



Learning Permutation Symmetry of a Gaussian Vector with `gips` in R

Adam Chojecki
Warsaw University
of Technology

Paweł Morgen
Warsaw University
of Technology

Bartosz Kołodziejek 
Warsaw University
of Technology

Abstract

The study of hidden structures in data presents challenges in modern statistics and machine learning. We introduce the `gips` package in R, which identifies permutation subgroup symmetries in Gaussian vectors. `gips` serves two main purposes: Exploratory analysis in discovering hidden permutation symmetries and estimating the covariance matrix under permutation symmetry. It is competitive to canonical methods in dimensionality reduction while providing a new interpretation of the results. `gips` implements a novel Bayesian model selection procedure within Gaussian vectors invariant under the permutation subgroup introduced in Graczyk, Ishi, Kołodziejek, and Massam (2022b, *The Annals of Statistics*).

Keywords: Bayesian model selection, permutation symmetry, high-dimensional statistics, dimensionality reduction, parameter sharing, R.

1. Introduction

The study of hidden structures in the data is one of the biggest challenges in modern mathematical statistics and machine learning (Hastie, Tibshirani, and Friedman 2009). Extracting meaningful information from high-dimensional datasets, where the number of variables p exceeds the number of observations n , poses a significant hurdle due to the curse of dimensionality.

One solution to the problem of an insufficient number of observations relative to the number of variables is to restrict to models with lower dimensionality. Graphical models have been introduced for this purpose (Lauritzen 1996), where a conditional independence structure (graph Markovian structure) is imposed on the distribution of a random vector. Such structures are conveniently described by graphs and allow for a reduction in the dimensionality of the problem. However, if the graph is not sparse enough, then such a procedure does

not allow for a reliable estimation of the covariance matrix. We note that the study of the covariance matrix is the basic way to describe the dependency structure of a random vector and provides a convenient way to quantify the dependencies between variables.

If the data is insufficient and some inference must be performed, one has to propose additional assumptions or restrictions. In such a situation, colored graphical models could be considered, where, in addition to conditional independence, certain equality conditions on the covariance matrix are imposed. Incorporating such equality conditions in colored graphical models is an example of parameter sharing. This concept allows for a reduction of dimensionality and can effectively incorporate domain knowledge into the model architecture. A notable example of parameter sharing, which possesses these advantages, is the convolution technique (Hastie *et al.* 2009) used in image and video processing, enabling efficient feature extraction and pattern recognition.

A rich family of such symmetry conditions can be expressed using the language of permutations. This idea was introduced in Andersson (1975); Andersson and Madsen (1998) and Højsgaard and Lauritzen (2008). In the latter paper, three types of such models (RCOP among them) were introduced to describe situations where some entries of concentration or partial correlation matrices are approximately equal. These equalities can be represented by a colored graph. The RCOP model, apart from the graph Markovian structure, permits additional invariance of the distribution with respect to some permutation subgroup. We say that the distribution of a p -dimensional random vector Z is invariant under permutation subgroup Γ on $V = \{1, \dots, p\}$ if $Z = (Z_i)_{i \in V}$ has the same distribution as $(Z_{\sigma(i)})_{i \in V}$ for any permutation $\sigma \in \Gamma$ (Andersson 1975). E.g., if $p = 4$ and Γ is the cyclic subgroup generated by a permutation $\sigma = (1, 2, 3)$ (written in the cycle notation), then the Γ -invariance is equivalent to

$$(Z_1, Z_2, Z_3, Z_4) \stackrel{d}{=} (Z_2, Z_3, Z_1, Z_4) \stackrel{d}{=} (Z_3, Z_1, Z_2, Z_4).$$

This property is called the permutation symmetry of the distribution of Z and imposes significant symmetry conditions on the model.

The case when the conditional dependency graph is unknown or known to be the complete graph was studied in Graczyk *et al.* (2022b). In that paper, the authors introduced a Bayesian model selection procedure for the case when Z is a Gaussian vector. In other words, by assuming a prior distribution on the parameters, they derived the posterior probability of a specific model. This allows one to find the permutation group under which (most likely) the data is invariant. Not only does this result in dimensionality reduction but also provides a simple and natural interpretability of the results. For example, if the distribution of Z is invariant under swapping its i th and j th entries, then one can say that both Z_i and Z_j play a symmetrical role in the model.

The concept of group invariance finds application in various domains and often leads to improved estimation properties. If the group under which the model is invariant is known, precise convergence rates for the regularized covariance matrix were derived in Shah and Chandrasekaran (2012), demonstrating significant statistical advantages in terms of sample complexity. Another noteworthy paper, Soloveychik, Trushin, and Wiesel (2016), explores group symmetries for estimating complex covariance matrices in non-Gaussian models, which are invariant under a known permutation subgroup. However, neither of these articles provides guidance on identifying the permutation subgroup when it is unknown, which is typically the case in practical applications.

Identifying the permutation subgroup symmetry can be interpreted as an automated way of extracting expert knowledge from the data. Discovering the underlying symmetries allows for a deeper understanding of the relationships and dependencies between variables, offering insights that may not be apparent through traditional analysis alone. The automated approach reduces the reliance on manual exploration and expert intervention.

In the present paper, we introduce an R (R Core Team 2024) package called **gips** (acronym derived from “Gaussian model invariant by permutation symmetry”, Chojecki, Morgen, and Kołodziejek 2025), which implements the model selection procedure described in Graczyk *et al.* (2022b). The **gips** package, presented in this paper, serves two purposes:

1. Discovering hidden permutation symmetries among variables (exploratory analysis).
2. Estimating covariance matrix under the assumption of known permutation symmetry.

Both points are limited to the Gaussian setting. To the best of our knowledge, there are currently no other software packages available (in R or any other programming language) that address the topic of finding permutation symmetry. Our approach focuses on zero-mean Gaussian vectors, although the method can be applied to centered data and, if the sample size n is reasonably large, to standardized data as well, see Section 2.7. Package **gips** (Chojecki *et al.* 2025) is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=gips>.

Let Z be a Gaussian vector with a known mean. If we assume full symmetry of the model, see Section 2.2, meaning that the distribution of Z is invariant under any permutation, then the maximum likelihood estimator (MLE) of the covariance matrix requires only a single sample ($n_0 = 1$) to exist. Somewhat surprisingly, the same phenomenon applies when the normal sample is invariant under a cyclic subgroup generated by a cycle of length p . While it is natural to consider permutation symmetries alongside conditional independence structures, we follow Graczyk *et al.* (2022b) and assume no conditional independencies among the variables. Such an approach already enables a substantial reduction in dimensionality, accompanied by a readily interpretable outcome. The development of the method to incorporate non-trivial graph Markovian structures is a topic for future research, and we will consider expanding the package if a new theory emerges. The first step towards generalizing the theory to homogeneous graphs has already been taken in Graczyk, Ishi, and Kołodziejek (2022a). Additionally, a simple heuristic can be employed to identify non-trivial Markovian structures using our model, see Graczyk *et al.* (2022b, Section 1.2), Graczyk, Ishi, Kołodziejek, and Massam (2022c, Section 4.1), and Section 3.1 in this paper.

Although there are no other software packages available for finding permutation symmetries in data, we have made the decision to compare the results of our model with canonical methods commonly used to tackle high-dimensional problems, namely ridge and graphical LASSO (GLASSO) estimation and model selection (implemented, for instance, in R packages: **huge**, Jiang *et al.* 2021, and **rags2ridges**, Peeters, Bilgrau, and Van Wieringen 2022, 2023). These methods correspond to estimation with constraints or, conversely, to Bayesian estimation with Gaussian or Laplace priors, respectively (James, Witten, Hastie, and Tibshirani 2021, Section 6.2.2). We demonstrate that **gips** is competitive with these widely used approaches in terms of dimensionality reduction properties, and moreover, it offers interpretability of the results in terms of permutation symmetries.

Furthermore, it is worth noting that due to the discrete nature of the problem, we believe that finding permutation symmetry cannot be adequately addressed by penalized likelihood methods, which are generally much faster than Bayesian methods. Although other methods (which do not have available implementations to our best knowledge) allow for model selection within colored graphical Gaussian models, none of them are applicable to permutation invariant models (RCOP models). Compared to other models (such as RCON, RCOR in Højsgaard and Lauritzen 2008, which correspond to different type of restrictions/symmetries), RCOP models possess a more elegant algebraic description and offer a natural interpretation (Gehrmann 2011; Gao and Massam 2015; Massam, Li, and Gao 2018; Li, Gao, and Massam 2020; Li, Sun, Wang, and Gao 2021; Ranciati, Roverato, and Luati 2021).

The replication code is available in the supplementary files and at https://github.com/PrzeChoj/gips_replication_code.

1.1. Overview of the paper

The paper is organized as follows. The Introduction consists of four subsections. In the next subsection, we present two low-dimensional toy examples that illustrate the use of **gips**. In the subsequent subsection, we discuss the potential for successfully exploiting group symmetry in many natural real-life problems. In the final subsection of the Introduction, we argue that it is both necessary and sufficient to focus on cyclic symmetries, which are more tractable.

Section 2 provides the necessary methodological background on permutation symmetries and its possible interpretations. Moreover, we define the Bayesian model proposed in Graczyk *et al.* (2022b), specialized to cyclic subgroups. We also introduce an Markov chain Monte Carlo (MCMC) algorithm that allows the estimation of the maximum a posteriori (MAP) within our Bayesian model, and we discuss the issue of centering and standardizing input data.

Section 3 is dedicated to numerical simulations. We present a high-dimensional example using breast cancer data from Miller *et al.* (2005). Additionally, we use a heuristic approach from Graczyk *et al.* (2022b) for identifying the graphical model invariant under permutation symmetries (RCOP model from Højsgaard and Lauritzen (2008)) and apply this procedure to the real-life example. In the subsequent subsections, we examine the impact of hyperparameters on model selection and compare **gips** with competing packages that facilitate dimensionality reduction.

Finally, in Section 4 we draw some conclusions.

An example to Section 1.4 is presented in Appendix A. Mathematical details behind the Bayesian model are relegated to the Appendix B.

1.2. Toy examples

We illustrate the concept of permutation symmetry using the **gips** package in two simple use cases. These examples demonstrate how permutational symmetry can enhance the data mining process. A similar procedure was successfully applied to the Frets' heads dataset (Graczyk *et al.* 2022b, Section 4.2) and the mathematical marks dataset (Graczyk *et al.* 2022a, Section 4).

In the first example, we use **aspirin** dataset from the **HSAUR2** package. By examining the covariance matrix, we manually choose a reasonable permutation symmetry. Additionally, we

employ the **gips** package to demonstrate that our algorithm generates reasonable estimates. For the second example, we utilize the **oddbooks** dataset from the **DAAG** package. We showcase how one can incorporate expert field knowledge in the analysis. We use the **gips** to find the permutation symmetry and interpret the result.

A standard personal computer (PC) can execute the entire code in this section within 10 seconds.

Aspirin dataset

This dataset consists of information about a meta-analysis of the efficacy of Aspirin (versus placebo) in preventing death after a myocardial infarct.

We renumber the columns for better readability:

```
R> data("aspirin", package = "HSAUR2")
R> Z <- aspirin
R> Z[, c(2, 3)] <- Z[, c(3, 2)]
R> names(Z) <- names(Z)[c(1, 3, 2, 4)]
R> head(Z, 4)
```

```
      dp da  tp  ta
1    67 49 624 615
2    64 44 771 758
3   126 102 850 832
4    38 32 309 317
```

Each of the $n = 7$ rows in Z corresponds to a different study, and the $p = 4$ columns represent the following: **dp**: Number of deaths after placebo, **da**: Number of deaths after Aspirin, **tp**: Total number of subjects treated with placebo, **ta**: Total number of subjects treated with Aspirin.

Initially, we calculate the empirical covariance matrix S :

```
R> n <- nrow(Z)
R> p <- ncol(Z)
R> S <- cov(Z)
```

Note that since $n = 7$ is greater than $p = 4$, S is the standard MLE of Σ in the (unrestricted) Gaussian model. The heatmap of the S matrix is shown in Figure 1.

We observe significant similarities between the empirical covariances of variables **tp** (column 3) and **ta** (column 4). They exhibit comparable variances ($S[3,3] \approx S[4,4]$), and their covariances with the other variables also show resemblance ($S[1,3] \approx S[1,4]$ and $S[2,3] \approx S[2,4]$).

By definition, the distribution of a random vector $Z = (Z_1, Z_2, Z_3, Z_4)^\top$ is invariant under the permutation (3, 4) if the distributions of $(Z_1, Z_2, Z_3, Z_4)^\top$ and $(Z_1, Z_2, Z_4, Z_3)^\top$ coincide. When Z follows a centered Gaussian distribution, this property can be expressed purely in terms of its covariance matrix, leading to the following conditions: $\text{VAR}[Z_3] = \text{VAR}[Z_4]$, $\text{COV}[Z_1, Z_3] = \text{COV}[Z_1, Z_4]$, and $\text{COV}[Z_2, Z_3] = \text{COV}[Z_2, Z_4]$. We observe that the structure

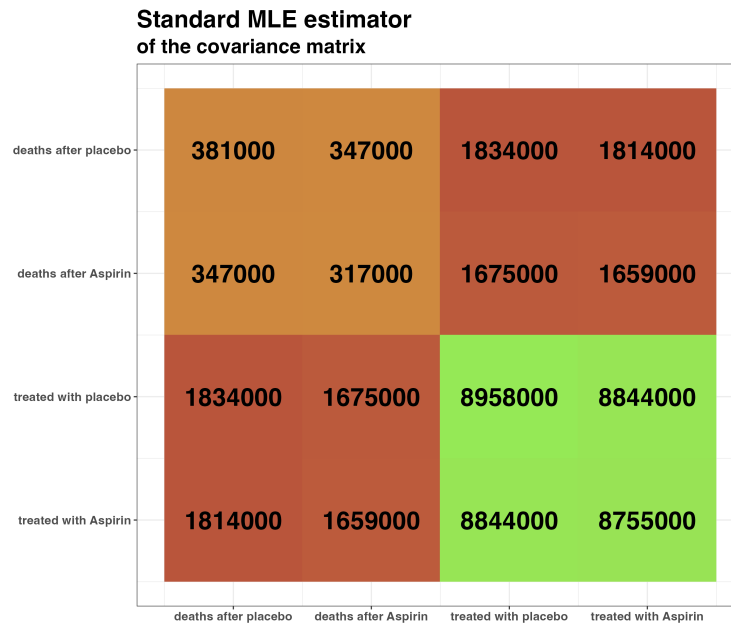


Figure 1: Standard MLE of the covariance matrix of the `aspirin` dataset. We can see that some entries of \mathbf{S} have similar colors, which suggests a lower dimensional model with equality constraints.

of \mathbf{S} closely corresponds to that of the covariance matrix of a random vector invariant under the permutation $(3, 4)$. By observing that $\mathbf{S}[1, 1] \approx \mathbf{S}[2, 2]$, $\mathbf{S}[1, 3] \approx \mathbf{S}[2, 3]$, and $\mathbf{S}[1, 4] \approx \mathbf{S}[2, 4]$, we can also argue that the data is invariant under the permutation $(1, 2)$ or even $(1, 2)(3, 4)$.

We want to emphasize that such manual exploration becomes infeasible for larger values of p due to the massive number and complexity of possible relationships. A priori, it is unclear which scenario is preferable (one can compare Bayesian information criterion (BIC), but the MLE does not always exist). The `gips` package uses the Bayesian paradigm (described in detail in Section 2.5) to precisely quantify posterior probabilities of considered permutation groups. The workflow in `gips` is as follows: First, use the `gips()` function to define an object of the class ‘`gips`’ that contains all the necessary information for the model. Next, use the `find_MAP()` function with an optimizer of your choice to find the permutation that provides the maximum a posteriori estimate. Finally, we use the `project_matrix()` function to obtain the MLE of the covariance matrix in the invariant model, which will serve as a more stable covariance estimator. The process can be summarized as follows:

```
R> g <- gips(S, n)
R> g_MAP <- find_MAP(g, optimizer = "BF", save_all_perms = TRUE,
+   return_probabilities = TRUE)
R> g_MAP
```

The permutation $(1, 2)(3, 4)$:

- was found after 17 posteriori calculations;
- is 3.374 times more likely than the $()$ permutation.

According to the output of `find_MAP()`, the permutation $(1,2)(3,4)$ best reflects the symmetries of the models and is over 3 times more probable (under our Bayesian setting) than the identity permutation $()$, which corresponds to no symmetry. The invariance with respect to the permutation $(3,4)$ arises from the fact that the samples of patients treated with aspirin and placebo had similar sizes (treatment-control choice of experimenters) and is not related to the effects of aspirin. On the other hand, the invariance with respect to the permutation $(1,2)$ signifies the lack of aspirin treatment effect. The permutation $(1,2)(3,4)$ corresponds to both of these effects. We emphasize that this study is an exploratory analysis rather than a statistical test.

We can easily calculate probabilities of all symmetries using a built-in function:

```
R> get_probabilities_from_gips(g_MAP)
```

$(1,2)(3,4)$	$(3,4)$	$(1,2)$	$()$	$(1,4)$	$(1,3)$
5.107108e-01	1.695605e-01	1.663982e-01	1.513854e-01	4.341644e-04	4.047690e-04
$(2,4)$	$(2,3)$	$(1,3,2,4)$	$(1,3)(2,4)$	$(1,4)(2,3)$	$(1,3,4)$
3.797581e-04	3.607292e-04	1.240381e-04	7.410652e-05	7.406484e-05	2.197791e-05
$(1,2,4)$	$(1,2,3)$	$(2,3,4)$	$(1,2,4,3)$	$(1,2,3,4)$	
2.026609e-05	1.813565e-05	1.782315e-05	7.676231e-06	7.528912e-06	

or compare two permutations of interest:

```
R> compare_posteriories_of_perms(g_MAP, "(34)")
```

The permutation $(1,2)(3,4)$ is 3.012 times more likely than the $(3,4)$ permutation.

```
R> compare_posteriories_of_perms(g_MAP, "(12)")
```

The permutation $(1,2)(3,4)$ is 3.069 times more likely than the $(1,2)$ permutation.

```
R> compare_posteriories_of_perms(g_MAP, "()")
```

The permutation $(1,2)(3,4)$ is 3.374 times more likely than the $()$ permutation.

Note that for $p = 4$, there are $p! = 24$ different permutations, but only 17 distinct symmetries are reported above. This is because some permutations correspond to the same symmetry. More precisely, it is the group generated by a permutation σ and not σ itself that identifies the symmetry. For example $\sigma_1 = (1, 2, 3)$ and $\sigma_2 = (1, 3, 2)$ generate the same group.

We also note that given the small number of variables ($p = 4$), the space of possible permutation symmetries is also small. Consequently, we were able to compute the exact posterior probabilities of our Bayesian model for every single permutation symmetry. The number of permutation symmetries grows superexponentially with p , e.g., for $p = 10$ its cardinality is approximately 1 million (see sequence A051625, [OEIS 2025](#)). Thus, for larger p we recommend




Figure 2: The MLE of the covariance matrix of the `aspirin` dataset in the model invariant under the permutation $(1, 2)(3, 4)$. We can see that the entries corresponding to similar colors in Figure 1 are exactly equal.

using the implemented Metropolis-Hastings algorithm to approximate these probabilities, see Section 2.6.

Assuming that the data actually come from a distribution invariant under the permutation $(1, 2)(3, 4)$, we can provide a new estimate for the covariance matrix. Formally, we project the matrix S onto the space of positive definite matrices that are invariant under the permutation $(1, 2)(3, 4)$ (for further details, refer to Section 2.4). In practice, we enforce the desired equalities by averaging.

```
R> S_projected <- project_matrix(S, g_MAP)
```

One can easily plot the found covariance estimator with a line:

```
R> plot(g_MAP, type = "heatmap")
```

It is shown in Figure 2 (we made cosmetic modifications to this plot; the exact code is provided in the replication code available online).

The `S_projected` matrix can now be interpreted as a more stable covariance matrix estimator, see, e.g., Shah and Chandrasekaran (2012); Soloveychik *et al.* (2016).

Books dataset

This dataset consists of information about thickness (mm), height (cm), width (cm), and weight (g) of 12 books.

```
R> data("oddbooks", package = "DAAG")
R> head(oddbooks, 4)
```

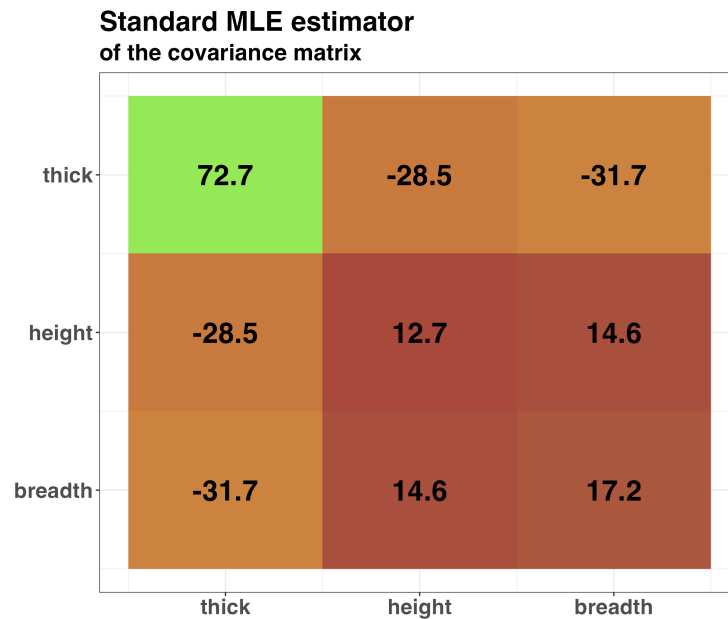



Figure 3: Standard MLE of the covariance matrix of the `oddbooks` dataset. We can see that some entries of S have similar colors, which suggests a lower dimensional model with equality constraints.

```

      thick height breadth weight
1      14   30.5   23.0   1075
2      15   29.1   20.5    940
3      18   27.5   18.5    625
4      23   23.2   15.2    400

```

We will only consider relationships between the thickness, height, and width.

```
R> Z <- oddbooks[, c(1, 2, 3)]
```

One can suspect that books from this dataset were printed with a $\sqrt{2}$ aspect ratio, as in the popular A-series paper size. Therefore, we can utilize this domain knowledge in the analysis and unify the data for height and width:

```
R> Z$height <- Z$height / sqrt(2)
```

Let us see the standard MLE of the covariance matrix:

```
R> S <- cov(Z)
```

We can plot this covariance matrix to see if we would notice any connection between variables. Figure 3 was obtained with the code below (we made cosmetic modifications to this plot; the exact code is provided in the replication code available online):

```
R> g <- gips(S, number_of_observations)
R> plot(g, type = "heatmap")
```




Figure 4: Estimator of the covariance matrix of the `oddbooks` dataset after the projection on the permutational symmetry group (2,3). We can see that now similar colors from Figure 3 are exactly equal.

We can see that some entries of S have similar colors, which suggests a lower dimensional model with equality constraints. In particular, the covariance between `thick` and `height` is very similar to the covariance between `thick` and `breadth`, and the variance of `height` is similar to the variance of `breadth`. Those are not surprising, given the data interpretation (after the `height` rescaling that we did).

Let us examine the posterior probabilities returned by `gips`:

```
R> g_MAP <- find_MAP(g, optimizer = "BF", return_probabilities = TRUE,
+   save_all_perms = TRUE)
R> get_probabilities_from_gips(g_MAP)
```

```
          (2,3)          ()          (1,3)          (1,2,3)          (1,2)
5.660781e-01 4.339087e-01 6.728772e-06 4.683290e-06 1.862353e-06
```

We see that the a posteriori distribution is maximized by a permutation (2,3). The MLE of the covariance matrix in the model invariant under the permutation (2,3) is presented in Figure 4.

1.3. Motivation behind permutation symmetries

We argue that it is natural to expect certain symmetries in various applications, which strengthens the need for tools to investigate permutation symmetry within the data.

For example, there are natural symmetries in the data from gene expression. Specifically, the expression of a given gene is triggered by the binding of transcription factors to gene

transcription factor binding sites. Transcription factors are proteins produced by other genes, often referred to as regulatory genes. Within the gene network, it is common for multiple genes to be triggered by the same regulatory genes, suggesting that their relative expressions depend on the abundance of the regulatory proteins (i.e., gene expressions) in a similar manner (Graczyk *et al.* 2022b). Extracting permutation symmetries can be utilized to identify genes with similar functions or groups of genes with similar interactions or regulatory mechanisms. This approach is particularly useful in unraveling the structures of gene regulatory networks (Kotiang and Eslami 2020).

Furthermore, in examples of social networks, such as those influenced by geographical or social group clusters, additional symmetries must be taken into account, as mentioned in Gao and Massam (2015). In the study of the human brain's dynamics, it is believed that the left and right hemispheres possess a natural symmetric structure (Ranciati *et al.* 2021).

The discovery of hidden symmetries can greatly contribute to understanding complex mechanisms. Extracting patterns from gene expression profiles can offer valuable insights into gene function and regulatory systems (Toh and Horimoto 2002). Clustering genes based on their expression profiles can aid in predicting the functions of gene products with unknown purposes and identifying sets of genes regulated by the same mechanism.

1.4. Arbitrary permutation symmetries vs cyclic permutation symmetries

As observed in Graczyk *et al.* (2022b), performing model selection within an arbitrary permutation subgroup is a highly challenging task. This difficulty arises not only due to theoretical reasons but also because of computational complexity issues arising when p is large. Informally speaking, finding the parameters of an arbitrary permutation group becomes virtually impossible for large values of p . In Graczyk *et al.* (2022b), a general model was developed; however, it was specifically applied to cyclic subgroups. Such subgroups are generated by a single permutation, and by restricting the analysis to them, efficient methods can be devised to conduct the model selection procedure. All the technical details regarding these methods will be presented in the subsequent sections.

Furthermore, we argue that cyclic subgroups form a sufficiently rich family, as mentioned in Graczyk *et al.* (2022b, Section 4.1). Since these subgroups correspond to simpler symmetries, they are also more easily interpretable. Although our procedure exclusively explores cyclic subgroups, it can still provide valuable information even when the true subgroup is not cyclic, as discussed in Graczyk *et al.* (2022c, Section 3.3). In fact, if the posterior probabilities (which are calculated with **gips**) are high for multiple groups, it is reasonable to expect that the data will exhibit invariance under the group containing those subgroups. We present a simple example in the Appendix A.

2. Methodological background

After providing an informal introduction, let us proceed to define the key concepts and present the theory behind the **gips** package in a formal manner. Definitions in this section are accompanied with code in **gips** package. The running example is for $p = 5$ and $n = 10$. A standard PC can execute all the code in this section within 20 seconds (except for the final chunk of code in Section 2.6 which runs for 2 minutes).

```
R> p <- 5
R> n <- 10
```

2.1. Permutations

Fix $p \in \{1, 2, \dots\}$. Let \mathfrak{S}_p denote the symmetric group, the set of all permutations on the set $V = \{1, \dots, p\}$, with function composition as the group operation.

Each permutation $\sigma \in \mathfrak{S}_p$ can be represented in a cyclic form. For example, if σ maps 1 to 2, 2 to 1, and leaves 3 unchanged, then we can express σ as $(1, 2)(3)$. It is sometimes convenient to exclude cycles of length 1 from this representation. The identity permutation is denoted as id or $()$. The number of cycles denoted as C_σ , remains the same across different cyclic representations of σ . It is important to note that cycles of length 1 are included when calculating C_σ .

We say that a permutation subgroup $\Gamma \subset \mathfrak{S}_p$ is cyclic if $\Gamma = \{\sigma, \sigma^2, \dots, \sigma^N\} =: \langle \sigma \rangle$ for some $\sigma \in \mathfrak{S}_p$, where N is the smallest positive integer such that $\sigma^N = \text{id}$. Then, N is the order of the subgroup Γ . If p_i denotes the length of the i th cycle in a cyclic decomposition of $\sigma \in \mathfrak{S}_p$, then N is equal to the least common multiple of $p_1, p_2, \dots, p_{C_\sigma}$.

If $\Gamma = \langle \sigma \rangle$, then we say that σ is a generator of Γ . It is worth noting that a cyclic subgroup may have several generators. Specifically, $\langle \sigma \rangle = \langle \sigma^k \rangle$ for all $k = 1, \dots, N - 1$, where k is coprime with N . We identify each cyclic permutation subgroup by its generator, which is the smallest permutation according to lexicographic order. For further topics on permutation groups, readers may refer to [Gallian \(2017\)](#).

2.2. Permutation symmetry

Let Γ be an arbitrary subgroup of \mathfrak{S}_p . We say that the distribution of $Z = (Z_i)_{i \in V}$ is invariant under a subgroup Γ if Z has the same distribution as $(Z_{\sigma(i)})_{i \in V}$ for all $\sigma \in \Gamma$. If Z is a multivariate random variable following a centered Gaussian distribution $N_p(0, \Sigma)$, then this invariance property can be expressed as a condition on the covariance matrix. Specifically, the distribution of Z is invariant under Γ if and only if for all $i, j \in V$:

$$\Sigma_{ij} = \Sigma_{\sigma(i)\sigma(j)} \quad \text{for all } \sigma \in \Gamma. \quad (1)$$

When $\Gamma = \mathfrak{S}_p$, the above conditions imply that all diagonal entries of Σ are the same, and similarly, the off-diagonal entries are the same (see the top left panel of Figure 5). On the other hand, if Γ is the trivial subgroup, i.e., $\Gamma = \{\text{id}\}$, then (1) does not impose any restrictions on the entries of Σ . If Γ is non-trivial, the sample size n required for the MLE to exist is lower than p , as discussed in Section 2.4.


Let $\text{Sym}(p; \mathbb{R})$ and $\text{Sym}^+(p; \mathbb{R})$ denote the space of $p \times p$ symmetric matrices and the corresponding cone of positive definite matrices, respectively. For a subgroup $\Gamma \subset \mathfrak{S}_p$, we define the colored space as the space of symmetric matrices invariant under Γ :

$$\mathcal{Z}_\Gamma := \{S \in \text{Sym}(p; \mathbb{R}) : S_{ij} = S_{\sigma(i)\sigma(j)} \text{ for all } \sigma \in \Gamma\},$$

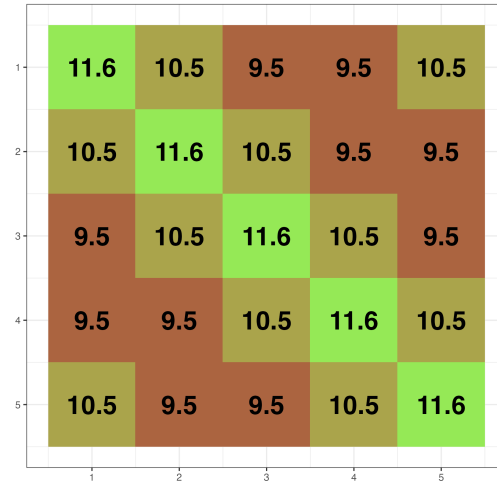
We also define the colored cone of positive definite matrices valued in \mathcal{Z}_Γ as:

$$\mathcal{P}_\Gamma := \mathcal{Z}_\Gamma \cap \text{Sym}^+(p; \mathbb{R}).$$


Example of a matrix **with full** permutation symmetry



Example of a matrix **with long** permutation symmetry



Example of a matrix **with short** permutation symmetry



Example of a matrix **without** permutation symmetry




Figure 5: Examples of covariance matrices with symmetries: Full $\Gamma = \mathfrak{S}_p$ permutation symmetry (top left), long permutation symmetry $\Gamma = \langle(12345)\rangle$ (top right), short permutation symmetry $\Gamma = \langle(123)\rangle$ (bottom left), no permutation symmetry (bottom right). They all are projections of the same matrix on corresponding colored spaces.

The set \mathcal{P}_Γ contains all possible covariance matrices of Gaussian vectors invariant under subgroup Γ . The dimension of the space \mathcal{Z}_Γ corresponds to the number of free parameters in the covariance matrix. The dependence structure of a Gaussian vector is fully described by the covariance matrix Σ . When certain entries of Σ are identical, we refer to them as having the same color. There are colors that correspond to equalities among the diagonal elements of Σ , and there are independent colors that correspond to equalities among the off-diagonal elements of Σ . Thus, in the context of colored models, $\dim(\mathcal{Z}_\Gamma)$ can be interpreted as the number of distinct colors.

In **gips**, we can easily find the number of free parameters in the model invariant under a cyclic subgroup as follows (S is a matrix from the bottom right of Figure 5):

```
R> g <- gips(S, n, perm = "(12345)", was_mean_estimated = FALSE)
R> summary(g)$n_parameters
```

3

Notice there were exactly 3 different numbers in the matrix from the top right of Figure 5.

It is important to note that the mapping $\Gamma \mapsto \mathcal{Z}_\Gamma$ is not one-to-one. In particular, for $p = 3$, we have $\mathcal{Z}_{\langle(1,2,3)\rangle} = \mathcal{Z}_{\mathfrak{S}_3}$. A notable property of cyclic subgroups is that they correspond to different colored spaces. More precisely, if $\mathcal{Z}_{\langle\sigma\rangle} = \mathcal{Z}_{\langle\sigma'\rangle}$ for some $\sigma, \sigma' \in \mathfrak{S}_p$, then $\langle\sigma\rangle = \langle\sigma'\rangle$ (Graczyk *et al.* 2022b, Lemma 15).

2.3. Recovering useful information from permutation symmetry

The permutation symmetries can cover a very complicated relationships, even when restricting to cyclic subgroups. It is therefore of interest to extract useful information from a given permutation symmetry, especially when the number of variables p is large. Let p_1, \dots, p_{C_σ} be the sizes of cycles in a cycle representation of σ . The order N of $\Gamma = \langle\sigma\rangle$ equals the least common multiple of p_1, \dots, p_{C_σ} .

If N is divisible by 2, then the permutation $\sigma^{N/2}$ consists of two-cycles only, i.e.,

$$\sigma^{N/2} = (i_1, j_1) \dots (i_k, j_k).$$

Define two sets of indices $I = (i_1, \dots, i_k)$ and $J = (j_1, \dots, j_k)$. Let $Z_I = (Z_i)_{i \in I}$, $Z_J = (Z_j)_{j \in J}$ and $Z_{\text{rest}} = (Z_i)_{i \in V \setminus (I \cup J)}$. If the distribution of $Z = (Z_i)_{i \in V}$ is invariant under Γ , then

$$(Z_I, Z_J, Z_{\text{rest}}) \stackrel{d}{=} (Z_J, Z_I, Z_{\text{rest}}),$$

meaning that the vectors Z_I and Z_J play the same role in a model.

More generally, if N is divisible by $\ell > 1$, then there exists $k = k_\ell \in \mathbb{N}$ such that

$$\sigma^{N/\ell} = (i_1^1, i_2^1, \dots, i_\ell^1) \dots (i_1^k, i_2^k, \dots, i_\ell^k).$$

Define $Z^\alpha = (Z_i)_{i=i_\alpha^1, \dots, i_\alpha^k}$ for $\alpha = 1, \dots, \ell$. Then, the Γ -invariance gives us

$$(Z^1, Z^2, \dots, Z^\ell, Z_{\text{rest}}) \stackrel{d}{=} (Z^2, Z^3, \dots, Z^1, Z_{\text{rest}}) \stackrel{d}{=} \dots \stackrel{d}{=} (Z^\ell, Z^1, \dots, Z^{\ell-1}, Z_{\text{rest}}),$$

where Z_{rest} are the remaining components of Z .

A concrete example of the application of this idea can be found in the Section 3.1.

2.4. MLE in the Gaussian model invariant under permutation symmetry

Let $Z^{(1)}, \dots, Z^{(n)}$ be an independent and identically distributed (i.i.d.) sample from $N_p(0, \Sigma)$. The presence of equality restrictions in (1) reduces the number of parameters to estimate in permutation invariant models. Consequently, the sample size required for the MLE of Σ to exist is lower than p for non-trivial subgroup $\Gamma \subset \mathfrak{S}_p$. Assuming $\Sigma \in \mathcal{P}_\Gamma$, where $\Gamma = \langle\sigma\rangle$ is a

cyclic subgroup, Graczyk *et al.* (2022b, Corollary 12) establishes that the MLE of Σ exists if and only if

$$n \geq n_0 := C_\sigma. \quad (2)$$

In particular, when $\sigma = \text{id}$, no restrictions are imposed on Σ , and we recover the well-known condition that the sample size n must be greater than or equal to the number of variables $p = C_{\text{id}}$. However, if σ consists of a single cycle, i.e., $C_\sigma = 1$, the MLE always exists. This remarkable observation is crucial in high-dimensional settings.

In **gips**, we can compute n_0 as follows:

```
R> g <- gips(S, n, perm = "(12345)", was_mean_estimated = FALSE)
R> summary(g)$n0
```

1

```
R> g <- gips(S, n, perm = "()", was_mean_estimated = FALSE)
R> summary(g)$n0
```

5

If (2) is satisfied, the MLE of Σ is given by

$$\hat{\Sigma} = \pi_\Gamma \left(\frac{1}{n} \sum_{i=1}^n Z^{(i)} \cdot (Z^{(i)})^\top \right),$$

where π_Γ denotes the orthogonal projection onto the colored space \mathcal{Z}_Γ . It is defined as

$$\pi_\Gamma(X) = \frac{1}{\#\Gamma} \sum_{\sigma \in \Gamma} \sigma \cdot X \cdot \sigma^\top,$$

where each permutation σ is identified with its corresponding permutation matrix. We note that if X is positive definite, then $\pi_\Gamma(X)$ is also positive definite. The resulting matrix $\pi_\Gamma(X)$ is often referred to as the regularized matrix since the mapping averages the entries of X that correspond to the same orbits of Γ : for $\{i, j\} \subset V$ define its Γ -orbit by $O_{ij}^\Gamma = \{\{\sigma(i), \sigma(j)\} : \sigma \in \Gamma\}$. Then, for any $\{u, v\} \in O_{ij}^\Gamma$ one has

$$\pi_\Gamma(X)_{uv} = \frac{1}{\#O_{ij}^\Gamma} \sum_{\{k, l\} \in O_{ij}^\Gamma} X_{kl}. \quad (3)$$

In **gips**, the projection $\pi_{(\text{perm})}(\mathbf{S})$ of a matrix \mathbf{S} onto $\mathcal{Z}_{(\text{perm})}$ is calculated as follows:

```
R> S_projected <- project_matrix(S, perm)
```

where `perm` can be the permutation of a form "(12345)", or an object of a 'gips' class.

2.5. Bayesian model selection procedure

Now we shift our focus to methods aimed at discovering permutation symmetries in the data. The model introduced in Graczyk *et al.* (2022b) is considered. In this model, the multivariate

Gaussian sample $Z^{(1)}, \dots, Z^{(n)}$ given $\{K = k, \Gamma = c\}$ consists of i.i.d. $N_p(0, k^{-1})$ random vectors.

Let Γ be a discrete random variable uniformly distributed over the set $\mathcal{C} := \{\langle \sigma \rangle : \sigma \in \mathfrak{S}_p\}$ of cyclic subgroups of \mathfrak{S}_p . It is assumed that K given $\{\Gamma = c\}$ follows the Diaconis-Ylvisaker conjugate prior (Diaconis and Ylvisaker 1979) distribution, defined by its density

$$f_{K|\Gamma=c}(k) = \frac{1}{I_c(\delta, D)} \text{Det}(k)^{(\delta-2)/2} e^{-\frac{1}{2} \text{Tr}[D \cdot k]} \mathbf{1}_{\mathcal{P}_c}(k),$$

where $\delta > 1$ and $D \in \mathcal{P}_c$ are the hyperparameters, and $I_c(\delta, D)$ is the normalizing constant.

It was derived in Graczyk *et al.* (2022b) that the posterior probability is proportional to

$$\mathbb{P}(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)}) \propto \frac{I_c(\delta + n, D + U)}{I_c(\delta, D)}, \quad c \in \mathcal{C}, \quad (4)$$

where $U = \sum_{i=1}^n Z^{(i)} \cdot (Z^{(i)})^\top$. In order to utilize quotient (4), it is necessary to calculate or approximate the ratios of the normalizing constants. An efficient method for calculating these constants for cyclic subgroups was introduced in Graczyk *et al.* (2022b). This method relies on the block decomposition of the colored space \mathcal{Z}_Γ and is implemented in the **gips** package. Further technical details are provided in Appendix B.

In **gips**, one can calculate the quotient on the right-hand side of quotient (4) for permutation $c = \langle (12345) \rangle$ as follows:

```
R> g <- gips(S, n, perm = "(12345)", was_mean_estimated = FALSE)
R> exp(log_posteriori_of_gips(g))
```

```
4.586821e-27
```

This is a very small number, but keep in mind that the posteriori probability of a subgroup c is proportional to this quantity (not equal). One can compare it with other subgroups to get interpretable result (under our Bayesian setting):

```
R> compare_posteriories_of_perms(g, "(123)")
```

```
The permutation (1,2,3,4,5) is 22.827 times more likely
than the (1,2,3) permutation.
```

Following the Bayesian paradigm, we work with the maximum a posteriori (MAP) estimator, which corresponds to the cyclic subgroup with the highest posterior probability, i.e., this estimator is defined as

$$\hat{\Gamma} = \arg \max_{c \in \mathcal{C}} \mathbb{P}(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)}). \quad (5)$$

While the choice of hyperparameters is not scale invariant, it is a common practice in similar models to set $\delta = 3$ and $D = I_p$ (Massam *et al.* 2018). The parameter $\delta = 3$ serves as the default parameter in our method, but we decided to set $D = \frac{\text{tr}(S)}{p} \cdot I_p$ as default and justify our choice in Section 2.7. In **gips**, one can pass the desired values of these parameters via `delta` and `D_matrix` arguments in `gips()` function.

Algorithm 1 Metropolis-Hastings algorithm

Let $T \in \mathbb{N}$.

Let σ_0 be an arbitrary permutation from \mathfrak{S}_p .

for $t = 1, 2, \dots, T$ **do**

 Sample x_t uniformly from the set \mathcal{T} and set $\sigma' = \sigma_{t-1} \circ x_t$.

 Accept the move $\sigma_t = \sigma'$ with probability

$$\min \left\{ 1, \frac{I_{\langle \sigma' \rangle}(\delta + n, D + U) \cdot I_{\langle \sigma_{t-1} \rangle}(\delta, D)}{I_{\langle \sigma' \rangle}(\delta, D) \cdot I_{\langle \sigma_{t-1} \rangle}(\delta + n, D + U)} \right\}.$$

 If the move is rejected, set $\sigma_t = \sigma_{t-1}$.

end for

In Section 3.2, we considered the influence of δ and d in $D = d \cdot I_p$. The role of the d parameter turns out to be quite similar to the role of the tuning parameters λ in LASSO (least absolute shrinkage and selection operator) methods. To conclude, smaller values of d tend to favor large symmetries. Therefore, through such exploratory analysis, users can adjust the parameter d to achieve a model that aligns most meaningfully with their preferences and requirements.

2.6. Searching for a MAP estimator

The quotient (4) enables the numerical evaluation of how well a given permutation symmetry (specifically, a cyclic group generated by a permutation) fits the data. Finding a cyclic subgroup with a high evaluation score is a challenging task for large values of p due to the vast size of the space of potential permutation symmetries.

Recall that \mathcal{C} is the space of cyclic subgroups of \mathfrak{S}_p . For small values of p (in our investigation, up to 9), it is possible to compute the posterior quotients (4) for all $c \in \mathcal{C}$ and determine $\hat{\Gamma}$ from (5) using exact calculations, i.e., for $c \in \mathcal{C}$ we have

$$\mathbf{P}(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)}) = \frac{I_c(\delta + n, D + U)}{I_c(\delta, D) \sum_{s \in \mathcal{C}} \frac{I_s(\delta + n, D + U)}{I_s(\delta, D)}}. \quad (6)$$

However, the cardinality of \mathcal{C} grows super-exponentially with p . Specifically, for $p = 150$, the cardinality of \mathcal{C} is approximately 10^{250} . This makes it computationally infeasible to calculate the quotients (4) for all $c \in \mathcal{C}$.

To address this challenge, we propose the use of a Monte Carlo Markov chain method. We define an irreducible Markov chain $(\sigma_t)_t$ that traverses an even larger space, \mathfrak{S}_p , and apply the Metropolis-Hastings algorithm to obtain preliminary estimates of the posterior probabilities. Subsequently, taking into account the fact that some permutations generate the same cyclic subgroup, we derive the estimates of the posterior probabilities using (7) below. Through the ergodic theorem, the Metropolis-Hastings algorithm provides statistical guarantees that the estimates will converge to the true values as the number of iterations tends to infinity.

A transposition is a permutation that swaps two elements while leaving other elements unchanged. In other words, each transposition is in the form (i, j) for some $i, j \in V$ where $i \neq j$. Let \mathcal{T} denote the set of all transpositions.

The algorithm produces a sequence of permutations $(\sigma_t)_{t=1}^T$, and we construct a corresponding sequence of cyclic subgroups $(\langle\sigma_t\rangle)_{t=1}^T$. The MAP estimator, which corresponds to the cyclic group with the highest posterior probability, is given by

$$\hat{\Gamma} = \arg \max_{c \in (\langle\sigma_t\rangle)_{t=1}^T} \mathbf{P} \left(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)} \right) = \arg \max_{c \in (\langle\sigma_t\rangle)_{t=1}^T} \frac{I_c(\delta + n, D + U)}{I_c(\delta, D)}.$$

where the maximum is taken over all permutations visited by the Markov chain constructed in the algorithm.

The Algorithm 1 is implemented in `find_MAP()` function with parameter `optimizer = "MH"`:

```
R> g <- gips(S, n, was_mean_estimated = FALSE)
R> g_MAP_MH_25 <- find_MAP(g, max_iter = 25, optimizer = "MH")
R> g_MAP_MH_25
```

The permutation (1,2,3,5):

- was found after 25 posteriori calculations;
- is 5.149 times more likely than the () permutation.

The algorithm in 25 steps found a quite long permutation. Always keep in mind the Algorithm 1 is only an approximate one. However, if one wants to have the true MAP, a brute-force search can be applied:

```
R> g_MAP_BF <- find_MAP(g, optimizer = "BF")
R> g_MAP_BF
```

The permutation (1,2,3,4,5):

- was found after 67 posteriori calculations;
- is 33.743 times more likely than the () permutation.

```
R> compare_posteriories_of_perms(g_MAP_BF, g_MAP_MH_25)
```

The permutation (1,2,3,4,5) is 6.553 times more likely than the (1,2,3,5) permutation.

If one is interested in estimating $\mathbf{P} \left(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)} \right)$ for an arbitrary $c \in \mathcal{C}$, the following approach can be used. For a permutation subgroup Γ , let $\varphi(\Gamma) = \Phi(\#\Gamma)$, where Φ is the Euler totient function, i.e., $\Phi(n) = \#\{k \in \{1, \dots, n\} : k \text{ and } n \text{ are coprime}\}$. In [Graczyk et al. \(2022b, Theorem 16\)](#), it is shown that as $T \rightarrow \infty$ and for $c \in \mathcal{C}$,

$$\hat{\pi}_c := \frac{\sum_{t=1}^T \mathbf{1}(\langle\sigma_t\rangle = c)}{\varphi(\langle c \rangle) \sum_{t=1}^T 1/\varphi(\langle\sigma_t\rangle)} \xrightarrow{a.s.} \mathbf{P} \left(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)} \right). \quad (7)$$

In practice, $\hat{\pi}_c$ serves as an approximation to $\mathbf{P} \left(\Gamma = c \mid Z^{(1)}, \dots, Z^{(n)} \right)$ for large T .

By default, **gips** does not save all the computed permutations, but only the best one. One can set the flag `save_all_perms = TRUE` to get the desired exact distribution:

```
R> g_MAP_BF_with_probs <- find_MAP(g, optimizer = "BF",
+   save_all_perms = TRUE, return_probabilities = TRUE)
R> head(get_probabilities_from_gips(g_MAP_BF_with_probs), 10)

(1,2,3,4,5)  (1,4)(2,3)  (1,2)(3,5)  (1,3)(4,5)  (1,2,4,5,3)
0.13478976  0.05532886  0.05487531  0.04410568  0.03076733
(1,5)(2,4)  (1,2,5,3)  (1,2,3,4)  (1,2,4)(3,5)  (1,3,4,5)
0.02969751  0.02928655  0.02692967  0.02638740  0.02501291
```

If one wants to estimate the distribution, e.g., when p is too large to search through entire space, one can do exactly the same with the `optimizer = "MH"`:

```
R> g_MAP_MH_20000 <- find_MAP(g, optimizer = "MH", max_iter = 20000,
+   save_all_perms = TRUE, return_probabilities = TRUE)
R> head(get_probabilities_from_gips(g_MAP_MH_20000), 10)

(1,2,3,4,5)  (1,4)(2,3)  (1,2)(3,5)  (1,3)(4,5)  (1,5)(2,4)
0.13832762  0.06678977  0.05718808  0.04167766  0.03102084
(1,2,3,4)  (1,2,5,3)  (1,2,4,5,3)  (1,2,4)(3,5)  (1,3,4,5)
0.02970193  0.02964917  0.02930625  0.02653653  0.02495384
```

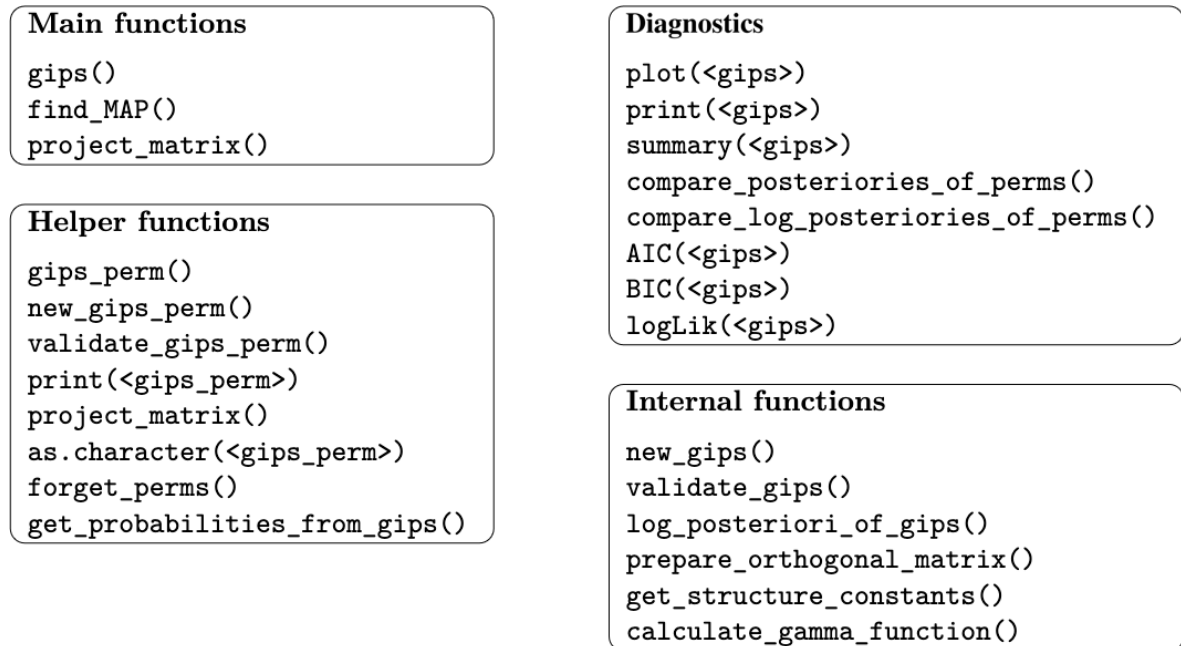
We can observe that the estimated probabilities are similar to the true ones. However, please note that the above code is intended solely to demonstrate convergence, and in practical scenarios, it is unreasonable to execute the function `find_MAP(optimizer = "MH", max_iter = my_max_iter)` for $\text{my_max_iter} > p!$ due to the faster and exact brute-force `find_MAP(optimizer = "BF")`. The `gips` package will show a warning in such a case.

2.7. Scaling, centering and standardizing data

We want to emphasize that the considered model of permutation symmetries is not scale-invariant in the following sense: If Z is invariant under a subgroup, a random vector $\text{diag}(\alpha) \cdot Z$, where $\alpha \in \mathbb{R}^p$, is generally not invariant under any permutation subgroup. Therefore, it is recommended to apply our procedure to data that have comparable scales and keep all variables in the same units. That is the reason for the division of `height` by $\sqrt{2}$ at the beginning of example *Books dataset* from Section 1.2.

It is worth noting that there are many examples of such data, such as gene expression data, where measurements are on the same scale due to being results of experiments of the same type and measured using the same gauges. For further references, see e.g., [Gehrmann \(2011\)](#); [Gao and Massam \(2015\)](#); [Massam et al. \(2018\)](#); [Li et al. \(2020, 2021\)](#); [Ranciati et al. \(2021\)](#). However, our model is scale-invariant under common scaling, i.e., if Z is invariant under Γ , then βZ for any $\beta \in \mathbb{R}$ is also invariant under Γ . Our practice shows that choosing β in a way that βZ has *average* unit variance often produces good results. Such scaling can be equivalently accomplished by choosing the hyperparameter $D = \frac{\text{tr}(S)}{p} \cdot I_p$, where S is the empirical covariance matrix of Z . This is the default parameter for `D_matrix` in `gips()` function.

While our Bayesian model is designed for a zero mean Gaussian sample, it can be easily extended to handle samples with arbitrary means. If $Z^{(1)}, \dots, Z^{(n)}$ is an i.i.d. sample from

Figure 6: The structure of the **gips** package.

$N_p(\mu, \Sigma)$, the user can center the data and take this into account by setting the parameter `was_mean_estimated = TRUE` in function `gips()`.

In cases where the sample size n is reasonably large, it is common to assume that the standardized normal sample (which follows a multivariate t -distribution) can be approximated by a Gaussian distribution. Therefore, for large n , one can standardize each variable and apply our model selection procedure to obtain reliable estimates. However, it is important to note that after standardization, the empirical covariance matrix will have a unit diagonal, which may favor cyclic subgroups whose generators consist of a single cycle, as they correspond to matrices with a constant diagonal. This is the reason why we do not recommend standardizing the data when the sample size n is small.

3. Package structure and usage

The package **gips** is available under the general public license (GPL ≥ 3) from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=gips>. Its documentation is available as **pkgdown** page at <https://przechoj.github.io/gips>. **gips** can be installed and loaded into the current R session using the following code:

```
R> install.packages("gips")
R> library("gips")
```

The primary use case of the **gips** package is to find a permutation subgroup with the maximum a posteriori probability in the Bayesian model introduced in the previous sections and estimate the covariance matrix in the model invariant under this permutation subgroup. Representations and operations on permutations are performed using the **permutations** package

(Hankin 2020). We decided to use this package due to its ease of transforming permutations and its compactness.

We start the description of the main functions implemented in the proposed package. The workflow in **gips** is as follows: First, use the `gips()` function to define an object of the ‘**gips**’ class that contains all the necessary information for the model.

```
R> g <- gips(S, number_of_observations, delta = 3, D_matrix = NULL,
+   was_mean_estimated = TRUE, perm = "")
```

Let Z be a data matrix of size `number_of_observations` \times p . The parameter `S` is the $p \times p$ empirical covariance matrix of Z and `number_of_observations` is the corresponding sample size. If one does not know the theoretical mean of the distribution data was sampled from, use `S = cov(Z)` and leave the flag `was_mean_estimated = TRUE` as default. If the theoretical mean is known to be 0, use `S = crossprod(Z) / number_of_observations`, and set the flag `was_mean_estimated = FALSE`,

Parameters `delta` and `D_matrix` are the hyperparameters of our Bayesian model. The domains of these parameters are the following: `delta > 1` and `D_matrix` has to be $p \times p$ positive definite matrix. The default value of `D_matrix` is `mean(diag(S)) * diag(p)`. The last parameter, `perm`, is an optional permutation on p elements. Can be of any form that the function `permutations::permutation()` can handle. This is the starting permutation for the Metropolis-Hastings.

Next, use the `find_MAP()` function with an optimizer of your choice to find the permutation that provides the maximum a posteriori estimate.

```
R> find_MAP(g, max_iter = NA, optimizer = NA, show_progress_bar = TRUE,
+   save_all_perms = FALSE, return_probabilities = FALSE)
```

The first parameter, `g`, is the object of ‘**gips**’ class. There are three optimizers implemented:

`optimizer = "BF"`: Brute-force search (default and only intended for $p \leq 9$).

`optimizer = "MH"`: The Metropolis-Hasting algorithm (default and recommended for $p > 9$).

`optimizer = "HC"`: The hill climbing algorithm.

The parameter `max_iter` is the number of steps for the Metropolis-Hastings and the hill climbing algorithm. For the Metropolis-Hastings algorithm, it has to be finite and greater than 2, while for hill climbing it can be also `Inf`.

The progress bar of the optimization process can be turned on and off by changing the Boolean parameter `show_progress_bar`.

To obtain the entire posterior distribution, the flag `return_probabilities` has to be set to `TRUE`. This flag can only be provided when also `save_all_perms = TRUE`, which saves a list of all permutations that were visited during optimization. In the case of `optimizer = "BF"`, the exact posterior probabilities are calculated. For other optimizers, their estimates are calculated. One can access these probabilities using the function

```
R> get_probabilities_from_gips(g)
```

where `g` is the optimized ‘`gips`’ object. If one is interested only in the maximum a posteriori estimate, it is better to set `return_probabilities = FALSE` in the `find_MAP()` function.

Finally, to obtain the MLE of the covariance matrix in the invariant model found by `find_MAP`, one projects the empirical covariance matrix on a colored space corresponding to chosen permutation:

```
R> project_matrix(S, perm)
```

The first argument, `S`, is the $p \times p$ covariance matrix used in the `gips()` function. The `perm` is either a permutation or the ‘`gips`’ object that describes the cyclic permutation symmetry.

3.1. Real life example

We obtained the results in this section with AMD EPYC 7413 on a single core, which took under 2 hours to compute.

Let us present the capabilities of `gips` package using breast cancer data from [Miller *et al.* \(2005\)](#). Following the approach of [Højsgaard and Lauritzen \(2008\)](#), we consider a set of $p = 150$ genes and $n = 58$ samples with a mutation in the p53 sequence. We numbered the variables alphabetically. Since $p > n$, only parsimonious models can be fitted at all. Data is available in `GEOQuery` package from Bioconductor. Code for downloading and minimal preprocessing is available in the replication code available online. We stress that the model space to search is here very large. It can be roughly estimated to be of magnitude 10^{250} .

```
R> Z <- breast_cancer
R> dim(Z)
```

```
[1] 58 150
```

We note, that we have fewer observations than variables. Let us search for permutation symmetries. Create ‘`gips`’ object and run `find_MAP()` function on it. We set `D_matrix = diag(p)` as in [Graczyk *et al.* \(2022c, Section 4.1\)](#).

```
R> S <- cov(Z)
R> g <- gips(S, 58, D_matrix = diag(p), was_mean_estimated = TRUE)
R> set.seed(2022)
R> g_MAP <- find_MAP(g, max_iter = 150000, optimizer = "MH")
```

To acquire knowledge about the optimization process, we can call the `summary()` function on the object of the ‘`gips`’ class. We note that there is no guarantee that our solution via the Metropolis-Hastings algorithm is the true one and in order to get more reliable estimates one should increase the number of iterations.

```
R> summary(g_MAP)
```

The optimized ‘`gips`’ object.

Permutation:

(1,10,83,61,69,37,137,106) (2,42,19,16,43,49,24,82,34,139,140,52,26,98,17,100,
97,145) (3,9,11,71,120,101,126,76) (4,8,89) (5,148) (6,30,149,107,65,78,60,127)
(7,133,36,95) (12,103,92,146,138,144,84,62,58,77,111,122,66,129,93,59,41,81,35,
64,86,117,63,150,70,75,108) (13,50,57,132,114,22,116,125,74,72,91,90,113,130,
124) (14,110,46,29) (15,51,56,48,53,25,45,119) (18,68,99) (20,79,21) (23,131,27,67,
38,128,147,112,102) (28,73,44,135,105,96,104,39) (31,40,118,115,143) (32,33,123,
134,121,88) (47,109,94,136) (141,142)

Log_posteriori:

3626.114

Times more likely than starting permutation:

7.865e+549

The number of observations:

58

The mean in the `S` matrix was estimated.

Therefore, one degree of freedom was lost.

There are 57 degrees of freedom left.

n0:

25

The number of observations is bigger than n0 for this permutation,
so the gips model based on the found permutation does exist.

The number of free parameters in the covariance matrix:

611

BIC:

8741.694

AIC:

7482.764

Optimization algorithm:

Metropolis_Hastings

Number of log_posteriori calls:

150000

Optimization time:

1.428308 hours

Acceptance rate:

0.001953333333333333

Log_posteriori calls after the found permutation:
36814

The resulting permutation consists of $C_\sigma = 24$ cycles (including the cycles of length 1) and $n_0 = 25$. Recalling (2), this means that 25 observations are enough for the MLE to exist within such model and this is the case here as $n = 58 \geq n_0$. The dimension of the space $\mathcal{Z}_{\langle\sigma\rangle}$, i.e., the number of free parameters in the covariance matrix, is 611. The order of the group $\hat{\Gamma} = \langle\sigma\rangle$ is 1080 and following the principles from Section 2.3 we have

$$\sigma^{1080/2} = (1, 69)(3, 120)(6, 65)(9, 101)(10, 37)(11, 126)(15, 53)(25, 51)(28, 105)(30, 78)(39, 135) \\ (44, 104)(45, 56)(48, 119)(60, 149)(61, 106)(71, 76)(73, 96)(83, 137)(107, 127).$$

Define the two sets of indices:

$$I = (1, 3, 6, 9, 10, 11, 15, 25, 28, 30, 39, 44, 45, 48, 60, 61, 71, 73, 83, 107), \\ J = (69, 120, 65, 101, 37, 126, 53, 51, 105, 78, 135, 104, 56, 119, 149, 106, 76, 96, 137, 127).$$

The Γ -invariance tells us that

$$(Z_I, Z_J, Z_{\text{rest}}) \stackrel{d}{=} (Z_J, Z_I, Z_{\text{rest}})$$

and therefore one can say that $Z_I = (Z_i)_{i \in I}$ and $Z_J = (Z_i)_{i \in J}$ play the same role in the model. Similarly, we may analyze the action of

$$\sigma^{1080/3} = (12, 77, 35)(41, 75, 62)(58, 81, 108)(59, 70, 84)(63, 138, 129)(64, 103, 111)(66, 117, 146) \\ (86, 92, 122)(93, 150, 144),$$

which gives us

$$(Z^1, Z^2, Z^3, Z_{\text{rest}}) \stackrel{d}{=} (Z^2, Z^3, Z^1, Z_{\text{rest}}) \stackrel{d}{=} (Z^3, Z^1, Z^2, Z_{\text{rest}}),$$

where $Z^\alpha = (Z_i)_{i \in J_\alpha}$, $\alpha = 1, 2, 3$, with

$$J_1 = (12, 41, 58, 59, 63, 64, 66, 86, 93), \quad J_2 = (77, 75, 81, 70, 138, 103, 117, 92, 150), \\ J_3 = (35, 62, 108, 84, 129, 111, 146, 122, 144)$$

and $Z_{\text{rest}} = (Z_i)_{i \notin J_1 \cup J_2 \cup J_3}$.

We can interpret this result as an indication of hidden symmetry in genes and evidence that our procedure can be used as an exploratory tool for finding such symmetries.

We also carry out the heuristic procedure introduced in Graczyk *et al.* (2022b, Section 1.2) for finding a graphical model which is invariant under the above symmetry. We threshold the entries of the partial correlation matrix at the level $\alpha = 0.05608621$ and construct undirected graph $G = (V, E)$ with $V = \{1, \dots, p\}$ and $E = \{\{i, j\}: i, j \in V, i \neq j\}$ defined by

$$\{i, j\} \in E \quad \text{if and only if} \quad \frac{|k_{ij}|}{\sqrt{k_{ii}k_{jj}}} \geq \alpha,$$




Figure 7: The undirected colored graph corresponds to constructed RCOP model. The variances of vertices with the same color are equal, and the covariances of edges with the same color are also equal. Additionally, entries of the precision matrix corresponding to missing edges are 0.

where (k_{ij}) are the entries of the estimated precision matrix $\hat{K} = \hat{\Sigma}^{-1}$. The constructed dependency graph is depicted in Figure 7. The graph is non-decomposable, it has 3324 edges (compared to $p(p-1)/2 = 11175$ edges in the full graph with $p = 150$ vertices) and the size of its biggest clique is 21. Although not straightforward to identify in this example from a figure, a group of variables Z_I and Z_J play the same role in the model (as explained in Section 2.3) if there exists an automorphism of the colored graph that maps vertices from the set I to the set J and vice versa. Note that the vertices from the sets I and J are necessarily of the same color, but need not form a clique.

We found the MLE of the covariance matrix in the corresponding colored graphical model using the `ggmfit()` function from the `gRim` package. Note that the maximum likelihood equation for this model can be solved by first taking appropriate averages of the elements in the Wishart matrix (projecting the empirical covariance matrix onto the corresponding colored space) and then solving the equations for corresponding graphical Gaussian model without symmetry restrictions (Højsgaard and Lauritzen 2008).

By deleting edges, the number of parameters was reduced from 611 to 271, resulting in a decrease in the log-likelihood from -3130 to -3353 . This reduction leads to a lower BIC compared to the former model, decreasing from 8742 to 7807. These findings indicate that the simpler model provides a better description of the data and highlight the relevance of the entire approach.

3.2. Hyperparameters' influence

We obtained the results in this section with AMD EPYC 7413 on 24 cores, which took under 5 minutes to compute.




Figure 8: True covariance matrices corresponding to the three scenarios: Left panel – no structure \mathbf{S} , middle panel – moderate structure $\pi_{\langle(1,2,3,4)\rangle}(\mathbf{S})$, right panel – large structure $\pi_{\langle(1,2,\dots,8)\rangle}(\mathbf{S})$.

The Bayesian model introduced in Section 2.5 depends on two parameters of the a priori distribution, a scalar δ and a matrix D . In the following section, we present the effect of these hyperparameters on the a posteriori distribution. Despite having an explicit formula for the quotient (4), it is too complex to allow for direct analysis. Furthermore, this formula inherently depends on the data, which further complicates the study.

Therefore, drawing conclusions about the influence of hyperparameters on the method’s outcome is difficult and must be done with caution. They directly influence the shape of the a posteriori distribution and, therefore, change both the theoretical MAP and the difficulty of the optimization problem.

We consider only the low-dimensional setting $p = 8$ because only then are we able to efficiently calculate posterior probabilities for all cyclic subgroups. On a standard PC, it takes about 4 minutes to calculate the entire posterior distribution. Moreover, there is no rationale to suggest that the influence of the hyperparameters would be significantly different for larger p . Comparisons are conducted across three different scenarios.

First, we generate an empirical covariance matrix \mathbf{S} from the Wishart distribution on $\text{Sym}(p; \mathbb{R})$ with the scale parameter I_p and the shape parameter p . Then, the true covariance matrices for three scenarios are defined as the projections (recall (3)) of \mathbf{S} onto the spaces invariant under the following cyclic permutation subgroups:

No structure: $\langle \text{id} \rangle$.

Moderate structure: $\langle (1, 2, 3, 4) \rangle$.

Large structure: $\langle (1, 2, \dots, 8) \rangle$.

The number of free parameters (or just the dimension) of the model without structure is $p + p(p - 1)/2 = 36$, while the model with moderate structure has dimension 17, and the model with large structure has 5 dimensions. For each of these three scenarios, we simulate $n = 30$ samples from $N_p(0, \Sigma)$, where Σ is the true covariance matrix for a given scenario (see Figure 8).

In order to investigate the characteristics of the a posteriori distribution, we considered δ to be 3 and 30, and D of the form $d \cdot I_p$ for $d \in \{0.1, 1, 10, 100\}$. Recall that $\delta = 3$ and $D = I_p$ are the default parameters. Since there is no natural total ordering of cyclic subgroups, we

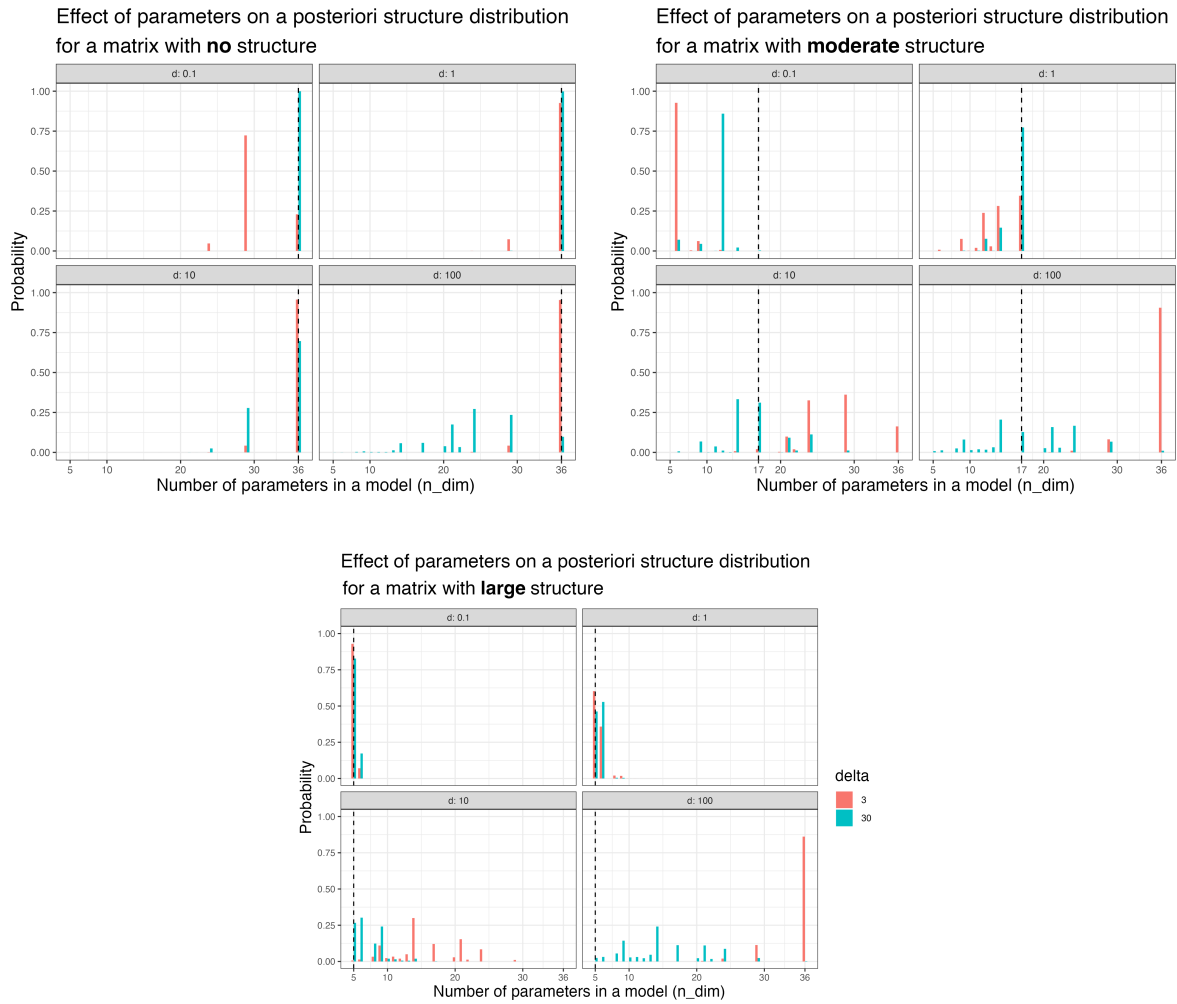


Figure 9: The posterior distribution (6) for $\delta \in \{3, 30\}$ and $D = d \cdot I_p$ for $d \in \{0.1, 1, 10, 100\}$ for three different scenarios. The black vertical line indicates the dimension of the true model.

identify each subgroup with the dimension of the model it generates, i.e., $\dim(\mathcal{Z}_\Gamma)$. In this way, for the sake of this analysis, all models that have the same dimension are merged.

The obtained distributions are shown at Figure 9.

Comments:

- The influence of parameter $D = d \cdot I_p$ is best reflected in the moderate structure scenario. As the value of parameter d decreases, larger symmetries (corresponding to lower-dimensional models) are preferred. Increasing the parameter δ shifts the distribution and reduces the difference between probabilities and therefore increases the entropy of the distribution.
- On the other hand, the posterior distribution for large values of parameter $D = d \cdot I_p$ becomes independent of the data. In fact, the plots for $d = 100$ are very similar across all the considered scenarios.

- In each of the scenarios, the Bayesian model correctly identifies the true model when $\delta = 3$ and $D = 1 \cdot I_p$. The value of the default parameter D for all scenarios is very similar, being equal to $0.77 \cdot I_p$. This demonstrates the model’s good properties for the default values of hyperparameters.

3.3. Comparison with other methods

We obtained the results in this section with AMD EPYC 7413 on 45 cores, which took around 1 hour 30 minutes to compute.

As mentioned in the Introduction, although there are no other software packages available for finding permutation symmetries in data, we have made the decision to compare the results of our model with canonical methods commonly used to tackle high-dimensional problems. In this section, we will compare method from **gips** package with methods implemented in **huge** (GLASSO) and **rags2ridges** (ridge) packages. Both **huge** and **rags2ridges** are based on matrix penalization and include a hyperparameter λ , which controls the strength of the penalty. They also both have implemented hyperparameter search techniques, which we will utilize.

Methodology

We conducted the comparison across different sample sizes and across strengths of the symmetry structure of true covariance matrices, similarly as in the previous section. For $p = 50$, we utilized matrices that are invariant under the following permutation subgroups:

No structure: $\langle \text{id} \rangle$.

Moderate structure: $\langle (1, 2, \dots, 25) \rangle$.

Large structure: $\langle (1, 2, \dots, 50) \rangle$.

For each of these scenarios, we constructed the true covariance matrices in the following way: First, we sampled a positive definite matrix from the Wishart distribution on $\text{Sym}(p; \mathbb{R})$ with a scale parameter of I_p and a shape parameter of p . We then projected this matrix onto the colored space corresponding to a given scenario. Next, we thresholded the inverse of this matrix by setting 25% of the off-diagonal entries with the smallest absolute values to zero. The inverse of such matrix served as our true covariance matrix. It is important to note that this approach does not always produce a positive definite matrix, which was indeed the case in the “no structure” scenario. In this particular case, we added $0.1 \cdot \text{diag}(p)$ to the realization of the Wishart distribution. This adjustment ensured the construction of a proper covariance matrix.

In the initial analysis, we also considered covariance matrices whose inverses did not contain any zeros. However, to our surprise, the results for matrices with and without zeros in their inverses were very similar for all the methods. Therefore, we decided to focus only on covariance matrices corresponding to nontrivial conditional dependence structures, as they are expected to favor likelihood penalization methods more.

For each of the scenarios, we considered three different sample sizes $n \in \{10, 20, 40\}$. Therefore, we have a total of $3 \cdot 3 = 9$ settings for this experiment. The comparison method for each setting is as follows:

1. Fix a sample size n and true covariance matrix Σ .
2. Generate a sample Z from $N_p(0, \Sigma)$ of size n .
3. Estimate the covariance matrix using:
 - **gips** package: `find_MAP()` function with `optimizer = "MH"` and `n_iter = 300000`,
 - **rags2ridges** package: `ridgeP()` function with `lambda` parameter found by the `optPenalty.kCVauto()` function with `lambdaMin = 0.001`, `lambdaMax = 100` range,
 - **huge**: Function `huge()` with parameters `method = "glasso"` and `nlambda = 40` and `lambda.min.ratio = 0.02` and the tuning parameter `lambda` was selected using default parameters of `huge.select()` function (rotation information criterion).
4. Record the log-likelihood and evaluate estimation using the Frobenius norm.
5. Repeat 2–4 10 times and aggregate results.

Recall that the Frobenius norm of a $p \times p$ matrix $M = (m_{ij})_{i,j}$ is defined by

$$\|M\|_F = \sqrt{\sum_{i=1}^p \sum_{j=1}^p |m_{ij}|^2}.$$

When M is a difference between the true covariance matrix and its estimate, $\|M\|_F^2$ is proportional to the mean squared error (MSE).

Results

Each of the three methods produces an estimator of the covariance matrix. We calculate the corresponding (negative) log-likelihood in our Gaussian model and present the results in Figure 10.

Comments:

- In all configurations, the **gips** method outperformed **huge**. However, it should be emphasized that the main purpose of the GLASSO method is the model selection within graphical models rather than the estimation of the covariance matrix. Typically, when possible, estimation is performed within the selected model and such an approach leads to systematically smaller bias.
- The comparison with **rags2ridges** is more interesting, as the **gips** method yielded weaker results when there was no structure and better results when the symmetry structure was large. This behavior was expected as **gips** is designed to look for these structures in the data.
- We can see that the results of **gips** were very unstable when there was no structure in the underlying ground truth matrix. This behavior is expected as **gips** will more likely find some non-existing structure when n is much smaller than p . Each method gains stability when the sample size increases.




Figure 10: Negative log-likelihoods of covariance matrix estimations. 10 runs for each configuration.




Figure 11: Frobenius norm (on a logarithmic scale) of the difference of the estimate and the true covariance matrix. 10 runs for each configuration.

The Frobenius norm of the difference of the estimate and the true covariance matrix is shown in Figure 11 with logarithmic Y axis.

Comments:

- When there is no symmetry structure in the data, the three methods considered generate estimates that are very similar in terms of the Frobenius norm.
- Generally, the bigger the symmetry structure, the better the quality of the estimation. Both **gips** and **huge** perform similarly in all scenarios, while **rags2ridges** performs significantly better in the scenario with a large structure.

We acknowledge that the proposed method of comparison is not systematic enough to draw conclusions in full generality. On one hand, the results act in support of the theory: **gips** method is a viable choice if we suspect, that the true matrix has some structure. On the other hand, it is difficult to recognize it post-hoc by comparing the method's performance using log-likelihood (or possibly other measures) usable in real-world cases, when the true covariance matrix is unknown.

From a practical perspective, it is worth repeating, that the **gips** method's output provides not only a projected covariance matrix but also an interpretation in the language of permutation symmetries of the data.

Finally, we note that both **huge** and **rags2ridges** methods execute within a few seconds, while it takes approximately 20 minutes to run the **gips** method for one scenario for $p = 50$ and 300 000 iterations of the Metropolis-Hastings algorithm.

4. Summary and discussion

In this paper, we have presented **gips**, an R package for learning permutation symmetry (i.e., identifying exchangeable variables) from Gaussian multivariate sample. Its usage revolves around setting up the main object of class 'gips' and then calling main functions `find_MAP` and `project_matrix`. Implementations of methods well known to R users, such as `summary` or `plot`, are also available. The applicability of **gips** spans a diverse array of scenarios where permutation symmetries are anticipated, as detailed in Section 1.3. It also proves advantageous in a high-dimensional setting, where conventional methods fail.

Our model provides interpretable results which are competitive to popular dimensionality reduction and covariance matrix estimation methods in Gaussian models. It excels when the underlying probability distribution has a rich symmetry structure, since it is able to detect and use that characteristic directly. We emphasize that, in the field of model selection methods for colored graphs, there is currently no competition for our package. Known methods are, to the best of our knowledge, not implemented in a publicly available package and the model presented in Graczyk *et al.* (2022b) and implemented in **gips** is the only one that allows for the search of permutation symmetries.

The **gips** package is under active maintenance and will continue to be developed to incorporate more advanced features. One potential avenue for future development is the inclusion of a model selection procedure within Gaussian graphical models that are invariant under permutation symmetry (RCOP models). By providing this package, we intend to stimulate the discussion about the concept of permutation symmetries and the investigation of its potential applications for statisticians and the R community. We invite everyone to a discussion about potential directions of development at <https://github.com/PrzeChoj/gips/issues>.

Computational details

The replication code is available at <https://jstatsoft.org> and on GitHub at https://github.com/PrzeChoj/gips_replication_code.

The results in this paper were obtained using R 4.4.1 with the **gips** 1.2.3 package. R itself and most packages used are available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/>. The only exception is package **GEOquery** 2.66.0 (Davis

and Meltzer 2007) that was used to obtain the data for Section 3.1. This package is available on **Bioconductor** repository (Gentleman *et al.* 2004) at <http://www.Bioconductor.org/>. For **gips**'s dependencies, we used **numbers** 0.8-5 (Borchers 2022), **permutations** 1.1-2 (Hankin 2020), **rlang** 1.1.1 (Henry and Wickham 2023), and **utils** 4.2.2.

For packages in Section 3.3, we used **rags2ridges** 2.2.6, **huge** 1.3.5. The remaining packages used are **Biobase** 2.58.0, **GEOquery** 2.66.0, **BiocManager** 1.30.21, **MASS** 7.3-60 (Venables and Ripley 2002), **ggplot2** 3.4.2 (Wickham 2016), **magrittr** 2.0.3 (Bache, Wickham, and Henry 2022), **parallel** 4.2.2, **dplyr** 1.1.2 (Wickham, François, Henry, Müller, and Vaughan 2023), **stringi** 1.7.12 (Gągolewski 2022), **gRim** 0.2.10 (Højsgaard, Edwards, and Lauritzen 2012). For producing the Figure 7 we used **Cytoscape** v3.10.0 (Shannon *et al.* 2003).

Acknowledgments

Research was funded by (POB Cybersecurity and Data Science) of Warsaw University of Technology within the Excellence Initiative: Research University. This research was carried out with the support of the Laboratory of Bioinformatics and Computational Genomics and the High Performance Computing Center of the Faculty of Mathematics and Information Science Warsaw University of Technology. The authors thank the reviewer for remarks that led to the improvement of the interpretation of the results of **gips**.

References

- Andersson S (1975). “Invariant Normal Models.” *The Annals of Statistics*, **3**, 132–154. doi: [10.1214/aos/1176343004](https://doi.org/10.1214/aos/1176343004).
- Andersson S, Madsen J (1998). “Symmetry and Lattice Conditional Independence in a Multivariate Normal Distribution.” *The Annals of Statistics*, **26**(2), 525–572. doi: [10.1214/aos/1028144848](https://doi.org/10.1214/aos/1028144848).
- Bache SM, Wickham H, Henry L (2022). **magrittr**: *A Forward-Pipe Operator for R*. doi: [10.32614/CRAN.package.magrittr](https://doi.org/10.32614/CRAN.package.magrittr). R package version 2.0.3.
- Borchers HW (2022). **numbers**: *Number-Theoretic Functions*. doi: [10.32614/CRAN.package.numbers](https://doi.org/10.32614/CRAN.package.numbers). R package version 0.8-5.
- Chojecki AP, Morgen P, Kołodziejek B (2025). **gips**: *Gaussian Model Invariant by Permutation Symmetry*. doi: [10.32614/CRAN.package.gips](https://doi.org/10.32614/CRAN.package.gips). R package version 1.2.3.
- Davis S, Meltzer P (2007). “**GEOquery**: A Bridge between the Gene Expression Omnibus (GEO) and **Bioconductor**.” *Bioinformatics*, **14**, 1846–1847. doi: [10.1093/bioinformatics/btm254](https://doi.org/10.1093/bioinformatics/btm254).
- Diaconis P, Ylvisaker D (1979). “Conjugate Priors for Exponential Families.” *The Annals of Statistics*, **7**(2), 269–281. doi: [10.1214/aos/1176344611](https://doi.org/10.1214/aos/1176344611).
- Gągolewski M (2022). “**stringi**: Fast and Portable Character String Processing in R.” *Journal of Statistical Software*, **103**(2), 1–59. doi: [10.18637/jss.v103.i02](https://doi.org/10.18637/jss.v103.i02).

- Gallian JA (2017). *Contemporary Abstract Algebra*. 9th edition. Cengage Learning.
- Gao X, Massam H (2015). “Estimation of Symmetry-Constrained Gaussian Graphical Models: Application to Clustered Dense Networks.” *Journal of Computational and Graphical Statistics*, **24**(4), 909–929. doi:10.1080/10618600.2014.937811.
- Gehrmann H (2011). “Lattices of Graphical Gaussian Models with Symmetries.” *Symmetry*, **3**(3), 653–679. doi:10.3390/sym3030653.
- Gentleman RC, Carey VJ, Bates DM, Bolstad B, Dettling M, Dudoit S, Ellis B, Gautier L, Ge Y, Gentry J, Hornik K, Hothorn T, Huber W, Iacus S, Irizarry R, Leisch F, Li C, Maechler M, Rossini AJ, Sawitzki G, Smith C, Smyth G, Tierney L, Yang JYH, Zhang J (2004). “**Bioconductor**: Open Software Development for Computational Biology and Bioinformatics.” *Genome Biology*, **5**(10), R80. doi:10.1186/gb-2004-5-10-r80.
- Graczyk P, Ishi H, Kołodziejek B (2022a). “Graphical Gaussian Models Associated to a Homogeneous Graph with Permutation Symmetries.” *Physical Sciences Forum*, **5**(1). ISSN 2673-9984. doi:10.3390/psf2022005020.
- Graczyk P, Ishi H, Kołodziejek B, Massam H (2022b). “Model Selection in the Space of Gaussian Models Invariant by Symmetry.” *The Annals of Statistics*, **50**(3), 1747–1774. doi:10.1214/22-aos2174.
- Graczyk P, Ishi H, Kołodziejek B, Massam H (2022c). “Supplement to “Model Selection in the Space of Gaussian Models Invariant by Symmetry”” *The Annals of Statistics*, **50**(3), 1–16. doi:10.1214/22-aos2174supp.
- Hankin RKS (2020). “Introducing the **permutations** Package.” *SoftwareX*. doi:10.1016/j.softx.2020.100453.
- Hastie T, Tibshirani R, Friedman J (2009). *The Elements of Statistical Learning*. Springer Series in Statistics, 2nd edition. Springer-Verlag. doi:10.1007/978-0-387-84858-7.
- Henry L, Wickham H (2023). **rlang**: *Functions for Base Types and Core R and tidyverse Features*. doi:10.32614/CRAN.package.rlang. R package version 1.1.1.
- Højsgaard S, Edwards D, Lauritzen S (2012). *Graphical Models with R*. Use R! Springer-Verlag. doi:10.1007/978-1-4614-2299-0.
- Højsgaard S, Lauritzen S (2008). “Graphical Gaussian Models with Edge and Vertex Symmetries.” *Journal of the Royal Statistical Society B*, **70**(5), 1005–1027. doi:10.1111/j.1467-9868.2008.00666.x.
- James G, Witten D, Hastie T, Tibshirani R (2021). *An Introduction to Statistical Learning: With Applications in R*. Springer Texts in Statistics, 2nd edition. Springer-Verlag. doi:10.1007/978-1-0716-1418-1.
- Jiang H, Fei X, Liu H, Roeder K, Lafferty J, Wasserman L, Li X, Zhao T (2021). **huge**: *High-Dimensional Undirected Graph Estimation*. doi:10.32614/CRAN.package.huge. R package version 1.3.5.

- Kotiang S, Eslami A (2020). “A Probabilistic Graphical Model for System-Wide Analysis of Gene Regulatory Networks.” *Bioinformatics*, **36**(10), 3192–3199. doi:10.1093/bioinformatics/btaa122.
- Lauritzen S (1996). *Graphical Models*, volume 17 of *Oxford Statistical Science Series*. The Clarendon Press, New York.
- Li Q, Gao X, Massam H (2020). “Bayesian Model Selection Approach for Coloured Graphical Gaussian Models.” *Journal of Statistical Computation and Simulation*, **90**(14), 2631–2654. doi:10.1080/00949655.2020.1784175.
- Li Q, Sun X, Wang N, Gao X (2021). “Penalized Composite Likelihood for Colored Graphical Gaussian Models.” *Statistical Analysis and Data Mining*, **14**(4), 366–378. doi:10.1002/sam.11530.
- Massam H, Li Q, Gao X (2018). “Bayesian Precision and Covariance Matrix Estimation for Graphical Gaussian Models with Edge and Vertex Symmetries.” *Biometrika*, **105**(2), 371–388. doi:10.1093/biomet/asx084.
- Miller LD, Smeds J, George J, Vega VB, Vergara L, Ploner A, Pawitan Y, Hall P, Klaar S, Liu ET, Bergh J (2005). “An Expression Signature for P53 Status in Human Breast Cancer Predicts Mutation Status, Transcriptional Effects, and Patient Survival.” *Proceedings of the National Academy of Sciences of the United States of America*, **102**(38), 13550–13555. doi:10.1073/pnas.0506230102.
- OEIS (2025). “A051625: Number of ‘Labeled’ Cyclic Subgroups of Symmetric Group S_n .” The On-Line Encyclopedia of Integer Sequences, URL <https://oeis.org/A051625>.
- Peeters CFW, Bilgrau AE, Van Wieringen WN (2022). “**rags2ridges**: A One-Stop- ℓ_2 -Shop for Graphical Modeling of High-Dimensional Precision Matrices.” *Journal of Statistical Software*, **102**(4), 1–32. doi:10.18637/jss.v102.i04.
- Peeters CFW, Bilgrau AE, Van Wieringen WN (2023). **rags2ridges**: *Ridge Estimation of Precision Matrices from High-Dimensional Data*. doi:10.32614/CRAN.package.rags2ridges. R package version 2.2.7.
- Ranciati S, Roverato A, Luati A (2021). “Fused Graphical Lasso for Brain Networks with Symmetries.” *Journal of the Royal Statistical Society C*, **70**(5), 1299–1322. doi:10.1111/rssc.12514.
- R Core Team (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Shah P, Chandrasekaran V (2012). “Group Symmetry and Covariance Regularization.” *Electronic Journal of Statistics*, **6**, 1600–1640. doi:10.1214/12-ejs723.
- Shannon P, Markiel A, Ozier O, Baliga NS, Wang JT, Ramage D, Amin N, Schwikowski B, Ideker T (2003). “**Cytoscape**: A Software Environment for Integrated Models of Biomolecular Interaction Networks.” *Genome Research*, **13**(11), 2498–2504. doi:10.1101/gr.1239303.

- Soloveychik I, Trushin D, Wiesel A (2016). “Group Symmetric Robust Covariance Estimation.” *IEEE Transactions on Signal Processing*, **64**(1), 244–257. ISSN 1053-587X. doi: [10.1109/tsp.2015.2486739](https://doi.org/10.1109/tsp.2015.2486739).
- Toh H, Horimoto K (2002). “Inference of a Genetic Network by a Combined Approach of Cluster Analysis and Graphical Gaussian Modeling.” *Bioinformatics*, **18**(2), 287–297. doi: [10.1093/bioinformatics/18.2.287](https://doi.org/10.1093/bioinformatics/18.2.287).
- Venables WN, Ripley BD (2002). *Modern Applied Statistics with S*. 4th edition. Springer-Verlag, New York.
- Wickham H (2016). **ggplot2**: *Elegant Graphics for Data Analysis*. Springer-Verlag, New York. doi: [10.1007/978-0-387-98141-3](https://doi.org/10.1007/978-0-387-98141-3).
- Wickham H, François R, Henry L, Müller K, Vaughan D (2023). **dplyr**: *A Grammar of Data Manipulation*. doi: [10.32614/CRAN.package.dplyr](https://doi.org/10.32614/CRAN.package.dplyr). R package version 1.1.2.

A. Example to Section 1.4

The standard PC can run all the code in this appendix within 2 seconds.

Consider an i.i.d. sample $(Z^{(i)})_{i=1}^n$ from $N_p(0, I_p)$ for $p = 4$ and $n = 50$ and let S be its empirical covariance matrix. The distribution of Z is clearly invariant under any permutation. Let us examine the output of `gips`.

```
R> p <- 4
R> n <- 50
R> set.seed(2022)
R> Z <- matrix(rnorm(n * p), ncol = p)
R> S <- cov(Z)
R> g <- gips(S, n)
R> g_MAP <- find_MAP(g, optimizer = "BF", return_probabilities = TRUE,
+   save_all_perms = TRUE)
R> get_probabilities_from_gips(g_MAP)

      (1,3,2,4)  (1,2,3,4)  (1,2,4,3)  (1,3,4)  (2,3,4)  (1,2,4)
2.542477e-01 2.393022e-01 2.124555e-01 1.837534e-01 6.234428e-02 2.517374e-02
      (1,2,3)  (1,4)(2,3)  (1,3)(2,4)  (1,2)(3,4)  (3,4)  (1,4)
1.971380e-02 9.287411e-04 6.006721e-04 4.634810e-04 3.743338e-04 2.542462e-04
      (1,3)  (2,4)  (2,3)  (1,2)  ()
1.811636e-04 8.977424e-05 8.207691e-05 3.465056e-05 2.263418e-07
```

We observe that the symmetries with the highest probability correspond to the long cycles, and these probabilities are very close to each other. This suggests that the data is invariant under each of these symmetries. The only model invariant under these three symmetries is the full-symmetry model, which is invariant under any permutation (both the diagonal and off-diagonal of the covariance matrix are constant).

B. Formulas for structure constants

In this appendix, we outline the steps required to find the ingredients necessary for the calculation of the normalizing constants

$$I_\Gamma(\delta, D) = \int_{\mathcal{P}_\Gamma} \text{Det}(k)^{(\delta-2)/2} e^{-\frac{1}{2}\text{Tr}[D \cdot k]} dk, \quad \delta > 1, D \in \text{Sym}^+(p; \mathbb{R}).$$

for arbitrary cyclic subgroup Γ . These constants are indispensable for our model selection procedure.

We note that the formulas for normalizing constants for an arbitrary subgroup $\Gamma \subset \mathfrak{S}_p$ are presented in [Graczyk *et al.* \(2022b\)](#). Here, we specialize these formulas to cyclic subgroups, which allows for significant simplification.

Let p_i be the length of the i -th cycle in the cyclic decomposition of $\sigma \in \mathfrak{S}_p$, and let $\{i_1, \dots, i_{C_\sigma}\}$ be a complete system of representatives of the cycles of σ . Furthermore, let $(e_i)_{i=1}^p$ be the standard basis of \mathbb{R}^p .

1. For $c = 1, \dots, C_\sigma$, calculate $v_1^{(c)}, \dots, v_{p_c}^{(c)} \in \mathbb{R}^p$ as

$$\begin{aligned} v_1^{(c)} &:= \sqrt{\frac{1}{p_c}} \sum_{k=0}^{p_c-1} e_{\sigma^k(i_c)}, \\ v_{2\beta}^{(c)} &:= \sqrt{\frac{2}{p_c}} \sum_{k=0}^{p_c-1} \cos\left(\frac{2\pi\beta k}{p_c}\right) e_{\sigma^k(i_c)} \quad (1 \leq \beta < p_c/2), \\ v_{2\beta+1}^{(c)} &:= \sqrt{\frac{2}{p_c}} \sum_{k=0}^{p_c-1} \sin\left(\frac{2\pi\beta k}{p_c}\right) e_{\sigma^k(i_c)} \quad (1 \leq \beta < p_c/2), \\ v_{p_c}^{(c)} &:= \sqrt{\frac{1}{p_c}} \sum_{k=0}^{p_c-1} \cos(\pi k) e_{\sigma^k(i_c)} \quad (\text{if } p_c \text{ is even}). \end{aligned}$$

2. Construct an orthogonal matrix U_Γ by arranging column vectors $\{v_k^{(c)}\}$, $1 \leq c \leq C_\sigma$, $1 \leq k \leq p_c$, in the following way: we put $v_k^{(c)}$ earlier than $v_{k'}^{(c')}$ if

- (i) $\frac{\lfloor k/2 \rfloor}{p_c} < \frac{\lfloor k'/2 \rfloor}{p_{c'}}$, or
- (ii) $\frac{\lfloor k/2 \rfloor}{p_c} = \frac{\lfloor k'/2 \rfloor}{p_{c'}}$ and $c < c'$, or
- (iii) $\frac{\lfloor k/2 \rfloor}{p_c} = \frac{\lfloor k'/2 \rfloor}{p_{c'}}$ and $c = c'$ and k is even and k' is odd.

3. Let N be the order of Γ . For $\alpha = 0, 1, \dots, \lfloor \frac{N}{2} \rfloor$ calculate

$$\begin{aligned} r_\alpha^* &= \#\{c \in \{1, \dots, C_\sigma\} : \alpha p_c \text{ is a multiple of } N\}, \\ d_\alpha^* &= \begin{cases} 1 & (\alpha = 0 \text{ or } N/2), \\ 2 & (\text{otherwise}). \end{cases} \end{aligned}$$

In the definition of r_α^* , we treat 0 as a multiple of N , and thus $r_0^* = C_\sigma$.

Then, we set $L = \#\{\alpha : r_\alpha^* > 0\}$, $r = (r_\alpha^* : r_\alpha^* > 0)$ and $d = (d_\alpha^* : r_\alpha^* > 0)$. The parameters $(r_i, d_i)_{i=1}^L$ are called the structure constants.

The constructed orthogonal matrix U_Γ possesses a notable property. According to [Graczyk et al. \(2022b, Theorem 5\)](#), it performs a block decomposition of the colored space \mathcal{Z}_Γ in the following sense: For each $S \in \mathcal{Z}_\Gamma$, we have

$$U_\Gamma^\top \cdot S \cdot U_\Gamma = \begin{pmatrix} x_1 & & \\ & \ddots & \\ & & x_L \end{pmatrix}, \quad (8)$$

where $x_i \in \text{Sym}(r_i d_i; \mathbb{R})$, $i = 1, \dots, L$.

For any $S \in \text{Sym}^+(p; \mathbb{R})$ and $\delta \in \mathbb{R}$, we define a function

$$\gamma_\Gamma(S, \delta) = \prod_{i=1}^L \text{Det}(x_i)^{-(\delta+r_i-3)/2-1/d_i},$$

where $x_i \in \text{Sym}(r_i d_i; \mathbb{R})$ are the diagonal blocks of a decomposition (8) of $\pi_\Gamma(S)$ (recall (3)).

Finally, by Graczyk *et al.* (2022b, Theorem 9), integral $I_\Gamma(\delta, D)$ is convergent if $(\delta - 2)/2 > \max_{i=1}^L \{-1/d_i\}$ and D is positive definite. The expression $\max_{i=1}^L \{-1/d_i\}$ equals $-1/2$ unless $L = 1$ (which corresponds to the trivial subgroup $\{\text{id}\}$), in which case it is equal to -1 . Thus, for all $\delta > 1$ and $D \in \mathcal{P}_\Gamma$ we have

$$I_\Gamma(\delta, D) = e^{-A_\Gamma(\delta-2)/2 - B_\Gamma} \gamma_\Gamma \left(\frac{1}{2} D, \delta \right) \prod_{i=1}^L \Gamma_i(1 + d_i(\delta + r_i - 3)/2),$$

where

$$A_\Gamma = \sum_{i=1}^L r_i d_i \log d_i, \quad B_\Gamma = \frac{1}{2} \sum_{i=1}^L r_i (1 + (r_i - 1)d_i/2) \log d_i,$$

$$\Gamma_i(\lambda) = (2\pi)^{r_i(r_i-1)d_i/4} \prod_{k=1}^{r_i} \Gamma(\lambda - (k-1)d_i/2).$$

Affiliation:

Adam Chojecki, Paweł Morgen, Bartosz Kołodziejek

Warsaw University of Technology

Faculty of Mathematics and Information Science

Koszykowa 75

00-662 Warsaw, Poland

E-mail: adam.prze.choj@gmail.com, seriousmorgen@protonmail.com,

bartosz.kolodziejek@pw.edu.pl