visualizations are typically relied upon. **BGGM** includes several options for plotting, each of which is implemented as a `S3` generic. For example, the partial correlations and their respective uncertainty can be plotted with:

```

R> Y <- BGGM::bfi[,1:10]
R> fit_sampling <- estimate(Y, analytic = F)
R> plot_1A <- plot(fit_sampling, ci_width = 0.95,
+                 width = 0.1,
+                 size = 2) +
+   coord_cartesian() +
+   theme(axis.text.x = element_text(angle = 90)) +
+   ggtitle("A") +
+   theme(plot.title = element_text(size = 25))
```

This is displayed in Figure 1 (panel A). Further, as seen with `ggtitle(.)`, the plot can be

Figure 1: Example plots implemented with the S3 generic `plot(.)`. (A) Partial correlations (edges) and their uncertainty (y -axis). The x -axis is the specific relation—e.g., 2-1 is the edge between variables 1 and 2. (B) An example heatmap of the selected edges. The gradient color corresponds to edge strength (effect size). (C) An example “network” plot using interval exclusion of zero to determine the graphical structure. (D) An example “network” plot using the *region of practical equivalence* to determine the graphical structure. In contrast to (C), the smaller edges are not included in the graph. (E) An example “network” plot using the *region of practical equivalence* for determining conditional independence (i.e., evidence for practically null values).

further customized with **ggplot2** (Wickham 2016). This is possible with all the plotting functions (i.e., `plot(.)`). **BGGM** also includes two options for visualizing E . A heatmap is plotted with:

```
R> E <- select(fit_sampling, ci_width = 0.95)
R> plot_1B <- plot(E, type = "heatmap", lower_tri = TRUE) +
+   ggtitle("B") +
+   theme(plot.title = element_text(size = 25))
```

Here `lower_tri = TRUE` controls which partial correlations are plotted. In this case, only the lower triangular elements are included in the plot. This can be changed with `lower_tri = FALSE`. This plot (`type = "heatmap"`) is displayed in Figure 1 (panel B). The colors correspond to the direction of the partial correlations. The transparency and line width reflect edge strength (effect size). On the other hand, a “network” plot can be obtained with:

```
R> plot_1C <- plot(E, type = "network",
+   layout = 'circle',
+   node_outer = 10,
+   node_inner = 9,
+   node_text_size = 6) +
+   ggtitle("C") +
+   theme(plot.title = element_text(size = 25))
```

This is also displayed in Figure 1 (panel C).

A key feature of **BGGM** is extending inference beyond identifying non-zero partial correlations. The region of practical equivalence can be used for this purpose, as it allows for determining which relations are practically zero. Of course, the user must define the null region. In this case, we follow Cohen’s guidelines (Cohen and Cohen 1983), wherein a standardized effect that is smaller than 0.1 in the absolute sense is considered to be negligible. This is implemented with:

```
R> Y <- BGGM::bfi[,1:10]
R> fit_sample <- estimate(Y, samples = 5000, analytic = F)
R> E <- select(fit_sample, rope = 0.1, prob = 0.95)
R> head(E, nrow = 10, summarize = T, digits = 2)
```

BGGM: Bayesian Gaussian Graphical Models

Type: Selected Graph (Sampling)

Probability: 0.95

Region of Practical Equivalence: [-0.1, 0.1]

```

Connectivity: 60 %
---
Call:
select.estimate(x = fit_sample, rope = 0.1, prob = 0.95)
---
Estimates:

  egde post_mean post_sd pr_out pr_in
1--2   -0.2401   0.018   1.00 0.0000
1--3   -0.1077   0.019   0.65 0.3506
2--3    0.2862   0.017   1.00 0.0000
1--4   -0.0075   0.019   0.00 1.0000
2--4    0.1645   0.019   1.00 0.0002
3--4    0.1779   0.019   1.00 0.0000
1--5   -0.0091   0.019   0.00 1.0000
2--5    0.1560   0.019   1.00 0.0006
3--5    0.3590   0.017   1.00 0.0000
4--5    0.1213   0.019   0.86 0.1368
---

```

The argument `prob = 0.95` requires that 95 % of the posterior density be in or out the rope to be considered practically or different from zero. With this decision rule, as seen with the `head(.)`, edges 1--4 and 1--5 are practically equivalent to zero. This inference is made possible with **BGGM**.

In this case, `plot(.)` returns two objects: (1) the selected edges with (Equation 13); (2) those for which there is support for the null values (Equation 12). These plots are displayed in Figure 1 (panels D and E). This is implemented with:

```

R> plts <- plot(E, type = "network",
+   layout = 'circle',
+   node_outer = 10,
+   node_inner = 9,
+   node_text_size = 6)
R> plot_1D <- plts$plot_nonzero +
+   ggtitle("D") +
+   theme(plot.title = element_text(size = 25))
R> plot_1E <- plts$plot_zero +
+   ggtitle("E") +
+   theme(plot.title = element_text(size = 25))

```

We emphasize that GGMs are often thought to capture conditionally *independent* relations—i.e., evidence for the null hypothesis of no effect, conditional on the other variables in the model. However, the dominant approach assesses conditional *dependence* ($\rho_{ij} \neq 0$), and

then sets the off-diagonal elements to zero otherwise. **BGGM**, as seen in panel E, can explicitly answer the question of conditional independence either with the rope, or as we show below (Section 4), with the Bayes factor.

For the remainder of this work, we focus on the function calls and corresponding visualizations. The more lengthy `summary(.)` outputs are omitted, and instead we use `head(.)` where applicable. However all functions in **BGGM** are implemented as S3 generics, and each includes `summary(.)`, `print(.)`, and `plot(.)`.

3.3. Edge differences

Differences between partial correlations are often tested in GGMs; for example, with a classical (i.e., frequentist) approach that is implemented in **bootnet**. One contribution of **BGGM** is providing Bayesian analogs for commonly used methods, as well as extensions to those methods. In this case, our package uses posterior probabilities to determine which edges are practically *equivalent*. This is implemented with:

```
R> Y <- BGGM::bfi[,1:10]
R> fit_sample <- estimate(Y, analytic = F)
R> edge_difference <- edge_compare(fit_sample,
+                               contrast = list("1--5 - 1--3",
+                                               "1--2 - 1--6",
+                                               "1--4 - 1--7",
+                                               "1--5 - 1--10",
+                                               "1--2 - 1--9"),
+                               ci_width = 0.95,
+                               rope = 0.1)
R> head(edge_difference, nrow = 4)
BGGM: Bayesian Gaussian Graphical Models
---
Type: Edge comparison(s)
Credible Interval: 95 %
Region of Practical Equivalence: [-0.1, 0.1]
---
Call:
edge_compare.estimate(x = fit_sample, contrast = list("1--5 - 1--3",
  "1--2 - 1--6", "1--4 - 1--7", "1--5 - 1--10",
  "1--2 - 1--9"), ci_width = 0.95, rope = 0.1)
---
Estimates:
```

contrast	post_mean	post_sd	pr_out	pr_in
1--5 - 1--3	0.0893	0.0325	0.372	0.628
1--2 - 1--6	-0.2968	0.0271	1	0
1--4 - 1--7	-0.0873	0.0295	0.3356	0.6644

```

1--5 - 1--10    0.0094  0.0268  6e-04  0.9994
---
```

This output includes the posterior mean and standard deviation for each difference. Further, `pr_out` is the proportion of samples included between ± 0.1 . This can be interpreted as the posterior probability of practical equivalence, which has been defined with the argument `rope = 0.1`. Note that care must be taken when specifying the contrasts, as an error will arise if they are not in the proper format. The object `edge_difference` can then be plotted, but this is omitted to keep the presentation concise.

3.4. Prediction

The following is based on the correspondence between the elements of Θ and multiple regression. Detailed proofs are provided in [Stephens \(1998\)](#), and further information can be found in [Kwan \(2014\)](#). In the context of GGMs, using regression to select edges is referred to as “neighborhood” selection ([Meinshausen and Bühlmann 2006](#)). This is implemented in `glasso`, `huge`, and `mgm`. On the other hand, the method described in [Williams \(2018\)](#) work directly with either the posterior distribution (Equation 10) or the maximum a posteriori (Equation 6).

The technique can be described with traditional regression notation, in which p multiple regression models are fitted. First let each node V_i be defined as $\mathbf{y}$ —i.e., the scores of the n subjects on the j th variable/node. Each node is regressed on the remaining $p - 1$ variables, which estimates the potential “neighborhood” for each variable

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta}_j + \boldsymbol{\varepsilon}, \quad (14)$$

where $\boldsymbol{\varepsilon}$ is an n -dimensional vector, with the mean as a vector of zeroes, and the covariance matrix as $\sigma_j^2 \mathbf{I}_n$. Here $\mathbf{X}$ is a $n \times (p - 1)$ design matrix, that excludes the i th node, and $\boldsymbol{\beta}$ is $(p - 1) \times 1$ vector. The intercept is excluded, due to standardizing the data, so $\boldsymbol{\beta}^{(j)}$ contains $p - 1$ regression coefficients. To be clear, $\boldsymbol{\beta}_j$ denotes the vector of coefficients for the j th regression model, where the individual elements are defined as β_{ij} . The residuals are assumed to follow $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_j^2)$, where σ_j^2 is the residual variance for the j th node. The regression coefficients and error variances then correspond to the off-diagonal and diagonal elements of Θ —i.e.,

$$\theta_{ij} = \frac{-\beta_{ij}}{\sigma_j^2} \text{ and } \theta_{jj} = \frac{1}{\sigma_j^2}, \quad (15)$$

where θ_{ij} denotes the covariance corresponding to i th row and j th column of Θ . The diagonal of Θ is then denoted θ_{jj} . Consequently, for all posterior samples, $s = 1, \dots, S$,

$$\beta_{ij}^{(s)} = \frac{-\theta_{ij}^{(s)}}{\theta_{jj}^{(s)}} \text{ and } \sigma_j^{2(s)} = \frac{1}{\theta_{jj}^{(s)}} \quad (16)$$

results in the posterior distribution for each regression coefficient and residual variance. It follows that **BGGM** can also be used for the purpose of Bayesian multiple regression—i.e.,

```
R> Y <- BGGM::bfi[,1:10]
R> fit <- estimate(Y, samples = 5000)
R> coefficients(fit, node = 1, ci_width = 0.95)

BGGM: Bayesian Gaussian Graphical Models
---
Type: Inverse to Regression
---
Call:
BGGM::beta_summary(x = fit, node = node, ci_width = ci_width,
  samples = samples)
---
Estimates:
```

node	post_mean	post_sd	2.5%	97.5%
2	-0.278	0.022	-0.322	-0.237
3	-0.125	0.023	-0.173	-0.078
4	-0.015	0.021	-0.054	0.027
5	-0.017	0.022	-0.060	0.024
6	0.057	0.021	0.016	0.099
7	0.080	0.022	0.038	0.123
8	0.045	0.020	0.006	0.086
9	0.142	0.022	0.101	0.184
10	-0.028	0.021	-0.069	0.010

```
---
```

Here `node = 1` indicates which node is summarized. This correspondence allows for computing measures of prediction, including Bayesian R^2 and Bayesian leave-one-out cross validation, each of which has a measure of uncertainty. Furthermore, when a computationally option is desirable, **BGGM** includes an analytic expression for prediction error based on (6). This is also known as the predicted residual sums of squares (PRESS; Allen 1971).

Bayesian R^2

In the following we include 10 variables in the model. In-sample Bayesian R^2 is implemented with:

```
R> Y_train <- BGGM::bfi[1:100,1:10]
R> fit_train <- estimate(Y_train, samples = 5000)
R> train_R2 <- predict(fit_train, ci_width = 0.95,
```

Figure 2: Example “network” predictability plots implemented with the `S3` generic `plot(.)`. The nodes are ordered by which had the best performance. (A) Bayesian R^2 for training and test data sets. The edges are first selected in the former, and then used to predict the latter. (B) Bayesian (approximate) leave-one-out cross-validation. (C) PRESS leave-one-out cross-validation. The order is the same as (B), but only includes a point estimate. These plots highlight a key aspect of the Bayesian predictive methods in **BGGM**, in that they include a measure of uncertainty.

```

+                               samples = 1000, measure = "R2")
R> head(train_R2, nrow = 2)

BGGM: Bayesian Gaussian Graphical Models
---
Type: In-sample predictive accuracy
Measure: Variance Explained (R2)
---
Call:
predict.estimate(fit = fit_train, ci_width = 0.9, samples = 1000,
  measure = "R2")
---
Estimates:

  node post_mean  post_sd      5%      95%
    1 0.1132451 0.0856017 0.002985273 0.2804324
    2 0.2425554 0.1120839 0.053173789 0.4184764
---
```

Here `ci_width = 0.95` indicates the decision rule for setting coefficients to zero. Similarly, out-of-sample Bayesian R^2 is computed with:

```
R> Y_test <- BGGM::bfi[101:2000,1:10]

R> test_R2 <- predict(fit_train, ci_width = 0.90,
+                   test_data = Y_test,
+                   samples = 1000, measure = "R2")
```

The work flow is completed by visualizing Bayesian R^2 for each node—i.e.,

```
R> plot_2A <- plot(train_R2, test_R2, order = "test")
```

Here the nodes have been ordered (`order = "test"`) by which has the best out-of-sample performance (Figure 2; panel A). **BGGM** also includes the mean squared error (`measure = "mse"`).

Leave-one-out cross-validation

Bayesian leave-one-out cross-validation can be executed as follows:

```
R> Y <- BGGM::bfi[1:1000,1:10]
R> fit_sample <- estimate(Y, samples = 5000)
R> bayes_loo <- loocv(fit_sample)
R> plt_2B <- plot(bayes_loo) +
+             theme_classic() +
+             ylab("Bayesian Leave-One-Out")
```

Similarly, by setting `analytic = T`, leave-one-out prediction error can be computed analytically. This is implemented with:

```
R> Y <- BGGM::bfi[1:1000,1:10]
R> fit_analytic <- estimate(Y, analytic = T)
R> press_loo <- loocv(fit_analytic)
R> plt_2C <- plot(press_loo) +
+             theme_classic() +
+             ylab("PRESS: Leave-One-Out") +
+             scale_y_continuous(expand = c(0, 0),
+             limit = c(0, 1000))
```

The plots for Bayesian R^2 , Bayesian leave-one-out, and the PRESS version are displayed in Figure 2. This highlights the difference between the leave-one-out methods, in that the Bayesian version has a measure of uncertainty (although the order is the same).

4. Hypothesis testing

The following methods were introduced in Williams and Mulder (2019). That work not only presented an exploratory approach using the Bayes factor, but it also proposed

methodology for confirmatory hypothesis testing in GGMs. The latter provides an alternative to data driven model selection that is commonplace in the GGM literature, and in particular, it allows for comparing theoretical models. In particular, **BGGM** allows for testing competing sets of order or equality constraints on multiple partial correlations in a GGM using the novel matrix- F prior (Mulder and Raúl Pericchi 2018). The following tests cannot be executed with **BDgraph** or **beam**.

4.1. Matrix- F distribution

The matrix- F was recently proposed as a flexible alternative to the inverse Wishart and Wishart prior for covariance and precision matrices, respectively (Mulder and Raúl Pericchi 2018). It overcomes one limitation of the Wishart prior (Section 3.1). Namely, as shown in Williams and Mulder (2019), the prior standard deviation is necessarily a function of the network size (i.e., $p^{-1/2}$). This limits its usefulness for Bayesian hypothesis testing in particular. We thus employ an encompassing matrix- F prior distribution for the precision matrix,

$$\Theta \sim F(\nu, \delta, \mathbf{B}), \quad (17)$$

where $\nu > p - 1$ and $\delta > 0$ are the first and second degrees of freedom, which control the behavior near the origin and in the tails, respectively, and $\mathbf{B}$ is a positive definite scale matrix. For completeness the prior density of the matrix- F prior is given in Appendix A.1 and further details about the encompassing prior approach for hypothesis testing can be found in Klugkist *et al.* (2005). The matrix- F prior can be written as a scale mixture of Wishart distributions with an inverse Wishart mixture distribution, i.e.,

$$\begin{aligned} \Theta | \Psi &\sim W(\nu, \Psi) \\ \Psi &\sim IW(\delta + p - 1, \mathbf{B}). \end{aligned} \quad (18)$$

Because the Wishart prior is conjugate, it follows that the matrix- F prior is conditionally conjugate. That is, the conditional posterior of Θ given Ψ has a Wishart distribution and the conditional posterior of Ψ given Θ has an inverse Wishart distribution (Appendix A.1). This makes the matrix- F prior computationally feasible for GGMs, in that the posterior can be obtained with a Gibbs sampler (Appendix A.1).

The hypotheses of interest are not directly formulated on Θ , but on the partial correlations ρ given in (2). To understand the *implied* marginal prior for ρ_{ij} , consider the fact that the matrix- F prior can *also* be written as a scale mixture of inverse Wishart distributions with a Wishart mixture distribution, i.e.,

$$\begin{aligned} \Theta | \Phi &\sim IW(\delta + p - 1, \Phi) \\ \Phi &\sim W(\nu, \mathbf{B}). \end{aligned} \quad (19)$$

Furthermore, due to Barnard *et al.* (2000) it is known that a covariance matrix having an inverse Wishart prior distribution with an identity scale matrix, i.e., $IW(\nu, \mathbf{I}_p)$, results in marginal priors for the bivariate correlations having $\text{beta}(\frac{\nu}{2}, \frac{\nu}{2})$ distributions in the interval

$(-1, 1)$. Consequently, if a precision matrix has an inverse Wishart prior distribution, i.e., $\Theta \sim IW(\delta + p - 1, \mathbf{I}_p)$, the partial correlations then follow a $beta(\frac{\delta}{2}, \frac{\delta}{2})$ distribution in the interval $(-1, 1)$, which is invariant to the dimension of the network (p). We therefore set $\mathbf{B} = \epsilon \mathbf{I}_p$ and $\nu = \epsilon^{-1}$, for a small value for ϵ (e.g., 0.001), so that $\Phi \approx \mathbf{I}_p$ and Θ is approximately distributed as $IW(\delta + p - 1, \mathbf{I}_p)$. Thus, instead of using the conjugate Wishart prior, which is too informative for large networks, we use the matrix $\mathbf{F}$ prior to construct an approximate inverse Wishart prior for the precision matrix, which is conditionally conjugate and appropriately controls the prior information for partial correlations for networks of any size.

Together, the prior for the precision matrix and the implied marginal prior for the partial correlations are specified as

$$\begin{aligned} \Theta &\sim F(\epsilon^{-1}, \delta, \epsilon \mathbf{I}_p) \\ \rho_{ij} &\sim beta(\frac{\delta}{2}, \frac{\delta}{2}) \text{ on } (-1, 1), \end{aligned} \tag{20}$$

for $i \neq j = 1, \dots, p$, respectively. The prior hyperparameter δ can be chosen such that the prior standard deviation corresponds with the expected deviation from zero in the case of a partial correlation would be unequal to zero. Because the prior standard equals $s_\rho = 1/\sqrt{\delta + 1}$, which is the standard deviation of a beta distribution, one can set the hyperparameter equal to $\delta = (s_\rho^2)^{-1} - 1$ by plugging in the anticipated deviation from zero of the partial correlations for s_ρ .

4.2. Structure learning

In **BGGM**, conditional relationships, and thus E , are determined by comparing a precise null model versus an unconstrained model-i.e.,

$$\begin{aligned} \mathcal{H}_0 &: \rho_{ij} = 0 \\ \mathcal{H}_u &: \rho_{ij} \neq 0 \end{aligned} \tag{21}$$

These hypotheses are nested in that $\mathcal{M}_u \subseteq \mathcal{M}_1$. This simplifies computing the Bayes factor, because it allows for using the Savage-Dickey ratio ([Dickey 1971](#); [Wagenmakers, Lodewyckx, Kuriyal, and Grasman 2010](#)). The hypothesis test in favor of the null hypothesis can be formulated as

$$\begin{aligned} BF_{0u} &= \frac{p(\mathbf{Y}|\mathcal{H}_0)}{p(\mathbf{Y}|\mathcal{H}_u)}, \\ &= \frac{\int f_0(\mathbf{Y}|\Sigma_0, \mathcal{H}_0) \pi_0(\Sigma_0|\mathcal{H}_0) d\Sigma_0}{\int f_u(\mathbf{Y}|\Sigma_u, \mathcal{H}_u) \pi_u(\Sigma_u|\mathcal{H}_u) d\Sigma_u}, \end{aligned} \tag{22}$$

where $f_i(\mathbf{Y}|\Sigma_i, \mathcal{H}_i)$ and $\pi_i(\Sigma_i|\mathcal{H}_i)$ are the likelihood functions and the prior of the parameters under each hypothesis $i \in \{0, 1\}$, in addition to Σ_0 denoting a covariance

matrix with $\theta_{ij} = 0$. Further, because the present approach works directly with the partial correlations, ρ_{ij} , and the corresponding *implied* prior distribution, this simplifies (22)–i.e.,

$$BF_{0u} = \frac{p(\mathbf{Y}|\mathcal{H}_0)}{p(\mathbf{Y}|\mathcal{H}_1)} = \frac{p(\rho_{ij} = 0|\mathbf{Y}, \mathcal{H}_1)}{p(\rho_{ij} = 0|\mathcal{H}_1)}, \quad (23)$$

where the off-diagonal elements have been standardized following (2). The nuisance parameters, μ_i and μ_j , are integrated out (as seen in the Appendix A.1). In other words, by only considering $\mathcal{H}_u$ with respect to ρ_{ij} , the Bayes factor can be computed as the unconstrained posterior density of ρ_{ij} evaluated at zero divided by the prior density also evaluated at zero (Mulder, Hoijtink, and de Leeuw 2012). Consequently, the necessary ingredients for computing (23) are posterior samples drawn from the full model with Markov chain Monte Carlo simulation. In **BGGM**, these samples are obtained with the computationally efficient Gibbs sampler described in the Appendix (A.1).

Visualizing scientific expectations

For Bayesian hypothesis testing in particular, it is important to *fully* understand the prior distribution $\mathcal{H}_u$. This is because it captures the predicted effect size, and it is used to compute the Bayes factor. This stands in contrast to the above estimation based methods, where E is determined with respect to the posterior distribution. To this end, **BGGM** includes functions to visualize the prior distribution–i.e.,

```
R> rho_sd <- c(0.1, 0.25, 0.5)
R> plt_3A <- hypothesis_plot(rho_sd = rho_sd) +
+       theme(panel.grid.major = element_blank()) +
+       ylab("Density")
```

`rho_sd = c(.)` defines a couple (prospective) values for the hypothesized standard deviation of the partial correlations. Note this is also customizable with **ggplot2**. The package **BGGM** allows for testing edge differences with the Bayes factor (Section 4.3). Accordingly, the *implied* prior distribution for the difference between two edges can also be plotted. This is implemented with:

```
R> rho_sd <- c(0.1, 0.25, 0.5)
R> plt_3B <- hypothesis_plot(rho_sd = rho_sd,
+       difference = T) +
+       theme(panel.grid.major = element_blank()) +
+       ylab("Density")
```

Note `difference = T`. The plots are displayed in Figure 3 (panels A and B).

Two-sided testing

The Bayes factor in (23) provides evidence for a two-sided hypothesis test. This is implemented with:

Figure 3: Example plots implemented with `hypothesis_plot()`. (A) Prior distributions for various standard deviations of the beta distribution (± 1). Larger values approach a uniform distribution, whereas smaller values increasingly concentrate around zero. (B) Implied prior distribution for edge differences. (C) Posterior and prior distribution plot. An estimated edge can be visualized with the chosen prior distribution. The points denote the heights of the densities at zero.

```
R> Y <- BGGM::bfi[,1:5]
R> fit_bf <- explore(Y, prior_sd = 0.5,
+                   iter = 5000,
+                   cores = 2)
R> summary(fit_bf)
```

BGGM: Bayesian Gaussian Graphical Models

Type: Hypothesis Testing (Exploratory)

Posterior Samples: 5000

Observations (n): 2709

Variables (p): 5

Edges: 10

Delta: 3

Call:

`explore.default(X = Y, prior_sd = 0.5, iter = 5000, cores = 2)`

Date: Wed May 15 11:40:49 2019

Note `summary()`, or alternatively `print()`, provides information about the fitted model, including that hypothesis testing (exploratory) was used, the number of observations (n) and variables (p), and the number of edges. Delta (δ) is the hyperparameter of

the matrix- F distribution. A value of 3 corresponds to `prior_sd = 0.5`. This output parallels the estimation based methods. Indeed, in **BGGM** all fitted objects include specific (what method was used) and general information (e.g., n and p) when printed.

The edge set E is then selected with:

```
R> E <- select(fit_bf,
+             BF_cut = 3,
+             alternative = "two.sided")
```

`BF_cut = 3` is the threshold for selecting non-zero and zero relations. `alternative = "two.sided"` specifies the two-sided hypothesis test. There are a variety of ways to summarize `E` and one option is with:

```
R> summary(E, summarize = T, log = T, digits = 2)
```

BGGM: Bayesian Gaussian Graphical Models

Type: Hypothesis Testing

Alternative: two.sided

Bayes Factor: 3

Call:

```
select.explore(x = fit_bf, BF_cut = 3, alternative = "two.sided")
```

Hypotheses:

$H_0: \rho = 0$

$H_1: \rho \neq 0$

Estimates:

edge	post_mean	post_sd	BF	10
1--2	-0.2405	0.018	76.4	
1--3	-0.1077	0.019	11.7	
2--3	0.2861	0.017	115.9	
1--4	-0.0075	0.019	-3.5	
2--4	0.1643	0.018	34.7	
3--4	0.1774	0.018	41.7	
1--5	-0.0086	0.019	-3.5	
2--5	0.1562	0.018	30.9	
3--5	0.3589	0.017	183.7	
4--5	0.1216	0.019	17.1	

note: `BF_10` is evidence in favor of H_1

The option `log = TRUE` controls the scale of the Bayes factor. In the case of large values, it is especially advantageous to summarize the Bayes factor on the logarithmic scale. Note that the same plotting options available for these methods; for example, those seen Figure 1 (Panels B, C, and E). They have the same implementation as the estimation based methods, and thus not shown here.

It is also possible to plot $\mathcal{H}_u$ and a selected posterior distribution. This visualizes how the Bayes factor is computed—i.e.,

```
R> plt_3C <- hypothesis_plot(fit = fit_bf,
+                             edge = "1--4", size = 3) +
+                             theme(legend.position = "top",
+                             panel.grid.minor = element_blank()) +
+ ylab("Density") +
+ xlab("")
```

This is displayed in Figure 3 (panel C). Here it can be seen that the Bayes factor is the ratio of density evaluated at 0. We anticipate this plotting option will be most useful, in applied settings, when explaining the Bayes factor approach for selecting the graph (or more generally partial correlations).

One-sided testing

The Bayesian approaches implemented in **BDgraph** and **beam** are effectively two-sided tests. However, in the case of psychopathology (for example), a researcher might expect that all relations are in the same direction. **BGGM** allows for directional hypothesis tests for determining E . To compute the one-sided Bayes factor, $\mathcal{H}_1 : \rho_{ij} > 0$, evidence between the unrestricted and one-sided is first assessed (BF_{1u}) as

$$BF_{1u} = \frac{\int_0^\infty p(\rho_{ij}|\mathbf{Y}, \mathcal{H}_u) d\rho_{ij}}{\int_0^\infty p(\rho_{ij}|\mathcal{H}_u) d\rho_{ij}} = \frac{\Pr(\rho_{ij} > 0|\mathbf{Y}, \mathcal{H}_u)}{\Pr(\rho_{ij} > 0|\mathcal{H}_u)}, \quad (24)$$

where the numerator and denominator denote the posterior and prior probability that ρ_{ij} is greater than zero under the unrestricted model $\mathcal{H}_u$. The (relative) evidence for $\mathcal{H}_1$ against $\mathcal{H}_0$ is then given by $BF_{10} = BF_{u0} \cdot BF_{1u}$. The former is the inverse of BF_{0u} that is given in (23). The Bayes factor for a negative effect can be obtained in a similar fashion. This is implemented with:

```
R> E <- select(fit_bf,
+              BF_cut = 3,
+              alternative = "greater")
R> summary(E)
```

BGGM: Bayesian Gaussian Graphical Models

Type: Hypothesis Testing

Alternative: greater

Bayes Factor: 3

Connectivity: 60 %

Call:

```
select.explore(x = fit_bf, BF_cut = 3, alternative = "greater")
```

Hypothesis:

H1: $\rho > 0$

Partial Correlations

	1	2	3	4	5
1	0	0.0000000	0.0000000	0.0000000	0.0000000
2	0	0.0000000	0.2861356	0.1643470	0.1561560
3	0	0.2861356	0.0000000	0.1773919	0.3589186
4	0	0.1643470	0.1773919	0.0000000	0.1216472
5	0	0.1561560	0.3589186	0.1216472	0.0000000

Adjacency (positive)

	1	2	3	4	5
1	0	0	0	0	0
2	0	0	1	1	1
3	0	1	0	1	1
4	0	1	1	0	1
5	0	1	1	1	0

Note all the included edges are now positive. Consequently, when using `plot(E,...)`, the connections will all be green. $\mathcal{H}_1 : \rho_{ij}$ can be tested by changing `alternative = greater` to `alternative = less`.

Exhaustive testing

A defining feature of Bayesian hypothesis testing is the ability to assess which theoretical model best predicts the data at hand (Lee 2018; Kass and Raftery 1995). However, in the absence of guiding theory, it is likely that a more exploratory approach is warranted. **BGGM** thus includes an exhaustive approach—i.e.,

$$\begin{aligned}\mathcal{H}_0 : \rho_{ij} &= 0 \\ \mathcal{H}_1 : \rho_{ij} &> 0 \\ \mathcal{H}_2 : \rho_{ij} &< 0,\end{aligned}\tag{25}$$

which covers the entire parameter space. The evidence for each can then be computed from the inverse of the Bayes factor given in (23)—i.e., BF_{u0} (evidence for the unrestricted model). The computation of each one-sided Bayes factor for $\mathcal{H}_1$ and $\mathcal{H}_2$ (BF_{10}) is given in (24). With the three Bayes factors in hand $BF_{tu}, t = 0, \dots, 2$, each compared to the unrestricted model $\mathcal{M}_u$, the posterior hypothesis probabilities can then be computed. For example, the probability of $\mathcal{H}_1$ is

$$p(\mathcal{H}_1|\mathbf{Y}) = \frac{p(\mathcal{H}_1)BF_{1u}}{p(\mathcal{H}_0)BF_{0u} + p(\mathcal{H}_1)BF_{1u} + p(\mathcal{H}_2)BF_{2u}}, \quad (26)$$

where $p(\cdot)$ is the prior probability for each hypothesis. **BGGM** assumes equal prior probabilities for the hypotheses. This results in the posterior probability for each hypothesis. In other words, this provides an *exhaustive* assessment of the previously stated hypotheses. The ratio of these probabilities, for example $p(\mathcal{H}_1|\mathbf{Y})/p(\mathcal{H}_2|\mathbf{Y})$, then provide the Bayes factor for which direction is most supported by the observed data. This is implemented with:

```
R> Y <- BGGM::bfi[,1:10]
R> fit_bf <- explore(Y, prior_sd = 0.5,
+                   iter = 5000,
+                   cores = 2)
R> E <- select(fit_bf,
+             hyp_prob = 0.90,
+             alternative = "exhaustive",
+             digits = 3)
R> head(E, summarize = T, nrow = 5)

BGGM: Bayesian Gaussian Graphical Models
---
Type: Hypothesis Testing
Alternative: exhaustive
Posterior probability: 0.9
---
Call:
select.explore(x = fit_bf, alternative = "exhaustive",
              hyp_prob = 0.9)
---
Hypotheses:
H0: rho = 0
H1: rho > 0
H2: rho < 0
---
Estimates:
```

Figure 4: Example “network” plots implemented with the S3 generic `plot(.)`. The graphical structures were determined with posterior hypothesis probabilities. $p(\mathcal{H}_0|\mathbf{Y})$ includes connections that had evidence the null hypothesis of conditional independence ($\rho_{ij} = 0$). $p(\mathcal{H}_1|\mathbf{Y})$ includes edges that had evidence for a positive effect $\rho_{ij} > 0$. $p(\mathcal{H}_2|\mathbf{Y})$ includes edges that had evidence for a negative effect $\rho_{ij} < 0$.

edge	post_mean	post_sd	$p(H_0 Y)$	$p(H_1 Y)$	$p(H_2 Y)$
1--2	-0.24047	0.0183	0.000	0.000	1.000
1--3	-0.10775	0.0193	0.000	0.000	1.000
2--3	0.28614	0.0175	0.000	1.000	0.000
1--4	-0.00751	0.0195	0.998	0.001	0.001
2--4	0.16435	0.0184	0.000	1.000	0.000

`hyp_prob = 0.90` is the posterior probability for concluding there is evidence for a particular hypothesis. The object `E` includes selected matrices for each hypothesis. These can then be visualized—i.e.,

```
R> plts <- plot(E, type = "network")
R> plot_H0 <- plts$plot_H0
R> plot_H1 <- plts$plot_H1
R> plot_H2 <- plts$plot_H2
```

The `type = 'network'` plots are displayed in Figure 4.

4.3. Edge differences

BGGM allows for testing edge differences with Bayesian hypothesis testing. The same options are available as for the function `select(.)`—i.e., `alternative = "two.sided"`,

alternative = "greater" (or "less"), and alternative = "exhaustive". The exhaustive approach is implemented with:

```
R> Y <- BGGM::bfi[,1:10]
R> fit_bf <- explore(Y, prior_sd = 0.5,
+                   iter = 5000,
+                   cores = 2)
R> edge_comp <- edge_compare(fit_bf,
+                             contrast = c("1--5 - 1--3",
+                                           "1--2 - 1--6",
+                                           "1--4 - 1--7",
+                                           "1--5 - 1--10",
+                                           "1--2 - 1--9"),
+                             alternative = "exhaustive")
R> head(x,nrow = 2, digits = 2)
BGGM: Bayesian Gaussian Graphical Models
---
Type: Edge comparison (Hypothesis Testing)
Alternative: exhaustive
---
Call:
edge_compare.explore(x = fit_bf, contrast = c("1--5 - 1--3",
"1--2 - 1--6", "1--4 - 1--7", "1--5 - 1--10",
"1--2 - 1--9"), alternative = "exhaustive")
---
Hypotheses:
H0: rho = 0
H1: rho > 0
H2: rho < 0
---
Estimates:
```

contrast	post_mean	post_sd	p(H0 Y)	p(H1 Y)	p(H2 Y)
1--5 - 1--3	0.09	0.032	0.13	0.87	0.002
1--2 - 1--6	-0.30	0.027	0.00	0.00	1.000

```
---
```

There are two plotting options for the returned object `edge_comp`. The first is ideal if few contrasts are tested, whereas the second can be used to visualize several contrasts in the same plot. These are implemented with:

```
R> plt <- plot(edge_comp, stack = F, spread = .75)
R> # one contrast
```

Figure 5: Example plots implemented with the **S3** generic `plot(.)`. This particular example looks at edge differences within the same GGM. Note that all functions in **BGGM** that provide evidence for competing hypotheses can be plotted. (A) The hypothesis probabilities for one contrast or edge difference 1-5 - 1-3. (B) When several contrasts are tested, a stacked bar graph can be used to visualize the hypothesis probabilities in the same figure.

```
R> plot_6A <- plt[[1]]
R> plot_6B <- plot(edge_comp, stack = T, spread = .75)
```

These are displayed in Figure 4.

4.4. Confirmatory hypothesis testing

Order constraints

Theory may suggest, for example, that a set of partial correlations is expected to be larger than another set of partial correlations, that there is a hypothesized order of edge magnitude for a given node (variable), or some edges are expected to be positive while others are predicted to be negative. These can be tested with **BGGM**. For example, consider the following hypotheses,

$$\begin{aligned} \mathcal{H}_1 : \rho_{1,2} > \rho_{1,3} > \rho_{1,4} \\ \mathcal{H}_c : \text{"not } \mathcal{H}_1 \text{"}, \end{aligned} \quad (27)$$

which can be expressed as

$$\mathbf{R}_1 \boldsymbol{\rho} > \mathbf{r}_1 \Leftrightarrow \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \rho_{1,2} \\ \rho_{1,3} \\ \rho_{1,4} \end{bmatrix} > \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Here $[\mathbf{R}_1|\mathbf{r}_1]$ is a matrix that specifies the order constraints under $\mathcal{H}_1$. The subspace under a constrained hypothesis $\mathcal{H}_1$ that satisfies the constraints on $\boldsymbol{\rho}$ will be denoted by Ω_1 . Because we use the encompassing prior approach, marginal likelihood computation can be avoided because the Bayes factor of each constrained hypothesis against the unconstrained hypothesis H_u can be expressed as the ratio of the posterior and prior probability that the order constraints of $\mathcal{H}_1$ hold under the larger, unconstrained model—i.e.,

$$B_{1u} = \frac{\Pr(\boldsymbol{\rho} \in \Omega_1 | \mathbf{Y}, \mathcal{H}_u)}{\Pr(\boldsymbol{\rho} \in \Omega_1 | \mathcal{H}_u)}. \quad (28)$$

The Bayes factors between the competing hypotheses, that are given in (27), can then be computed using the transitivity property of the Bayes factor, e.g., $BF_{1c} = \frac{BF_{1u}}{BF_{cu}}$. This relatively simple example generalizes to multiple hypotheses, e.g., $\mathcal{H}_t$, for $t = 1, \dots, T$, as well as equality constraints. The multivariate densities are computed with **mvtnorm**. Further details can be found in [Mulder \(2016\)](#) and [Williams and Mulder \(2019\)](#).

In **BGGM**, an order constrained hypothesis can be tested with:

```
R> hypothesis <- c("(1--2, 1--3, 1--4) < 0 < (1--6)")
R> test_order <- confirm(x = Y, hypothesis = hypothesis,
+                        prior_sd = 0.5, iter = 50000,
+                        cores = 2)
R> summary(test_order)
BGGM: Bayesian Gaussian Graphical Models
---
Type: Confirmatory Hypothesis Testing
---
Call:
confirm.default(x = Y, hypothesis = hypothesis, prior_sd = 0.5,
               iter = 50000, cores = 2)
---
Hypotheses:

H1 (1--2,1--3,1--4)<0<(1--6)
Hc 'not H1'
---
Posterior prob:

p(H1|Y) = 0.9808
p(Hc|Y) = 0.0192
---
Bayes factor matrix:
      H1    Hc
H1  1.00 0.02
Hc 51.15 1.00
```

note: equal hypothesis prior probabilities

There is strong support for $\mathcal{H}_1$ ($BF_{1c} = 51.15$). Note that `iter = 50000` (≈ 7 seconds), which ensures an adequate number of samples for accurately computing the Bayes factor. This object can then be plotted (`plot(test_order)`), which produces a plot similar to Figure 4.

Equality constraints

BGGM can also test exact equality constraints on sets of edges (i.e., partial correlations). One could predict that certain edges have the same strength (effect size) or that a set of edges are equal to zero. The latter allows for testing expectations about the underlying conditional independence structure, say, for a particular variable (node) in the model. For example,

$$\mathcal{H}_1 : (\rho_{1,2}, \rho_{1,3}, \rho_{1,4}) = 0 \quad (29)$$

$$\mathcal{H}_c : \text{"not } \mathcal{H}_1\text{"}, \quad (30)$$

which states that edges each of those edges are equal to zero. This is implemented with:

```
R> hypothesis <- "(1--2, 1--3, 1--4) = 0"
R> test_order <- confirm(x = Y, hypothesis = hypothesis,
                        prior_sd = 0.5, iter = 50000,
                        cores = 2)
R> summary(test_order)
BGGM: Bayesian Gaussian Graphical Models
---
Type: Confirmatory Hypothesis Testing
---
Call:
confirm.default(x = Y, hypothesis = hypothesis, prior_sd = 0.5,
               iter = 50000, cores = 2)
---
Hypotheses:

H1 (1--2,1--3,1--4)=0
Hc      'not H1'
---
Posterior prob:

p(H1|Y) = 0
p(Hc|Y) = 1
```

```

---
Bayes factor matrix:
      H1      Hc
H1  1 3.144036e+56
Hc  0 1.000000e+00
---
note: equal hypothesis prior probabilities

```

In this case the hypothesis was not supported. This powerful framework can seamlessly be extended to any number of edges, as well as testing several hypotheses that can include a combination of both order restrictions and equality constraints. As an example, `hypothesis <- c("(1-10, 1-4) = 0; (1-10, 1-4) > 0")`, each of which would be tested against “Not $\mathcal{H}_1$ or $\mathcal{H}_2$.” Further, the ratio, $p(\mathcal{H}_1|\mathbf{Y})/p(\mathcal{H}_2|\mathbf{Y})$, results in BF_{12} .

5. Comparing GGMs

BGGM includes two novel approaches for comparing GGMs across any number of groups. They were introduced in [Williams, Rast, Pericchi, and Mulder \(2019b\)](#). The methods can test entire structures, particular edges, and also allow for nodewise testing. One approach is based on the posterior predictive distribution, and tests the null hypothesis of group equality. Importantly, like classical methods (i.e., frequentist), this approach can only determine whether the null of group equality is rejected or retained. It cannot provide evidence for the null hypothesis of equality. To this end, **BGGM** includes an additional approach that is based on Bayesian model selection. This allows for assessing network invariances (i.e., the null hypothesis). These tests cannot be executed using **BDgraph** or **beam**.

5.1. Posterior predictive KL divergence

We first assume that each group $g \in \{1, \dots, G\}$ is a realization from the same multivariate normal distribution—i.e., the null model

$$\mathcal{M}_0 : \Theta_1 = \dots = \Theta_G. \quad (31)$$

The posterior for the common precision matrix $\Theta (= \Theta_1 = \dots = \Theta_G)$, given the observed data, can be written as $p(\Theta | \mathbf{Y}_1^{obs}, \dots, \mathbf{Y}_G^{obs}, \mathcal{M}_0)$. Under $\mathcal{M}_0$, a posterior draw (s) for $\Theta^{(s)}$ is, in fact, a posterior draw for the precision matrix in all groups. These samples are efficiently drawn with (10).

Next replicated data set are generated given those precision matrices, i.e.,

$$\begin{aligned}\boldsymbol{\Theta}_1^{(s)} &\rightarrow \mathbf{Y}_1^{rep(s)} \\ &\vdots \\ \boldsymbol{\Theta}_G^{(s)} &\rightarrow \mathbf{Y}_G^{rep(s)}.\end{aligned}$$

These replicated data sets are generated with the observed group sizes. Now the posterior expectation of a precision matrix for group g given $\mathbf{Y}_g^{rep}$ can be approximated as

$$E\{\boldsymbol{\Theta}_g^{rep}|\mathbf{Y}_g^{rep}\} = (n_g - 1)(\mathbf{Y}_g^{repT}\mathbf{Y}_g^{rep})^{-1}. \quad (32)$$

With these in hand, it is then possible to compare groups with a chosen test-statistic. Currently, the only options in **BGGM** is Kullback–Leibler divergence (KLD) for multivariate normal distributions (41).

To summarize, this method follows these steps:

1. Estimate $p(\boldsymbol{\Theta}|\mathbf{Y}_1^{obs}, \dots, \mathbf{Y}_G^{obs}, M_0)$ with (10).
2. For each posterior sample (s)
 - (a) $\boldsymbol{\Theta}_g^{(s)} \rightarrow \mathbf{Y}_g^{rep(s)}$, for $g \in \{1, \dots, G\}$.
 - (b) Compute $\boldsymbol{\Theta}_g^{rep(s)}$ as $(n - 1) \mathbf{S}^{-1}$, where $\mathbf{S}$ is $\mathbf{Y}_g^{rep(s)T} \mathbf{Y}_g^{rep(s)}$, for $g \in \{1, \dots, G\}$.
 - (c) Compute the predictive distance:
 $\text{KLD}(E\{\boldsymbol{\Theta}_{g_1}^{rep(s)}|\mathbf{Y}_{g_1}^{rep(s)}\}, E\{\boldsymbol{\Theta}_{g_2}^{rep(s)}|\mathbf{Y}_{g_2}^{rep(s)}\})$.
3. Compute the observed distance:
 $\text{KLD}(E\{\boldsymbol{\Theta}_{g_1}^{obs}|\mathbf{Y}_{g_1}^{obs}\}, E\{\boldsymbol{\Theta}_{g_2}^{obs}|\mathbf{Y}_{g_2}^{obs}\})$.
4. Compute the posterior predictive p -value.

Note that KLD is symmetrized by taking the average of both directions. Further, g_1 and g_2 were used to keep the notation manageable. This approach allows for comparing any number of groups.

Global predictive check

In the case of 3 groups, with the same covariance structure, the predictive check is implemented with:

```
R> Y1 <- MASS::mvrnorm(500, rep(0, 20), Sigma = diag(20))
R> Y2 <- MASS::mvrnorm(500, rep(0, 20), Sigma = diag(20))
R> Y3 <- MASS::mvrnorm(500, rep(0, 20), Sigma = diag(20))
```

Figure 6: Example posterior predictive check plots implemented with the `S3` generic `plot(.)`. The densities correspond to the predictive distribution of Kullback-Leibler divergence (KLD). The red region corresponds to a user defined critical region (e.g., $\alpha = 0.05$). Thus, when the observed KLD (black points) is in the critical region, the assumed model of group equality is not warranted for the GGMs under consideration. (A) The “global” approach tests all aspects of the graphical structures (i.e., the entire precision matrix). (B) The “nodewise” approach allows for testing each node in the GGMs.


```

R> ggm_ppc <- ggm_compare_ppc(Y1, Y2, Y3,
+                             type = "global", iter = 1000)
R> summary(ggm_ppc)
BGGM: Bayesian Gaussian Graphical Models
---
Type: GGM Comparison (Global Predictive Check)
---
Call:
ggm_compare_ppc.default(Y1, Y2, Y3, type = "global", iter = 1000)
---
Estimates:

      contrast      KLD p_value
Y_g1 vs Y_g2 0.3341682    0.930
Y_g1 vs Y_g3 0.4125857    0.339
Y_g2 vs Y_g3 0.4047698    0.450
---
note:
p_value = p(T(Y_rep) > T(y)|Y)
KLD = Kullback-Leibler divergence

```

Here the null hypothesis of group equality was retained. Of course, this can be plotted as well—i.e.,

```

R> plot(ggm_ppc, critical = 0.05) +
+   theme_minimal(base_size = 14) +
+   theme(legend.position = "none", panel.grid.minor.x = element_blank())

```

This is displayed in Figure 6 (panel A). The black points denote the observed KLD and the critical region (i.e., $\alpha = 0.05$) is shaded in red.

Nodewise predictive check

BGGM also provides a nodewise, or variable specific, predictive check. The idea is the same as the *global* check, but compares KLD of the predictive distributions for each node. Recall from (14) that each row in the precision matrix corresponds to a regression model. As such, after converting the elements to coefficients, the predicted values can be computed. Because the variables are centered, having mean zero, KLD is easily obtained from the variance of respective predicted values. Further details can be found in Williams *et al.* (2019a).

For two groups and $p = 5$ (as an example), the *nodewise* approach is implemented with:

```

R> Y1 <- MASS::mvrnorm(500, rep(0, 5), Sigma = diag(5))
R> Y2 <- MASS::mvrnorm(500, rep(0, 5), Sigma = diag(5))

```

```

R> ggm_ppc <- ggm_compare_ppc(Y1, Y2,
+                             type = "nodewise", iter = 1000)
R> summary(ggm_ppc)
BGGM: Bayesian Gaussian Graphical Models
---
Type: GGM Comparison (Nodewise Predictive Check)
---
Call:
ggm_compare_ppc.default(Y1, Y2, type = "nodewise", iter = 1000)
---
Estimates:

contrast: Y_g1 vs Y_g2

  node   KLD p_value
    1 0.064  0.578
    2 0.236  0.330
    3 0.005  0.895
    4 0.144  0.431
    5 0.427  0.196

---
note:
p_value = p(T(Y_rep) > T(y) | Y)
KLD = Kullback-Leibler divergence

```

Again the null is retained, makes sense because each group has the same covariance structure. This object can then be plotted with `plot(ggm_ppc)` (Figure 6, panel B)

5.2. Bayesian hypothesis testing

The predictive methods performed well in the simulations reported in [Williams *et al.* \(2019a\)](#). However, like classical significance testing, they cannot provide evidence for a null model that certain (or all) edges are equal across groups. **BGGM** also includes a GGM comparison approach with the Bayes factor. Here the aim is not to reject a null model (i.e., $\mathcal{M}_0$ that groups are the same), but compares models to assess the relative evidence in the data between competing hypotheses.

In the general case of more than two groups, the Bayes factor is computed from the respective multivariate prior and posterior densities with **mvtnorm**. This is accomplished with the following steps:

1. Get S prior and posterior draws for $\boldsymbol{\rho}$ by sampling from the matrix- F prior and by using the Gibbs sampler (Appendix A.1).

2. Apply a Fisher transformation to the drawn partial correlations, i.e.,

$$\eta_{ij,g}^{(s)} = F(\rho_{ij,g}^{(s)}) = \frac{1}{2} \log \left(\frac{1 + \rho_{ij,g}^{(s)}}{1 - \rho_{ij,g}^{(s)}} \right), \text{ for } s = 1, \dots, S. \quad (33)$$

3. Compute the Fisher transformed differences via $\boldsymbol{\xi}^{(s)} = \mathbf{R}_{ij}\boldsymbol{\eta}^{(s)}$, for draws $s = 1, \dots, S$. These transformed parameters are approximately normally distributed in the prior and posterior as shown below. The hypothesis test can be written as $H_{0,ij} : \boldsymbol{\xi} = \mathbf{0}$ versus $H_{1,ij} : \boldsymbol{\xi} \neq \mathbf{0}$.
4. Estimate the posterior mean vector $\boldsymbol{\mu}_{\boldsymbol{\xi},N}$ and covariance matrix $\boldsymbol{\Psi}_{\boldsymbol{\xi},N}$, and the prior covariance matrix $\boldsymbol{\Psi}_{\boldsymbol{\xi},0}$ from their respective posterior and prior samples. Note that the prior mean vector equals $\mathbf{0}$.
5. Estimate the Bayes factor using

$$B_{01,ij} \approx \frac{N(\mathbf{0}; \boldsymbol{\mu}_{\boldsymbol{\xi},N}, \boldsymbol{\Psi}_{\boldsymbol{\xi},N})}{N(\mathbf{0}; \mathbf{0}, \boldsymbol{\Psi}_{\boldsymbol{\xi},0})}, \quad (34)$$

where $N(\mathbf{0}; \boldsymbol{\mu}, \boldsymbol{\Psi})$ denotes a multivariate normal density with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Psi}$ evaluated at $\mathbf{0}$.

The approximate normality of the posterior in Step 4 can be understood from the well-known fact that the sampling distribution (i.e., the likelihood) of a Fisher transformed correlation is approximately normally distributed (Fisher 1921, 1915). Furthermore, the prior of a Fisher transformed partial correlation, $\eta_{ij,g}$, is also approximately normally distributed and the linear transformation $\boldsymbol{\xi} = \mathbf{R}_{ij}\boldsymbol{\eta}$ preserves the normal approximation (Mulder 2016).

Pairwise hypothesis testing

In this section, we present a Bayes factor for testing whether partial correlations between variable i and j are equal across groups—i.e.,

$$\mathcal{H}_{0,ij} : \rho_{ij,1} = \dots = \rho_{ij,G} \text{ vs. } H_{1,ij} : \text{“not } H_{0,ij}\text{”}.$$

The multivariate normal density is then evaluated after applying a linear transformation, which for the posterior mean vector $\boldsymbol{\mu}_{\boldsymbol{\xi},N}$, follows

$$\mathbf{R}_{ij}\boldsymbol{\rho} = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \rho_{ij,g1} \\ \rho_{ij,g2} \\ \rho_{ij,g3} \\ \rho_{ij,g4} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \quad (35)$$

After applying the same transformation to the prior samples, the Bayes factor is computed with (34). Further, in the case of pairwise hypothesis testing, this is applied to each edge (or partial correlation) in the model.

As an example ($p = 5$), this is implemented with:

```
R> Y1 <- MASS::mvrnorm(500, rep(0, 5), Sigma = diag(5))
R> Y2 <- MASS::mvrnorm(500, rep(0, 5), Sigma = diag(5))
R> Y3 <- MASS::mvrnorm(500, rep(0, 5), Sigma = diag(5))

R> bf_ggm <- ggm_compare_bf(Y1, Y2, Y3,
+                           prior_sd = .5,
+                           iter = 2500, cores = 2)

R> ggm_bf_sel <- select( bf_ggm, BF_cut = 3)

R> summary(ggm_bf_sel)
```

BGGM: Bayesian Gaussian Graphical Models

Type: GGM Comparison (Bayesian Hypothesis Testing)

Bayes Factor: 3

Call:

`select.ggm_compare_bf(x = bf_ggm, BF_cut = 3)`

Hypotheses:

H_0 $\rho_{g1_ij} = \rho_{g2_ij} = \rho_{g3_ij}$

H_1 'not H_0 '

Evidence for H_0 : 100 %

Evidence for H_1 : 0 %

Inconclusive: 0 %

Adjacency Matrix (H_0)

```
  1 2 3 4 5
1 0 1 1 1 1
2 1 0 1 1 1
3 1 1 0 1 1
4 1 1 1 0 1
5 1 1 1 1 0
---
```

Adjacency Matrix (H1)

```

  1 2 3 4 5
1 0 0 0 0 0
2 0 0 0 0 0
3 0 0 0 0 0
4 0 0 0 0 0
5 0 0 0 0 0
---
```

Here the models are first estimated with `ggm_compare_bf(.)`, which applies each of the steps listed above, and then the function `select(., BF_cut = 3)` is used select edges for which there is evidence for the hypotheses. In this case, the null was true for each edge and the method correctly supported $\mathcal{H}_0$. The resulting object can then be summarized or plotted. The latter can either be a heatmap or “network” plot, as seen in Figure 1, and thus not included here.

Confirmatory hypothesis testing

The confirmatory methods described above (Section 4.4) can be extended to any number of groups. This powerful approach for comparing GGMs is implemented in **BGGM**. For example, a researcher can test their expectations in, say, treatment vs. control groups, where certain edges are predicted to increase. This can then be compared to a model predicting those edges stayed the same—i.e.,

$$\begin{aligned}
 \mathcal{H}_1 : & (\rho_{1,2_{g1}} > \rho_{1,2_{g2}}), (\rho_{1,3_{g1}} > \rho_{1,3_{g2}}), (\rho_{2,3_{g1}} > \rho_{2,3_{g2}}) \\
 \mathcal{H}_2 : & (\rho_{1,2_{g1}} = \rho_{1,2_{g2}}), (\rho_{1,3_{g1}} = \rho_{1,3_{g2}}), (\rho_{2,3_{g1}} = \rho_{2,3_{g2}}) \\
 \mathcal{H}_c : & \text{“not } \mathcal{H}_1 \text{ or } \mathcal{H}_2 \text{”}
 \end{aligned} \tag{36}$$

Note that this method can be used for any number of groups or sets of competing inequality and equality restrictions. We chose two groups and three edges ($p = 3$) to keep the notation concise, as it allows for describing the linear transformations. The first hypothesis follows

$$\mathbf{R}_{\mathcal{H}_1} \boldsymbol{\rho} = \begin{bmatrix} 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 \end{bmatrix} \begin{bmatrix} \rho_{1,2_{g1}} \\ \rho_{1,3_{g1}} \\ \rho_{2,3_{g1}} \\ \rho_{1,2_{g2}} \\ \rho_{1,3_{g2}} \\ \rho_{2,3_{g2}} \end{bmatrix} > \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \tag{37}$$

whereas the second hypothesis ($\mathcal{H}_2$) is similarly defined but the inequality ($>$) is replaced with an equality restriction ($=$). This is implemented with:

```

R> Y1 <- MASS::mvrnorm(500, rep(0, 3), Sigma = diag(3))
R> Y2 <- MASS::mvrnorm(500, rep(0, 3), Sigma = diag(3))

R> ggms_confirm <- ggm_compare_bf(Y1, Y2, prior_sd = 0.5,
+                               hypothesis = "g1_1--2 > g2_1--2,
+                                           g1_1--3 > g2_1--3,
+                                           g1_2--3 > g2_2--3;
+                                           g1_1--2 = g2_1--2,
+                                           g1_1--3 = g2_1--3,
+                                           g1_2--3 = g2_2--3")

R> summary(ggms_confirm)

BGGM: Bayesian Gaussian Graphical Models
---
Type: Confirmatory Hypothesis Testing
---
Hypotheses:

H1 g1_1--2>g2_1--2,g1_1--3>g2_1--3,g1_2--3>g2_2--3
H2 g1_1--2=g2_1--2,g1_1--3=g2_1--3,g1_2--3=g2_2--3
Hc 'not H1-2'
---
Posterior prob:

p(H1|Y) = 1e-04
p(H2|Y) = 0.985
p(Hc|Y) = 0.0149
---
Bayes factor matrix:
      H1      H2      Hc
H1 1.000 19006.711 288.281
H2 0.000      1.000  0.015
Hc 0.003    65.931   1.000
---
note: equal hypothesis prior probabilities

```

In this case, because each group was assumed to have equal edges, $\mathcal{H}_2$ was supported by the data with a posterior probability of ≈ 0.985 . In addition to summarizing `ggms_confirm`, `plot(ggms_confirm)` produces a plot similar to that in Figure 5.

6. Conclusion

We presented the **BGGM** package for Gaussian graphical models. A key aspect of our package is to extend inference beyond learning the conditional independence structure—i.e., detecting non-zero effects. To this end, **BGGM** includes several recently introduced methods that allow for comparing edges within and between GGMs and assessing predictability with Bayesian methods (e.g., R^2). Further, for determining the conditional independence structure, **BGGM** includes a variety of approaches that include credible intervals, the region of practical equivalence, and the Bayes factor.

Two additional contributions of **BGGM** include confirmatory hypothesis testing (Williams and Mulder 2019) and the ability to compare Gaussian graphical models (Williams *et al.* 2019a). The former allows for moving beyond automated, data driven, model selection and to testing theoretical predictions in the form of inequality or equality constraints. The confirmatory approach can also be used to compare, say, sets of partial correlations between control and treatment groups. Further, any number of GGMs can also be compared with respect to the posterior predictive distribution. These are novel contributions to the GGM literature, and **BGGM** provides the software implementation.

We are committed to further developing **BGGM**. In the next version, all of the methods presented in this work will accommodate ordinal, binary, and mixed data. Another future goal is to make a **shinydashboard**, which will provide a user friendly interface for implementing the methodology included in **BGGM**.

References

- Allen DM (1971). “Mean Square Error of Prediction as a Criterion for Selecting Variables.” *Technometrics*, **13**(3), 469–475. doi:10.1080/00401706.1971.10488811.
- Anufriev M, Panchenko V (2015). “Connecting the dots: Econometric methods for uncovering networks with an application to the Australian financial institutions.” *Journal of Banking and Finance*, **61**, S241–S255. ISSN 03784266. doi:10.1016/j.jbankfin.2015.08.034.
- Baba K, Shibata R, Sibuya M (2004). “Partial Correlation and Conditional Correlation as Measures of Conditional Independence.” *Australian & New Zealand Journal of Statistics*, **46**(4), 657–664. ISSN 1369-1473. doi:10.1111/j.1467-842X.2004.00360.x.
- Baba K, Sibuya M (2005). “Equivalence of Partial and Conditional Correlation Coefficients.” *J. Japan Statist. Soc.*, **35**(1), 1–19.
- Barnard J, McCulloch R, Meng XL (2000). “Modelling Covariance Matrices in Terms of Standard Deviations and Correlations With Applications to Shrinkage.” *Statistica Sinica*, **10**(4), 1281–1311. ISSN 10170405, 19968507. doi:10.2307/24306780.

- Betancourt M (2018). “Calibrating Model-Based Inferences and Decisions.” pp. 1–35. URL <http://arxiv.org/abs/1803.08393>.
- Cohen J, Cohen P (1983). *Applied multiple regression/correlation analysis for the behavioral sciences*. 2 edition. L. Erlbaum Associates, Hillsdale. ISBN 0898592682.
- Das A, Sampson A, Lainscsek C, Muller L, Lin W, Doyle J, Cash S, Halgren E, Sejnowski T (2017). “Interpretation of the Precision Matrix and Its Application in Estimating Sparse Brain Connectivity during Sleep Spindles from Human Electroencephalography Recordings.” *Neural computation*, **29**(3), 603–642. ISSN 0899-7667. doi:10.1162/NECO{_}a{_}00936.
- Dempster A (1972). “Covariance Selection.” *Biometrics*, **28**(1), 157–175. ISSN 0006341X. doi:10.2307/2528966.
- Dickey JM (1971). “The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters.” *The Annals of Mathematical Statistics*, **42**(1), 204–223. ISSN 0003-4851. doi:10.1214/aoms/1177693507.
- Epskamp S, Borsboom D, Fried EI (2018a). “Estimating psychological networks and their accuracy: A tutorial paper.” *Behavior Research Methods*, **50**(1), 195–212. ISSN 15543528. doi:10.3758/s13428-017-0862-1.
- Epskamp S, Waldorp LJ, Mottus R, Borsboom D (2018b). “The Gaussian Graphical Model in Cross-Sectional and Time-Series Data.” *Multivariate Behavioral Research*, **53**(4), 453–480. doi:10.1080/00273171.2018.1454823.
- Fisher R (1915). “Frequency Distribution of the Values of the Correlation Coefficient in Samples from an Indefinitely Large Population.” *Biometrika*, **10**(4), 507–521.
- Fisher R (1921). “On the ”Probable Error” of a Coefficient of Correlation Deduced from a Small Sample.” *Metron*, **1**.
- Fried EI, Eidhof MB, Palic S, Costantini G, Dijk HMHv, Bockting CLH, Engelhard I, Armour C, Nielsen ABS, Karstoft Ki (2018). “Replicability and Generalizability of Posttraumatic Stress Disorder (PTSD) Networks : A Cross-Cultural Multisite Study of PTSD Symptoms in Four Trauma Patient Samples.” *Clinical Psychological Science*, **6**(3), 335–351. doi:10.1177/2167702617745092.
- Friedman J, Hastie T, Tibshirani R (2008). “Sparse inverse covariance estimation with the graphical lasso.” *Biostatistics*, **9**(3), 432–441. ISSN 14654644. doi:10.1093/biostatistics/kxm045.
- Gelman A, Carlin JB, Stern HS, Dunson BD, Vehtari A, Rubin DB (2014). *Bayesian Data Analysis*. CRC Press, Boca Raton. ISBN 9781439898208.
- Gelman A, Goodrich B, Gabry J, Ali I (2017). “R-squared for Bayesian regression models.” URL http://www.stat.columbia.edu/~gelman/research/unpublished/bayes_R2.pdf.

- Haslbeck JMB, Fried EI (2017). “How predictable are symptoms in psychopathological networks? A reanalysis of 18 published datasets.” *Psychological Medicine*, **47**(16), 2267–2276. ISSN 14698978. doi:10.1017/S0033291717001258. URL <https://doi.org/10.1017/S0033291717001258>.
- Haslbeck JMB, Waldorp LJ (2015). “MGM: Estimating Time-Varying Mixed Graphical Models in High-Dimensional Data.”
- Haslbeck JMB, Waldorp LJ (2018). “How well do network models predict observations? On the importance of predictability in network models.” *Behavior Research Methods*, **50**(2), 853–861. ISSN 15543528. doi:10.3758/s13428-017-0910-x. URL <http://link.springer.com/10.3758/s13428-017-0910-x>.
- Højsgaard S, Edwards D, Lauritzen S (2012). *Graphical Models with R*. ISBN 1461422981. doi:10.1007/978-1-4614-2299-0.
- Janková J, van de Geer S (2017). “Honest confidence regions and optimality in high-dimensional precision matrix estimation.” *Test*, **26**(1), 143–162. ISSN 11330686. doi:10.1007/s11749-016-0503-5.
- Javanmard A, Montanari A (2013). “Confidence Intervals and Hypothesis Testing for High-Dimensional Regression.” pp. 1–21. ISSN 15337928.
- Kass RE, Raftery AE (1995). “Bayes Factors.” *Journal of the American Statistical Association*, **90**(430), 773–795.
- Klugkist I, Kato B, Hoijsink H (2005). “Bayesian model selection using encompassing priors.” *Statistica Neerlandica*, **59**(1), 57–69. ISSN 00390402. doi:10.1111/j.1467-9574.2005.00279.x.
- Kruschke JK (2011). “Bayesian Assessment of Null Values Via Parameter Estimation and Model Comparison.” *Perspectives on Psychological Science*, **6**, 299–312. ISSN 1745-6916. doi:10.1177/1745691611406925.
- Kuismin M, Sillanpää M (2017). “Estimation of covariance and precision matrix, network structure, and a view toward systems biology.” *Wiley Interdisciplinary Reviews: Computational Statistics*, **9**(6), 1–13. ISSN 19390068. doi:10.1002/wics.1415.
- Kwan CCY (2014). “A Regression-Based Interpretation of the Inverse of the Sample Covariance Matrix.” *Spreadsheets in Education*, **7**(1). ISSN 1448-6156.
- Lauritzen SL (1996). *Graphical models*, volume 17. Clarendon Press.
- Leday GGR, Richardson S (2018). “Fast Bayesian inference in large Gaussian graphical models.” pp. 1–26.
- Ledoit O, Wolf M (2004). “A well-conditioned estimator for large-dimensional covariance matrices.” *Journal of Multivariate Analysis*, **88**(2), 365–411. ISSN 0047259X. doi:10.1016/S0047-259X(03)00096-4.

- Lee MD (2018). “Bayesian Methods in Cognitive Modeling.” In *Stevens’ Handbook of Experimental Psychology and Cognitive Neuroscience*, pp. 1–48. John Wiley & Sons, Inc., Hoboken, NJ, USA. doi:10.1002/9781119170174.epcn502.
- Li P, Redden DT (2015). “Small sample performance of bias-corrected sandwich estimators for cluster-randomized trials with binary outcomes.” *Statistics in Medicine*, **34**(2), 281–296. ISSN 10970258. doi:10.1002/sim.6344. URL <http://www.ncbi.nlm.nih.gov/pubmed/25345738><http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=PMC4268228>.
- Meinshausen N, Bühlmann P (2006). “High-dimensional graphs and variable selection with the Lasso.” *Annals of Statistics*, **34**(3), 1436–1462. ISSN 00905364. doi:10.1214/009053606000000281.
- Mohammadi A, Wit EC (2015a). “Bayesian structure learning in sparse Gaussian graphical models.” *Bayesian Analysis*, **10**(1), 109–138. ISSN 19316690. doi:10.1214/14-BA889.
- Mohammadi A, Wit EC (2015b). “BDgraph: An R Package for Bayesian Structure Learning in Graphical Models.” ISSN 0884-0431. doi:10.1359/JBMR.0301229.
- Mulder J (2016). “Bayes factors for testing order-constrained hypotheses on correlations.” *Journal of Mathematical Psychology*, **72**, 104–115. ISSN 10960880. doi:10.1016/j.jmp.2014.09.004.
- Mulder J, Hoijsink H, de Leeuw C (2012). “BIEMS: A Fortran 90 program for calculating {B}ayes factors for inequality and equality constrained model.” *Journal of Statistical Software*, **46**.
- Mulder J, Raúl Pericchi L (2018). “The Matrix-F Prior for Estimating and Testing Covariance Matrices.” *Bayesian Analysis*, (4), 1–22. ISSN 19316690. doi:10.1214/17-BA1092.
- Schäfer J, Strimmer K (2005). “A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics.” *Statistical Applications in Genetics and Molecular Biology*, **4**(1). ISSN 1544-6115. doi:10.2202/1544-6115.1175.
- Stephens G (1998). “On the Inverse of the Covariance Matrix in Portfolio Analysis.” *The Journal of Finance*, **53**(5), 1821–1827.
- Tibshirani R (1996). “Regression Selection and Shrinkage via the Lasso.” doi:10.2307/2346178.
- van Borkulo CD, Boschloo L, Kossakowski JJ, Tio P, Schoevers RA, Borsboom D, Waldorp LJ (2016). “Comparing network structures on three aspects: A permutation test.” *Manuscript submitted*, (March), 34. doi:10.13140/RG.2.2.29455.38569.
- Vehtari A, Gelman A, Gabry J (2017). “Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC.” *Statistics and Computing*, **27**(5), 1413–1432. ISSN 15731375. doi:10.1007/s11222-016-9696-4.

- Wagenmakers EJ, Lodewyckx T, Kuriyal H, Grasman R (2010). “Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method.” *Cognitive Psychology*, **60**(3), 158–189. ISSN 00100285. doi:[10.1016/j.cogpsych.2009.12.001](https://doi.org/10.1016/j.cogpsych.2009.12.001).
- Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4. URL <http://ggplot2.org>.
- Williams DR (2018). “Bayesian Inference for Gaussian Graphical Models: Structure Learning, Explanation, and Prediction.” doi:[10.31234/OSF.IO/X8DPR](https://doi.org/10.31234/OSF.IO/X8DPR).
- Williams DR, Mulder J (2019). “Bayesian Hypothesis Testing for Gaussian Graphical Models: Conditional Independence and Order Constraints.”
- Williams DR, Rast P (2018). “Back to the basics: Rethinking partial correlation network methodology.” pp. 1–15. doi:[10.17605/OSF.IO/FNDRU](https://doi.org/10.17605/OSF.IO/FNDRU).
- Williams DR, Rast P, Pericchi LR, Mulder J (2019a). “Comparing Gaussian Graphical Models with the Posterior Predictive Distribution and Bayesian Model Selection.” doi:<https://doi.org/10.31234/osf.io/yt386>.
- Williams DR, Rast P, Pericchi LR, Mulder J (2019b). “Comparing Gaussian Graphical Models with the Posterior Predictive Distribution and Bayesian Model Selection.” doi:<https://doi.org/10.31234/osf.io/yt386>.
- Williams DR, Rhemtulla M, Wysocki AC, Rast P (2019c). “On Nonregularized Estimation of Psychological Networks.” *Multivariate Behavioral Research*, **0**(0), 1–23. ISSN 00273171. doi:[10.1080/00273171.2019.1575716](https://doi.org/10.1080/00273171.2019.1575716). URL <https://doi.org/10.1080/00273171.2019.1575716>.
- Zerenner T, Friederichs P, Lehnertz K, Hense A (2014). “A Gaussian graphical model approach to climate networks.” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **24**(2), 023103. ISSN 1054-1500. doi:[10.1063/1.4870402](https://doi.org/10.1063/1.4870402).
- Zhang CH, Zhang SS (2014). “Confidence intervals for low dimensional parameters in high dimensional linear models.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **76**(1), 217–242. ISSN 13697412. doi:[10.1111/rssb.12026](https://doi.org/10.1111/rssb.12026). URL <http://doi.wiley.com/10.1111/rssb.12026>.
- Zhang R, Ren Z, Chen W (2018). “SILGGM: An extensive R package for efficient statistical inference in large-scale gene networks.” *PLoS Computational Biology*, **14**(8), e1006369. ISSN 15537358. doi:[10.1371/journal.pcbi.1006369](https://doi.org/10.1371/journal.pcbi.1006369).
- Zhao T, Liu H, Roeder K, Lafferty J, Wasserman L (2012). “The huge Package for High-dimensional Undirected Graph Estimation in R.” *Journal of machine learning research : JMLR*, **13**, 1059–1062. ISSN 1532-4435. doi:[10.1002/aur.1474](https://doi.org/10.1002/aur.1474). Replication.

A. Supplementary equations

A.1. Gibbs sampler

For a precision matrix Θ having a matrix- $F(\nu, \delta, \mathbf{B})$ prior distribution, the probability density function is given by

$$p^F(\Theta) = \frac{\Gamma_k(\frac{\nu+\delta+p-1}{2})}{\Gamma_k(\frac{\nu}{2})\Gamma_k(\frac{\delta+p-1}{2})|\mathbf{B}|^{\frac{\nu}{2}}} |\Theta|^{\frac{\nu-p-1}{2}} |\mathbf{I}_k + \Theta\mathbf{B}^{-1}|^{-\frac{\nu+\delta+p-1}{2}}. \quad (38)$$

Following [Mulder and Raúl Pericchi \(2018\)](#), this implies that the covariance matrix Σ follows a matrix- $F(\delta + p - 1, \nu - p + 1, \mathbf{B}^{-1})$, which can be written as a scale mixture of inverse Wishart distributions, i.e.,

$$\begin{aligned} \Sigma|\Psi &\sim IW(\nu, \Psi) \\ \Psi &\sim W(\delta + p - 1, \mathbf{B}^{-1}). \end{aligned} \quad (39)$$

Furthermore, the likelihood for n independent observations from the multivariate normal model is given by

$$p(\mathbf{Y}|\mu, \Sigma) \propto |\Sigma|^{-n/2} \exp\left\{-\frac{1}{2}\text{tr}\Sigma^{-1}\mathbf{S}\right\} \exp\left\{-\frac{n}{2}(\mu - \bar{\mathbf{y}})'\Sigma^{-1}(\mu - \bar{\mathbf{y}})\right\}, \quad (40)$$

where $\mathbf{S}$ is the sums of squares matrix and $\bar{\mathbf{y}}$ is the vector of sample means. Therefore when using a flat prior for the nuisance parameter μ , the posterior (including the parameter expansion with Ψ) can be written as

$$\begin{aligned} p(\mu, \Sigma, \Psi|\mathbf{Y}) &\propto p^{IW}(\Sigma|\Psi)p^W(\Psi)p(\mathbf{Y}|\mu, \Sigma) \\ &\propto |\Psi|^{(\nu+\delta-2)/2}|\Sigma|^{-(n+\nu+p+1)/2} \exp\left\{-\frac{1}{2}\text{tr}\Psi\Sigma^{-1}\right\} \\ &\quad \exp\left\{-\frac{1}{2}\text{tr}\Psi\mathbf{B}\right\} \exp\left\{-\frac{1}{2}\text{tr}\Sigma^{-1}\mathbf{S}\right\} \\ &\quad \exp\left\{-\frac{n}{2}(\mu - \bar{\mathbf{y}})'\Sigma^{-1}(\mu - \bar{\mathbf{y}})\right\} \\ \Rightarrow p(\Sigma, \Psi|\mathbf{Y}) &\propto |\Psi|^{(\nu+\delta-2)/2}|\Sigma|^{-(n+\nu+p)/2} \exp\left\{-\frac{1}{2}\text{tr}\Psi\Sigma^{-1}\right\} \\ &\quad \exp\left\{-\frac{1}{2}\text{tr}\Psi\mathbf{B}\right\} \exp\left\{-\frac{1}{2}\text{tr}\Sigma^{-1}\mathbf{S}\right\}. \end{aligned}$$

Hence a Gibbs sampler can then be formed using the following conditional posteriors

$$\begin{aligned} \Sigma|\Psi, \mathbf{Y} &\sim IW(n + \nu - 1, \mathbf{S} + \Psi) \\ \Psi|\Sigma, \mathbf{Y} &\sim W(\nu + \delta + p - 1, (\Sigma^{-1} + \mathbf{B})^{-1}). \end{aligned}$$

A posterior sample for the covariance matrix Σ , and thus of the precision matrix Θ and the partial correlations, can be obtained by iteratively sampling Σ and Ψ from their respective conditional posterior distribution.

A.2. KL divergence

For a multivariate normal distribution KLD is defined as

$$\Theta_{g1}\Theta_{g2} = \frac{1}{2} \left[\text{tr}(\Theta_{g1}^{-1}\Theta_{g2}) - \log(|\Theta_{g1}^{-1}\Theta_{g2}|) - p \right], \quad (41)$$

where p is the number of variables. Note that inverting Θ_{g1} results in the covariance matrix Σ_{g1} . Repeating this process for each posterior sample produces the predictive distribution of KLD. This distribution can be thought of as the amount of divergence (or relative entropy) we would expect to see assuming that the null model of group equality were *true*. This is then compared to the observed KL divergence.

Affiliation:

Donald R. Williams
 Department of Psychology
 University of California, Davis
 One Shields Ave, Davis, CA
 E-mail: drwilliams@ucdavis.edu
 URL: <https://www.donaldrwilliamsmethods.com/>