

A Bayesian Nonparametric Approach to Testing for Dependence Between Random Variables

Sarah Filippi* and Chris C. Holmes†

Abstract. Nonparametric and nonlinear measures of statistical dependence between pairs of random variables are important tools in modern data analysis. In particular the emergence of large data sets can now support the relaxation of linearity assumptions implicit in traditional association scores such as correlation. Here we describe a Bayesian nonparametric procedure that leads to a tractable, explicit and analytic quantification of the relative evidence for dependence vs independence. Our approach uses Pólya tree priors on the space of probability measures which can then be embedded within a decision theoretic test for dependence. Pólya tree priors can accommodate known uncertainty in the form of the underlying sampling distribution and provides an explicit posterior probability measure of both dependence and independence. Well known advantages of having an explicit probability measure include: easy comparison of evidence across different studies; encoding prior information; quantifying changes in dependence across different experimental conditions, and the integration of results within formal decision analysis.

Keywords: dependence measure, Bayesian nonparametrics, Pólya tree, hypothesis testing.

1 Introduction

Quantifying the evidence for dependence or testing for departures from independence between random variables is an increasingly important task and has been the focus of a number of studies in the past decade. A typical motivating example comes from the field of biology where a growing abundance of genetic, proteomic and transcriptomic data is being produced. In order to unravel the existing relationships between different molecular species (genes, proteins, ...) involved in a biological system, large datasets are commonly screened for evidence of association between the pairs of variables. This requires adequate statistical procedures to quantify the evidence of dependence (or lack of independence) between two samples of typically continuous random variables.

In this article, we propose a Bayesian nonparametric procedure to derive a probabilistic measure of dependency between two samples x and y without assuming a known form for the underlying distributions. In particular let \mathcal{M}_0 denote a model, or hypothesis, of independence and \mathcal{M}_1 a model, or hypothesis, of dependence. The posterior probability, $p(\mathcal{M}_1|x, y)$, is then a natural measure of the strength of evidence for dependence between the two samples against independence. The Bayes Factor quantifying the relative evidence in the data in favour of \mathcal{M}_1 over \mathcal{M}_0 is simply,

*Department of Statistics, University of Oxford, England, sarah.filippi@stats.ox.ac.uk

†Department of Statistics and Oxford-Man Institute, University of Oxford, England, cholmes@stats.ox.ac.uk

$$BF = \frac{p(x, y | \mathcal{M}_1)}{p(x | \mathcal{M}_0)p(y | \mathcal{M}_0)},$$

which is the ratio of the prior predictive probability of the observed data given the two competing hypotheses. This Bayes Factor implicitly avoids conditioning on the form of the unknown distribution functions. The role of Bayesian nonparametrics is to allow one to accommodate this uncertainty via a prior measure on the space of probability measures, for instance,

$$p(x, y | \mathcal{M}_1) = \int f(x, y | \mathcal{M}_1) \pi(dF | \mathcal{M}_1),$$

where $\pi(\cdot)$ is a Bayesian nonparametric prior with wide support over the space of probability measures on the joint sample space $\Omega_X \times \Omega_Y$; see for example Hjort et al. (2010).

We use Pólya tree priors (Lavine, 1992; Mauldin et al., 1992; Lavine, 1994) to model the unknown distributions of the data. We show that the use of such priors leads to an analytic derivation of the posterior probability of the dependence model $p(\mathcal{M}_1 | x, y)$. In particular, this measure of dependence involves a finite analytic calculation though the Pólya tree prior is defined over an infinite recursive partition. Pólya tree priors have previously been used to derive Bayesian nonparametric procedure for two samples hypothesis testing (Holmes et al., 2015; Ma and Wong, 2011) and extensions of these priors have been proposed to model distributions indexed by covariates (Trippa et al., 2011). The “two-sample testing” problem is different to that considered here in that it considers the same measurement, or outcome, Y , measured on independent samples under different conditions and tests for evidence of the “treatment” or covariate effect, whereas our paper is concerned with exploring evidence for statistical association between two joint measurement variables, $\{Y, X\}$, recorded together on a set of samples. However, our approach exploits a similar framework to the testing procedure from Holmes et al. (2015). In particular, our probabilistic measure necessitates the construction of Pólya tree priors on a two-dimensional space, and as discussed at the end of the paper, this engenders new challenges regarding the partitioning scheme to be employed. It is worth noting that Pólya trees offer a more appropriate nonparametric model than say Bayesian histograms, with Dirichlet priors (Leonard, 1973), as the recursive tree structure of the Pólya tree is indexed on the measurement variable, whereas the Dirichlet prior for histograms is for unordered categorical data and local dependence between measurement bins must be introduced via a hierarchical prior. Moreover the Pólya tree is defined via an infinite sequence partitioning, bypassing the need to truncate at some level, and, as noted above, our approach can compute the Bayes factors from the infinite sequence.

Numerous frequentist approaches have been developed for identifying associations between two samples (Shannon and Weaver, 1949; Cover and Thomas, 1991; Reshef et al., 2011; Gretton and Györfi, 2010) but to the best of our knowledge this is the first Bayesian nonparametric procedure to quantify the relative evidence of dependence vs independence. Being able to provide the posterior probability of the dependence model is attractive for a number of reasons. First, it can be combined with a variety of probabilistic approaches. In particular, it can be easily merged with a decision theory framework in order for optimal statistical decisions to be made in the face of uncertainty. Another important property of probabilistic measures is their interpretability.

Indeed, a given level $p = p(\mathcal{M}_1|x, y)$ of this measure is exactly the probability of a dependent generative model given the data and the probability of an independent generative model is simply $1 - p$. Over and above the standard arguments in favour of Bayesian inference, one explicit consequence of the coherence is that we can explicitly quantify the evidence for a change in dependence between two variables across two or more conditions. For example, if there is evidence that two dependent variables $\{X, Y\}$ become independent on application of a treatment, or across disease states. Answering such questions is problematic from a non-Bayesian perspective, as a null hypothesis of dependence is of higher dimension than the corresponding alternative hypothesis of independence which is nested *under* the null. This makes the quantification of a p-value extremely challenging. In Bayesian analysis, the symmetry of the model space makes for a simple and intuitive solution. In Section 4 we illustrate this issue using an application in cancer genetics concerned with the alteration of gene co-regulations with disease.

The remainder of the paper is organised as follows. We first introduce the Pólya tree priors in Section 2 and summarise their main properties. In Section 3, we describe our nonparametric procedure to test for dependencies between two samples. We then illustrate in Section 4 our approach on data generated from simple models and then we apply our procedure to two real world problems from biology. In the Supplementary material (Filippi and Holmes, 2016) we provide an empirical calibration comparing our method to that of other non-Bayesian approaches in the literature.

2 Polya tree priors

Pólya trees form a class of distributions for random probability measures F on some domain Ω (Ferguson, 1974) by considering a recursive partition of Ω into disjoint measurable spaces and constructing random measures on each of these spaces. A binary recursive partitioning is typically used for one-dimensional domains: Ω is divided in two disjoint sets C_0 and C_1 which themselves are divided in two other disjoint sets $C_0 = C_{00} \cup C_{01}$ and $C_1 = C_{10} \cup C_{11}$, and so on. The infinite recursive partition is denoted by $\mathcal{C} = \{C_j, j = 0, 1, 00, 01, 10, 11, \dots\}$; the partition at level k is comprised of 2^k sets C_j where j are all binary sequences of length k .

To better understand the probability measure constructed on these nested partitions, one can think of a particle going down the tree shown in Figure 1 (Left); at each junction j , usually represented in binary format, the particle has a random probability θ_j to choose the left branch. In Pólya trees, the random branching probability follows a Beta distribution, with $\theta_j \sim \text{Beta}(\alpha_{j,(0)}, \alpha_{j,(1)})$. Given the partition \mathcal{C} , the sequence of non-negative vectors $\mathcal{A} = \{\alpha_{j,(0)}, \alpha_{j,(1)}\}_j$ and the sequence of realisations of the random branching variables $\Theta = \{\theta_j\}_j$, it is possible to compute the likelihood for any set of observations x :

$$p(x|\Theta, \mathcal{C}, \mathcal{A}) = \prod_j \theta_j^{n_{j0}} (1 - \theta_j)^{n_{j1}}, \quad (1)$$

where the product is over the set of all partitions and n_{j0} and n_{j1} denote the number of observations that lie in the partitions C_{j0} and C_{j1} respectively. The Beta prior on the partition probability is conjugate to the Binomial likelihood and, integrating out θ_j for all j , we obtain that

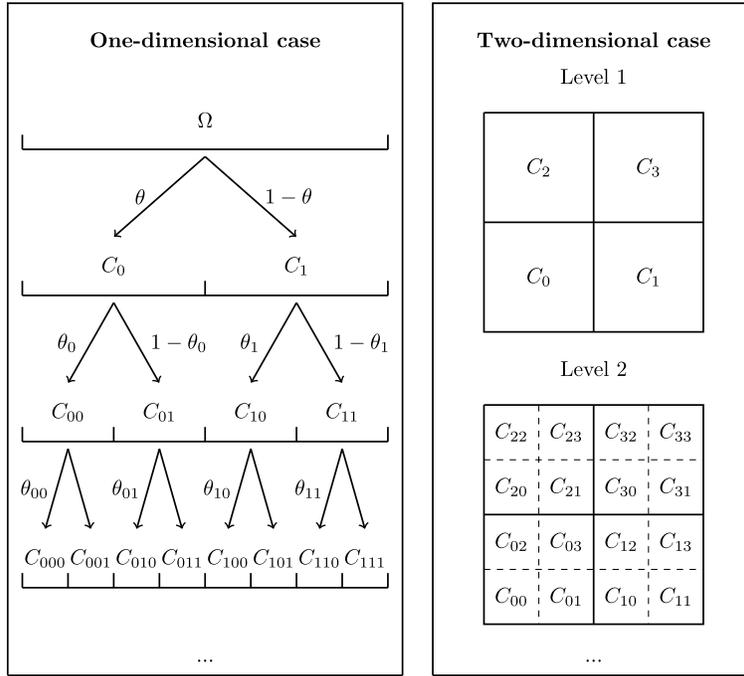


Figure 1: (Left) Construction of a Pólya tree distribution in the uni-dimensional case: at each junction j the particle has a random probability θ_j to choose the left branch and $1 - \theta_j$ to choose the right one. (Right) Illustration of the first two levels of the quadrant partitioning scheme in the two-dimensional case.

$$p(x|\mathcal{C}, \mathcal{A}) = \prod_j \frac{B(n_{j0} + \alpha_{j,(0)}, n_{j1} + \alpha_{j,(1)})}{B(\alpha_{j,(0)}, \alpha_{j,(1)})}, \tag{2}$$

where $B(\cdot, \cdot)$ refers to the Beta function. For more details on Pólya Tree priors, we refer the reader to Ferguson (1974); Lavine (1992); Mauldin et al. (1992); Lavine (1994); Ghosh and Ramamoorthi (2003); Wong et al. (2010).

In this paper, we are interested in testing independence between two samples x and y . We therefore need to consider the joint space $\Omega_X \times \Omega_Y$ of the two sample spaces. For reasons that will become obvious later on, we recursively subdivide this space into four rectangular regions. We start with partitioning $\Omega_X \times \Omega_Y$ in 4 quadrants, $\Omega_X \times \Omega_Y = C_0 \cup C_1 \cup C_2 \cup C_3$, and continue with nested partitions defined by $C_j = C_{j0} \cup C_{j1} \cup C_{j2} \cup C_{j3}$ for any base 4 number j . Thus the partition at level k is formed of 4^k sets C_j where j are all quaternary sequences of length k . We assume that the sets C_j are rectangular, i.e. can be written as a Cartesian product $D \times E$ where $D \subset \Omega_X$ and $E \subset \Omega_Y$. We arbitrarily choose to denote C_{j0} the left bottom region of the set C_j , C_{j1} the right bottom region, C_{j2} the left top region, and C_{j3} the right top region for all j . This recursive partition $\mathcal{C} = \{C_j, j = 0, 1, 2, 3, 01, 02, 03, 11, \dots\}$ is illustrated in Figure 1 (Right). Similarly to

the one-dimensional case, a probability measure can be constructed on this recursive partition by defining random branching probabilities in the recursive quaternary partition. In the following we will use different distributions for these branching probabilities.

3 A Bayesian nonparametric measure of dependence

3.1 The approach

Given a N sample (x, y) which are i.i.d. realisations of the random vector (X, Y) , we wish to evaluate the evidence for the competing hypotheses:

- \mathcal{M}_0 : X and Y are independent random variables;
- \mathcal{M}_1 : X and Y are dependent random variables.

We denote by F_{XY} the unknown joint probability distribution of (X, Y) and by F_X and F_Y the two unknown marginal distributions. Following a Bayesian approach, we aim at estimating the posterior probability

$$p(\mathcal{M}_1|x, y) \propto p(x, y|\mathcal{M}_1)p(\mathcal{M}_1),$$

where $p(\mathcal{M}_1)$ represents prior beliefs regarding the competing hypotheses. We specify our uncertainty in the distribution F_{XY} via a Pólya tree prior. Denoting by Ω_X and Ω_Y the domains of the probability measures F_X and F_Y respectively, we consider a recursive quaternary partition of $\Omega_X \times \Omega_Y$ into disjoint measurable sets as described in the previous section.

Under model \mathcal{M}_0 , we assume that samples x and y are independent. We can therefore think of the partitioning in terms of x-axis and y-axis separately. Let $\xi_{j,X}$ and $\xi_{j,Y}$ denote the independent random branching probabilities which determine the probability of going in the “left” region of C_j (i.e. $C_{j0} \cup C_{j2}$) and the “bottom” region of C_j (i.e. $C_{j0} \cup C_{j1}$) respectively. Similarly to the one-dimensional case, we assume that the random branching probabilities follow Beta distributions, $\xi_{j,X} \sim \text{Beta}(\alpha_{j,X,(0)}, \alpha_{j,X,(1)})$ and $\xi_{j,Y} \sim \text{Beta}(\alpha_{j,Y,(0)}, \alpha_{j,Y,(1)})$. By independence of $\xi_{j,X}$ and $\xi_{j,Y}$, the likelihood of the data given the partition \mathcal{C} , the sequence of random branching variables $\Xi = \{\xi_{j,X}, \xi_{j,Y}\}_j$, $\mathcal{A}_X = \{\alpha_{j,X,(0)}, \alpha_{j,X,(1)}\}_j$ and $\mathcal{A}_Y = \{\alpha_{j,Y,(0)}, \alpha_{j,Y,(1)}\}_j$ can be computed as follows

$$p(x, y|\Xi, \mathcal{C}, \mathcal{A}_X, \mathcal{A}_Y, \mathcal{M}_0) = \prod_j \xi_{j,X}^{n_{j0}+n_{j2}}(1 - \xi_{j,X})^{n_{j1}+n_{j3}} \xi_{j,Y}^{n_{j0}+n_{j1}}(1 - \xi_{j,Y})^{n_{j2}+n_{j3}},$$

where, for quaternary sequence $j \in \{0, 1, 2, 3, 01, \dots, 03, \dots, 31, \dots, 33, \dots\}$, n_j is the number of observations falling in C_j . Integrating the random branching probabilities out, we have

$$p(x, y|\mathcal{C}, \mathcal{A}_X, \mathcal{A}_Y, \mathcal{M}_0) = \prod_j \frac{B(n_{j0} + n_{j2} + \alpha_{j,X,(0)}, n_{j1} + n_{j3} + \alpha_{j,X,(1)})}{B(\alpha_{j,X,(0)}, \alpha_{j,X,(1)})} \times \frac{B(n_{j0} + n_{j1} + \alpha_{j,Y,(0)}, n_{j2} + n_{j3} + \alpha_{j,Y,(1)})}{B(\alpha_{j,Y,(0)}, \alpha_{j,Y,(1)})}. \quad (3)$$

Under the \mathcal{M}_1 hypothesis, we do not assume independence between samples x and y . In this case, for each set j , the random branching probability $\theta_j = (\theta_{j,(0)}, \theta_{j,(1)}, \theta_{j,(2)}, \theta_{j,(3)})$ is a random vector taking values in the simplex \mathbb{S}^3 . For every $i \in \{0, 1, 2, 3\}$, θ_{ji} is the probability that the particle falls in the quadrant C_{ji} . We use Dirichlet distributions with parameters $\alpha_j = (\alpha_{j,(0)}, \alpha_{j,(1)}, \alpha_{j,(2)}, \alpha_{j,(3)})$ for the random branching variables $\{\theta_j\}_j$ (Hanson, 2006). Hence the marginal likelihood is

$$p(x, y | \mathcal{C}, \mathcal{A}, \mathcal{M}_1) = \prod_j \frac{\tilde{B}(\tilde{n}_j + \alpha_j)}{\tilde{B}(\alpha_j)}, \quad (4)$$

where $\tilde{n}_j = (n_{j0}, n_{j1}, n_{j2}, n_{j3})$ and \tilde{B} is the multinomial Beta function defined as

$$\tilde{B}(\alpha_j) = \frac{\prod_{i=0}^3 \Gamma(\alpha_{j,(i)})}{\Gamma(\sum_{i=0}^3 \alpha_{j,(i)})},$$

where Γ designates the Gamma function.

Typically the values for the $\alpha_{j,(i)}$ are of the form ck^2 for α parameters at level k (Walker and Mallick, 1999); we recall that k is the depth of the set C_j and the length of the quaternary sequence j . We will follow this convention so that $a_j = \alpha_{j,(i)}$ for $i \in \{0, 1, 2, 3\}$. In addition, to ensure that the prior distributions are equivalent under the two models, we will assume that, for all j , $\alpha_{j,X,(0)} = \alpha_{j,(0)} + \alpha_{j,(2)} = 2a_j$, $\alpha_{j,X,(1)} = \alpha_{j,(1)} + \alpha_{j,(3)} = 2a_j$, $\alpha_{j,Y,(0)} = \alpha_{j,(0)} + \alpha_{j,(1)} = 2a_j$ and $\alpha_{j,Y,(1)} = \alpha_{j,(2)} + \alpha_{j,(3)} = 2a_j$.

To compare evidence in favours of both hypotheses, we compute the following ratio

$$\frac{p(\mathcal{M}_0 | x, y)}{p(\mathcal{M}_1 | x, y)} = \frac{p(x, y | \mathcal{M}_0) p(\mathcal{M}_0)}{p(x, y | \mathcal{M}_1) p(\mathcal{M}_1)},$$

where the first term is the Bayes factor which can be written as a product over all partitions:

$$\frac{p(x, y | \mathcal{M}_0)}{p(x, y | \mathcal{M}_1)} = \prod_j b_j, \quad (5)$$

where b_j is defined below. From equations (3) and (4) and expressing Beta and multinomial Beta functions in terms of Gamma functions, we have

$$b_j = \frac{\Gamma(n_{j0} + n_{j2} + 2a_j) \Gamma(n_{j1} + n_{j3} + 2a_j) \Gamma(n_{j0} + n_{j1} + 2a_j) \Gamma(n_{j2} + n_{j3} + 2a_j)}{\Gamma(n_{j0} + n_{j1} + n_{j2} + n_{j3} + 4a_j) \Gamma(n_{j0} + a_j) \Gamma(n_{j1} + a_j) \Gamma(n_{j2} + a_j) \Gamma(n_{j3} + a_j)} \\ \times \frac{\Gamma(4a_j) \Gamma(a_j)^4}{\Gamma(2a_j)^4}. \quad (6)$$

The product in (5) is defined over the infinite set of partitions. However for any set C_j containing zero or one data point (i.e. such that $n_j = n_{j0} + n_{j1} + n_{j2} + n_{j3} \leq 1$), $b_j = 1$. Therefore, only subsets with at least two data points contribute to this product. The Bayesian measure of the strength of evidence for dependence between the two samples against independence involves a finite analytic calculation even though (5) is over the infinite number of levels in the tree.

The procedure is described in Algorithm 1. For each set C_j containing more than one datapoint, the term b_j measures the relative evidence in favour of \mathcal{M}_0 given the number of datapoints falling in each of the four quadrants of C_j . Intuitively, for each set C_j we perform a Bayesian independence test based on the local two-by-two contingency table. Pólya tree priors provides us with a theoretical framework to perform these “local” independence tests at every level while taking into account potential dependences on neighbouring sets. In addition, it allows us to compute the Bayes Factor analytically without having to choose any arbitrary level or any truncation. The parameter a_j which decreases with the depth of the set C_j enables us to give more importance to “local” independence tests at the higher levels than at the deepest levels. In the next subsection, we investigate the impact of the choice of this parameter.

Algorithm 1 Bayesian nonparametric evidence for independence.

1. Fix the quadrant partitioning scheme; choose a constant c .
For every set C_j containing more than one data point, compute b_j defined in (6) with $a_j = ck^2$ where k is the depth of the set C_j .
2. Assuming equal prior belief for both hypotheses,

$$p(\mathcal{M}_1|x, y) = \frac{1}{1 + \prod_j b_j} = 1 - p(\mathcal{M}_0|x, y).$$

3.2 Sensitivity to choice of \mathcal{A}

The proposed procedure relies on a choice of the sequence of non-negative vector \mathcal{A} , \mathcal{A}_X and \mathcal{A}_Y . As discussed above, the α parameters are constant per level and such that $\alpha_{j,(i)} = ck^2$ where k is the depth of the set C_j . In addition, $\alpha_{j,X,(i)} = \alpha_{j,Y,(i)} = 2ck^2$. The parameter c controls the speed of divergence of the α parameters with the depth k and therefore the relative contribution of each level of the partition to the Bayes Factor. We have investigated the impact of the setting of c (see Figures S4 and S5 in Supplementary Material) and observe that small values of c typically favour the simpler model (\mathcal{M}_0) especially when the number of samples is small and there is not enough evidence to determine \mathcal{M}_1 . This is to be expected as Bayesian modelling encompasses a natural Occam factor in the prior predictive (see for example chapter 28 of MacKay (2003)). We have found $c = 5$ to be a reasonable canonical choice but practitioners are strongly advised to explore the setting for their own analysis.

3.3 Choice of the partition

Basic approach: partition centred on the median of the data

The inference resulting from a Pólya tree model is known to strongly depend on the specification of the partition \mathcal{C} over the data space (Paddock et al., 2003), and a multitude of quadrant partitioning scheme could be used in our procedure. As a default, partially subjective approach we suggest to construct a partition based on the quantiles

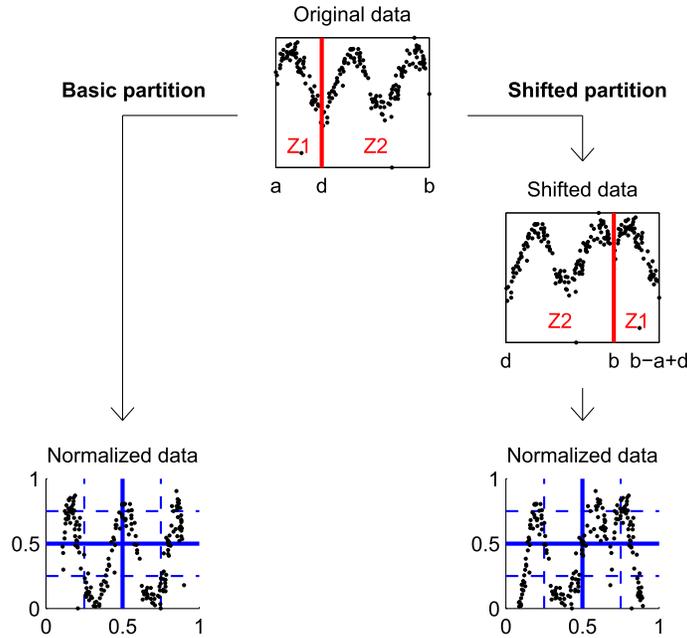


Figure 2: Construction of the partition. Both the basic approach and the empirical Bayes approach are illustrated on a simulated sinusoidal dataset with some outliers. Under the basic approach (Left column), the data are marginally transformed via the inverse of the c.d.f. of normal distributions. The empirical Bayes approach consists in shifting the central location of the partition by a factor δ and wrapping the data around: the data space is divided into two regions $\mathcal{Z}_1 = \{(x, y), x \leq \delta\}$ and $\mathcal{Z}_2 = \{(x, y), x > \delta\}$ which are then juxtaposed. The obtained “shifted” data are then normalised via the inverse of the CDF of a normal distribution. Quaternary recursive partitions of $[0, 1] \times [0, 1]$ are constructed by subdividing the normalised data into four rectangular quadrants of equal size (Bottom panels).

of two normal distributions (for the x - and y -axis respectively). In other words, both variables x and y are transformed through the inverse cumulative distributive function of a normal distribution; a quaternary recursive partition of $[0, 1] \times [0, 1]$ is then constructed by subdividing it into four rectangular quadrants of equal size (see Figure 2 (Left)) The mean and standard deviation of the normal distributions can be derived from empirical estimates of the location and spread of the two samples. In the next section we use the median and the median absolute deviation as this choice induces robustness to potential outliers.

A simple partial optimising procedure for partition centring

Two random variables X and Y are dependent if and only if the distribution of Y conditional on $X \in \mathcal{D}_X$ is different to the distribution of Y conditional on $X \in \mathcal{D}_X^c$ for

some \mathcal{D}_X a compact set of Ω_X and \mathcal{D}_X^C is its complement. In the previous paragraph, we suggest centring the partition on the median (m_x, m_y) of the data. For such choice of partition, at the top-level of the Pólya tree our procedure tests whether the distribution of Y conditional on $X \in \mathcal{D}_X = (-\infty, m_x]$ is equal to the distribution of Y conditional on $X \in \mathcal{D}_X^C$. The procedure performs the test symmetrically on the x-axis and the y-axis. Instead of focusing on the median, it would be more informative to test whether the distribution of Y conditionally on $X \in \mathcal{D}_X$ is equal to the distribution of Y conditionally on $X \in \mathcal{D}_X^C$ for any compact set $\mathcal{D}_X \in \Omega_X$.

In this section, we consider a simple partial optimisation of the partition-centring location by considering different compact sets \mathcal{D}_X . The approach involves shifting the central location of the partition as defined by the top-level split, and wrapping the data round to maintain balance in the number of points in each region. Consider a real number $\delta \in [a, b]$ where a and b are respectively the minimum and maximum values of the data on the x-axis. We denote by ψ_δ the transformation that divides the data space in two regions $\mathcal{Z}_1 = \{(x, y), x \leq \delta\}$ and $\mathcal{Z}_2 = \{(x, y), x \geq \delta\}$ and juxtaposes them as illustrated in Figure 2 (Right). More formally, $\psi_\delta : [a, b] \times \Omega_Y \rightarrow [\delta, b - a + \delta] \times \Omega_Y$ such that

$$\psi_\delta(x, y) = \begin{cases} (b - a + x, y) & \text{if } x \leq \delta, \\ (x, y) & \text{otherwise.} \end{cases}$$

The obtained “shifted” data are then transformed through the inverse cumulative distributive function of normal distributions and a quaternary recursive partition of $[0, 1] \times [0, 1]$ is constructed as in the basic approach. We denote the obtained partition by \mathcal{C}_δ .

We consider optimising the marginal evidence of dependence $p(\mathcal{M}_1|x, y)$ by maximising the Bayes factor as defined in equation (5) over all the partitions \mathcal{C}_δ for $\delta \in [a, b]$. The obtained probability of dependence is therefore

$$p(\mathcal{M}_1|x, y) = \frac{1}{1 + B_\delta},$$

where B_δ designates the Bayes Factor given the partition \mathcal{C}_δ . This approach is called the “empirical Bayes approach” in the rest of the paper. Note that optimising the central location of the partition with respect to $p(\mathcal{M}_1|x, y)$ will naturally tend to inflate the evidence for \mathcal{M}_1 . However, when testing many pairs of random variables for evidence of dependence we are mainly concerned with the ranking of the pairs for further analysis, rather than explicit quantification of the evidence, and partial optimisation of the partition may well help to produce more stable and accurate rankings as illustrated in Section 4.1.

4 Applications

In this section, we illustrate the performance of our Bayesian nonparametric procedure for detecting dependence across different datasets. We first test the procedure on simple models proposed by Kinney and Atwal (2014) and then apply it on two real examples from biology.

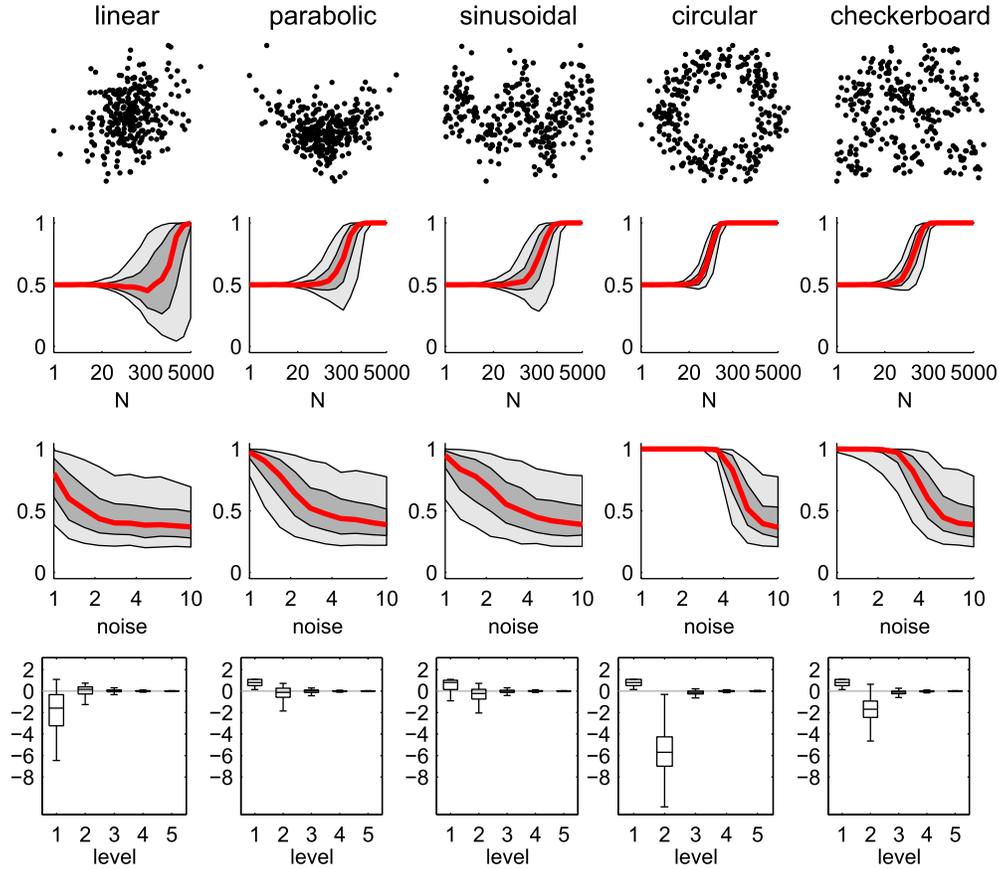


Figure 3: (Top row) Illustration of the five synthetic datasets and sampling distributions with $N = 300$ data and noise parameter $\sigma = 2$. (Middle rows 2 and 3) Frequency distribution (over 500 independent runs) of the probability of the hypothesis \mathcal{M}_1 for each model varying the number of data points (N) in plot row 2 and the level of noise (σ) in plot row 3. When varying the number of data points, the level of noise is fixed at $\sigma = 2$; when varying the level noise, the number of data points is fixed to $N = 300$. The red curve represents the median while the light and dark grey area designate the zone between the 5th and 95th percentiles and the inter-quartile region respectively. (Bottom row) Distribution (over 500 independent runs) of the contribution (B_k) of the 5 first levels in the Pólya Tree. A negative B_k indicates evidence against independence. We set $N = 150$ and $\sigma = 2$.

4.1 Illustration for synthetic datasets

We apply our Bayesian procedure on datasets generated under 5 different models proposed by Kinney and Atwal (2014) as illustrated in Figure 3 (Top row): a linear model ($y = 2x/3 + \eta$), a parabolic model ($y = 2x^2/3 + \eta$), a sinusoidal model ($y = 2 \sin(x) + \eta$),

a circular model ($x = 10 \cos(\theta) + \eta$ and $y = 10 \sin(\theta) + \eta$) and a model called “checkerboard” ($x = 10(i_x + \theta) + \eta$ and $y = 10(i_y + \theta) + \eta$ where $i_x \sim \mathcal{U}(\{0, 1, 2, 3\})$ and $i_y = \text{mod}(2u, i_x)$ with $u \sim \mathcal{U}(\{0, 1\})$); in addition $\theta \sim \mathcal{U}([0, 2\pi])$ and for each model we have i.i.d noise variables $\eta \sim \mathcal{N}(0, \sigma^2)$.

We generated 500 independent data sets of size N for each of the 5 models, and varied the level of noise (σ), and the number of data points (N). We used our procedure to compute the probability of the hypothesis \mathcal{M}_1 given each of these datasets. The partition structure was set using the robust mean and standard deviation of the data, and the parameter c was set equal to 5. The frequency distribution (over the 500 independent runs) of the probability of the hypothesis \mathcal{M}_1 for each model as a varying number of data points and the level of noise is shown in Figure 3 (Middle rows). The red curve represents the median while the light and dark grey area designate the zone between the 5th and 95th percentiles and the inter-quartile region respectively.

As expected, the probability that the two samples are dependent is equal to 0.5 for every model when $N = 1$ as we assumed equal prior belief of both hypotheses. The probability of \mathcal{M}_1 increases as the number of data points increases and is very close to one for every model if N is larger than 4000. It is interesting to note that when the number of samples is small there is not enough evidence to determine \mathcal{M}_1 and the Bayes Factor may favour the simpler model \mathcal{M}_0 . As mentioned previously, this it to be expected as Bayesian modelling encompasses a natural Occam factor in the prior predictive (see for example chapter 28 of MacKay (2003)). This effect is stronger for other smaller values of the parameter c (see Figures S4 and S5 in Supplementary Material).

The logarithm of the Bayes Factor defined in equation (5) can be decomposed in terms of levels in the recursive partition as follows

$$\log \left(\frac{p(x, y | \mathcal{M}_0)}{p(x, y | \mathcal{M}_1)} \right) = \sum_k \left(\sum_{j \text{ s.t. } C_j \text{ in level } k} \log(b_j) \right) = \sum_k B_k$$

and the contribution of each level in favour of the independence or dependence model (denoted by B_k for the level k) can be investigated. When B_k is close to 0, the contribution of level k is negligible, whereas large positive (resp. negative) B_k indicates stronger evidence in favour of independence (resp. dependence) at level k . In Figure 3 (Bottom row), we show the distribution (over 500 independent runs) of B_k for $k \in \{1, 2, 3, 4, 5\}$ for a sample of $N = 150$ data generated from the 5 different models with $\sigma = 2$. We observe that in the linear model, the dependence can already be detected at the first level, with B_1 being strongly negative for most of the generated datasets. However, for the four other models, the value of B_1 is mostly positive and the top level does not contain enough information to detect that there are dependencies between the two variables. In these examples, most of the information in favour of the dependence model is in the second level. Deeper levels contribute less to the decision; this is due to the form of parametrisation with the α parameters being proportional to k^2 . This decomposition of the evidence across levels is an attractive qualitative feature of the Pólya tree testing framework, which can assist the statistical analyst in better understanding the dependence structure.

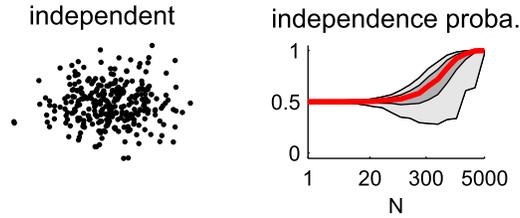


Figure 4: (Left) Illustration of the independent models sampling 300 data from two independent standard normal distributions. (Right) Distribution (over 500 independent runs) of the probability of the hypothesis \mathcal{M}_0 varying the number of data points (N). The red curve represents the median whereas the light and dark grey area designate the zone between the 5th and 95th percentiles and the inter-quartile region respectively.

The symmetric nature of the probability, $p(\mathcal{M}_0|\cdot) = 1 - p(\mathcal{M}_1|\cdot)$, allows us to explore the ability to detect independence. The probability of the independent hypothesis \mathcal{M}_0 given data sampled from two independent standard normal distributions is shown in Figure 4 for increasing number of data points (N). The probability of the independent hypothesis increases with N and is very close to one if N is larger than 500. Such a measure of independence between variables is problematic to compute for non-Bayesian methods, as we are testing for a simpler model nested within a more complex one. Frequentist approaches use a p-value which is conditional on \mathcal{M}_0 being true. Hence as it stands it cannot be used as evidence for \mathcal{M}_0 .

The performance of our Bayesian nonparametric approach varies between models both in terms of number of data points required to detect dependence and in terms of noise sensitivity: dependencies are detected for the circular and checkerboard models even for relatively high levels of noise and relatively small number of data points but less so for the linear model, which visually appears closer to independence (top plots in Figure 3). In addition, our approach necessitates a relatively high number of data points ($N \geq 300$) to detect dependence in the parabolic or the sinusoidal models even for a level of noise $\sigma = 2$. We can obtain some explanation for this behaviour by observing the bottom plots of Figure 3 where we see that for the parabolic and sinusoid examples there is a relative large evidence at level 1 in favour of independence. This tempers the contribution of the evidence at level 2 that detects the dependence structure present at this level of resolution. One way to try and address this issue is to change the partitioning scheme and apply the empirical Bayes approach described in Section 3.3. It consists in maximising the marginal evidence of the dependence model over different partition centrings. Figure 5 shows that the evidence for dependence is strongly increased for those two models when running the empirical Bayes approach. Centring the Pólya Tree partition on the model that maximises the evidence of \mathcal{M}_1 will clearly obscure the Bayesian interpretation of the posterior probability. However, when testing many pairs of variables in a common data set, the rank ordering of the pairs by the resulting Bayes Factor may still be a useful method for highlighting those showing greatest evidence. In the Supplementary Material, we provide an extensive, frequentist, power analysis of our approach against other popular non-Bayesian meth-

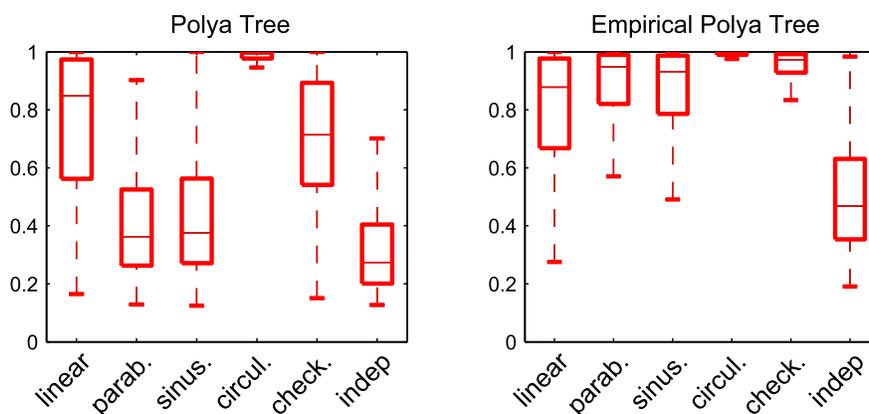


Figure 5: Distribution (over 500 independent runs) of the probability of dependence using our Bayesian nonparametric approach (Left) and our empirical Bayes approach which maximises the marginal probability of dependence over the shifted partition scheme with data wrapping (Right). Data are generated under the 5 illustrative examples as well as an independent generative model where both x and y are vectors of i.i.d. samples from a normal distribution with mean 0 and standard deviation 1. Here, $N = 150$ and $\sigma = 2$.

ods. We observe that the power of every method strongly varies from one generative model to another. Overall the power of the empirical Bayes approach appears to be competitive (see Figure S2 in Supplementary Material); however, on some datasets, our approach observes lower power than simpler non Bayesian state-of-the-art methods such as the Maximum Information Criterion (Reshef et al., 2011) or the Mutual Information estimated using the 20-nearest neighbours (Shannon and Weaver, 1949; Cover and Thomas, 1991). The sensitivity of results to the partitioning structure of the Pólya Tree is known to be a issue in Bayesian nonparametric model (Hjort et al., 2010), and is an open area of future research. Clearly careful diagnostic checking should be made by the statistician when employing any method, and permutation testing under the null could be employed here to highlight sensitivity issues with regards to power, although we do not explore this further.

4.2 Applications from molecular biology

Gene expression network form measurements at single-cell resolution

The field of biology contains numerous examples where a large amount of data has been produced and adequate measures to detect dependence between variables are required. Here we focus on an example from single-cell biology. Nowadays, the expression level of thousands of genes can be jointly measured at single-cell resolution, which allows biologists to precisely study the functional relationships between genes. In Wills et al. (2013), the expression of 96 genes affected by Wnt signalling have been measured in 288 single

cells. The authors provided evidence that many of these transcriptomic associations are masked when expression is averaged over bulk sequencing on many cells. In their study the authors investigated the relationships between these genes and constructed an expression network using the measurements at single-cell resolution. The expression network can illustrate potential functional relationships and dependencies that are interesting to elicit molecular pathways. The network is constructed by highlighting genes that have correlated or anti-correlated expressions between cells using Spearman correlation coefficient. We reproduce this network detecting dependences both with Spearman correlation and with our Bayesian procedure (see Figure 6 (Top row)). Our procedure detects many associations between genes that were not detected by simple correlation analysis: More than 250 pairs of genes have a probability of dependence higher than 0.95 and an absolute value of Spearman correlation lower than 0.5 whereas every (except one) link with absolute Spearman correlation higher than 0.5 has a probability of dependence larger than 0.95 (see Figure 6 (Bottom row, Left)). Some of these links that are only detected using our approach are between genes that are known to interact such as APC and DVL2, AXIN1 and GSK3B, DVL2 and LRP6 or AXIN1 and DACT1 (see Figure 6 (Bottom row)). Other detected links remain to be investigated.

Differential co-expression analysis

Networks have proved themselves to be important representation of biological systems where various molecules are interacting and functionally coordinating. A typical example is gene expression networks such as the one described in the previous subsection where nodes correspond to genes and edges represent interactions between genes. Interactions in biological networks can substantially change in response to different conditions. In particular, gene co-regulations may be altered with disease and an interaction between two genes could be present in some conditions and not in other. Differential co-expression analysis consist in identifying which interactions in gene expression network change from one condition to another (Hsu et al., 2015).

The main objectives of differential co-expression analysis is to identify couples of genes (x, y) such that the strength of dependence between x and y changes in response to different conditions. Our Bayesian procedure is perfectly suited for this type of problems which require methods able to detect both dependences and independences. Non-Bayesian testing procedures for independences typically only provide p-values to identify when the null hypothesis (here, the independence hypothesis) can be rejected. To the contrary our approach enables us to quantify the relative evidence of dependence versus independence. In particular, given the expression $\{x_i, y_i\}_{i=1, \dots, n}$ of two genes in n cells under condition A and the expression $\{\tilde{x}_i, \tilde{y}_i\}_{i=1, \dots, \tilde{n}}$ of the same two genes in \tilde{n} cells under another condition B, we can calculate the probability of a change of interactions between these two genes in response to conditions A and B as follows

$$p_{\text{diff}}(\{x_i, y_i\}_{i=1, \dots, n}, \{\tilde{x}_i, \tilde{y}_i\}_{i=1, \dots, \tilde{n}}) = \tag{7}$$

$$p(\mathcal{M}_1 | \{x_i, y_i\}_i) \left(1 - p(\widetilde{\mathcal{M}}_1 | \{\tilde{x}_i, \tilde{y}_i\}_i)\right) + p(\widetilde{\mathcal{M}}_1 | \{\tilde{x}_i, \tilde{y}_i\}_i) \left(1 - p(\mathcal{M}_1 | \{x_i, y_i\}_i)\right),$$

where \mathcal{M}_1 and $\widetilde{\mathcal{M}}_1$ denote the dependence model under respectively condition A and B.

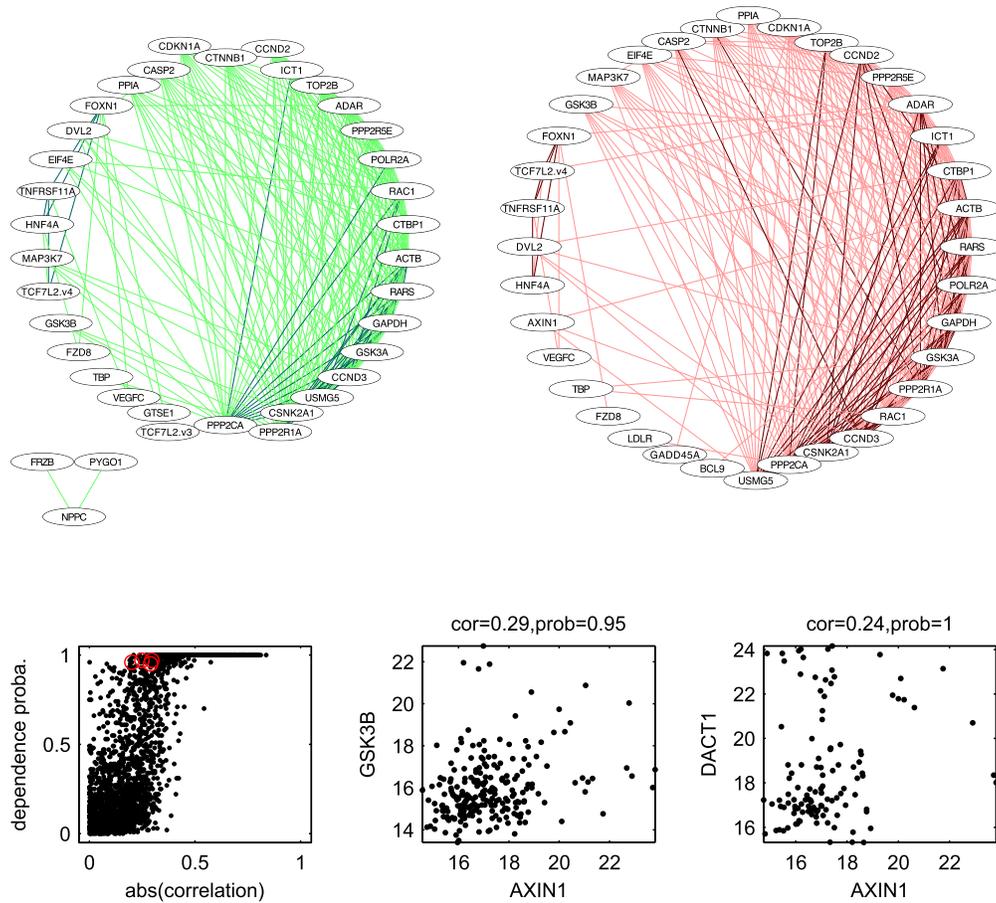


Figure 6: (Top row) Expression network constructed using correlation (Left) – where links with absolute correlation larger than 0.5 are shown and darker edges represent links with absolute correlation larger than 0.7 – and our nonparametric Bayesian procedure (Right) – where links with probability of dependences larger than 0.99999 are shown and darker edges represent links with probabilities equal to 1. Genes are ordered clockwise according to increasing number of detected links. (Bottom row – Left) Comparison of the probability of dependence computed following our approach and the absolute value of Spearman correlation for every pairs of genes. The 4 red circles indicate some pairs of genes with known interactions which have an absolute correlation lower than 0.5 but a probability of dependences larger than 0.95. (Bottom row – Middle and Right) Examples of gene expression data for two genes with known interaction.

In Curtis et al. (2012), a collection of around 2,000 breast cancer specimens from tumour banks in the UK and Canada is analysed and compared to a set of 144 normal cells. We propose to apply our algorithm to these gene expression dataset and make use

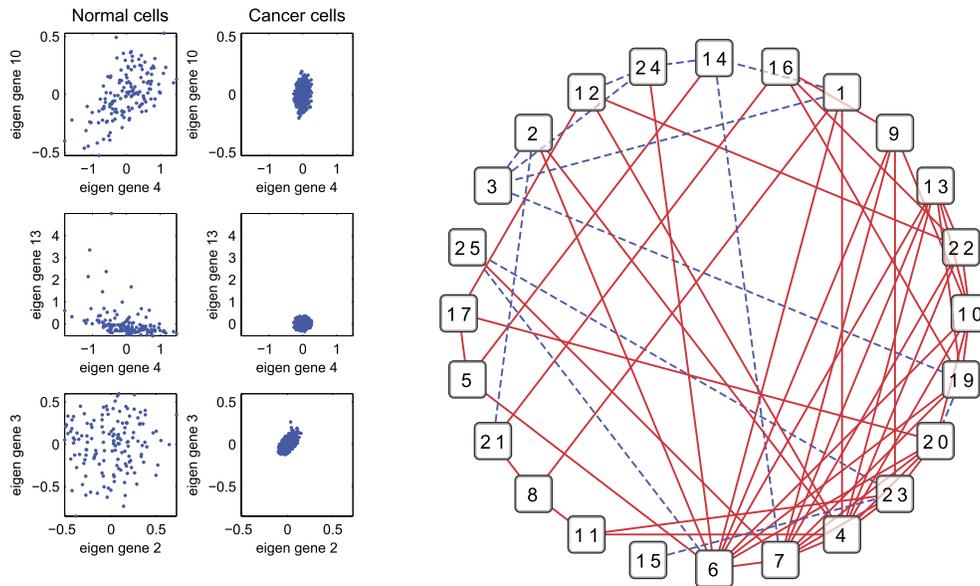


Figure 7: (Left) Expression of eigengenes for different modules (module 2, 3, 4, 10 and 13) in the 144 normal cells and in the 997 tumour cells. We observe that some modules are strongly interacting under one of the conditions (Normal or Breast Cancer) and not under the other. Each dot corresponds to the expression of two eigengenes in one cell. (Right) Differential co-expression network: nodes correspond to modules of genes; edges represent module interactions that significantly change between normal and cancer conditions. Red continuous edges correspond to interactions that are present in normal cells and vanished in tumour cells; blue dashed edges correspond to interactions that are only present in tumour cells.

of the probability in equation (7) in order to identify dysregulation in gene expression in response to breast cancer. We focus on comparing a subset of 997 tumour cells (called the discovery test in Curtis et al. (2012)) with the set of 144 normal cells; for each cell, the expression level of 48,803 probes is available. Following the approach proposed by Langfelder and Horvath (2007) we use the gene expression of the normal cells to identify 25 modules of correlated genes and determine so called “eigengenes” that represent the expression of genes in each module. We used the R implementation of this module detection and eigengene computation provided in the R package WGCNA (Langfelder and Horvath (2008)). We represent in Figure 7 (Left) the eigengene expression of some selected modules under the two conditions. By computing the probability in (7) for the 300 pairs of modules, we identify interactions between modules that significantly change in response to breast cancer. We find that the probability p_{diff} is larger than 0.95 for 69 module interactions, represented in Figure 7 (Right). Among those 69 interactions, 49 are interactions that were present in normal cells and vanished in tumour cells whereas 13 of the interactions only appear in tumour cells.

Module Number	Over-represented Pathway(s)
1	Cell Cycle, Mitotic
2	ATP sensitive Potassium channels
3	Gene Expression
4	Elastic fibre formation
5	IFN-alpha/beta pathways + Interferon Signalling*
6	ALK1 signalling events
7	Translation + Ribosome*
8	Bladder cancer – Homo sapiens (human)
9	Apoptotic cleavage of cell adhesion proteins
10	ERK/MAPK targets
11	Extracellular matrix organisation
12	AP-1 transcription factor network
13	Metabolism
14	Generation of second messenger molecules
15	The citric acid (TCA) cycle and respiratory electron transport
16	Thromboxane signalling + ADP signalling
17	Peptide ligand-binding receptors
18	Generic Transcription Pathway
19	Type I hemidesmosome assembly
20	Peptide chain elongation
21	PI3K-Akt signalling pathway – Homo sapiens (human)
22	Eukaryotic Translation Elongation
23	Immune System
24	Metabolism
25	Mitochondrial translation elongation

Table 1: Pathway enrichment results, using three pathway databases (PID, KEGG, and Reactome), for the 25 modules identified. For each module, the most significant pathway was selected based on adjusted p-values. Ties were resolved by taking the smaller pathway (for pathways with large discrepancy in size), or by the most significant by unadjusted p-values.

The pathway enrichment analysis enables us to implicate each of the 25 modules with established biological cascades and clinically-relevant pathways (see Table 1). The two modules with the highest degree in the differential co-expression network are: (a) the ALK1 signalling pathway, and (b) complexes associated with translation (e.g. ribosome). For the former, this likely indicates loss of regulation and signalling cross-talk; for the latter, since ribosomes are essential and translational genes are often tightly regulated, this hints at wide-spread transcriptomic perturbation. In particular, the high degree of the ribosome/translation module indicates that, in the cancerous state, more modules are increasingly dysregulated and out of sync with the more tightly regulated translation-involved modules. Our analysis shows that both of these high-degree modules are disconnected to numerous other important pathways such as the ERK/MAPK pathway, or genes involved in immune signalling (such as antigen presentation). More generally, this demonstrates that overlaying biological and clinically relevant annota-

tions on the differential co-expression network may be the basis for further research regarding transcriptomic alterations in breast tumours.

5 Discussion/conclusion

We have presented a novel Bayesian nonparametric approach that quantifies a probabilistic measure for the strength of evidence for dependence between two samples against that of independence. The procedure is based on Pólya tree priors that facilitate an analytic expression for the Bayes factor even though the Pólya tree prior is defined over an infinite recursive partition. We have applied our approach to simulated datasets as well as applications in molecular biology including single-cell gene-expression analysis and network analysis in cancer genetics.

The inference resulting from a Pólya tree model is known to strongly depend on the specification of the partition \mathcal{C} over the data space (Hjort et al., 2010). We have illustrated the sensitivity of the performance of our approach to the partitioning structure of the Pólya Tree by introducing an empirical procedure which consists in maximising the marginal evidence of the dependence model over different partition centrings. By definition, this empirical approach tends to inflate the evidence in favour of dependence which complicates the Bayesian interpretation of the posterior probability. However, the Bayes Factor obtained following the empirical Bayes approach can be used to accurately rank order different pairs of variables in a common data set for highlighting those showing greatest evidence in favour of the dependence model. Future work would involve exploring further methods to optimise the partitioning structure of the Pólya Tree such as incorporating random partitioning schemes.

Our probabilistic measure has some importance advantages over existing frequentist statistics such as Mutual Information or the Maximal Information Criterion (Shannon and Weaver, 1949; Cover and Thomas, 1991; Reshef et al., 2011) due to its interpretability in terms of a recursive partition of the data space, and its symmetry in that $p(\mathcal{M}_0|\cdot) = 1 - p(\mathcal{M}_1|\cdot)$. The explicit quantification of a probability allows for combining with other sources of information within a prior or meta-analysis. As shown in Section 4.2 the Bayesian approach provides a unified method for detecting both independence and dependence, something that is not possible without a fully probabilistic framework. The Bayesian probabilistic approach allows for the inclusion of substantive prior information on the plausibility of an association, which can be particularly useful for screening large biological data sets. There is also the possibility to embed the model within a hierarchical structure, borrowing strength coherently across categories, something that is challenging for existing approaches based on non-probabilistic methods.

Supplementary Material

Supplement Material of “A Bayesian Nonparametric Approach to Testing for Dependence Between Random Variables” (DOI: [10.1214/16-BA1027SUPP](https://doi.org/10.1214/16-BA1027SUPP); .pdf).

References

- Cover, T. M. and Thomas, J. (1991). *Elements of Information Theory*. Wiley, New York. MR1122806. doi: <http://dx.doi.org/10.1002/0471200611>. 920, 931, 936
- Curtis, C., Shah, S. P., Chin, S.-F., Turashvili, G., Rueda, O. M., Dunning, M. J., Speed, D., Lynch, A. G., Samarajiwa, S., Yuan, Y., et al. (2012). “The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups.” *Nature*, 486(7403): 346–352. 933, 934
- Ferguson, T. S. (1974). “Prior distributions on spaces of probability measures.” *The Annals of Statistics*, 615–629. MR0438568. 921, 922
- Filippi, S. and Holmes, C. C. (2016). “Supplementary Material of “A Bayesian Nonparametric Approach to Testing for Dependence Between Random Variables”.” *Bayesian Analysis*. doi: <http://dx.doi.org/10.1214/16-BA1027SUPP>. 921
- Ghosh, J. K. and Ramamoorthi, R. (2003). *Bayesian Nonparametrics*, volume 1. Springer. MR1992245. 922
- Gretton, A. and Györfi, L. (2010). “Consistent nonparametric tests of independence.” *The Journal of Machine Learning Research*, 11: 1391–1423. MR2645456. 920
- Hanson, T. E. (2006). “Inference for mixtures of finite Polya tree models.” *Journal of the American Statistical Association*, 101(476). MR2279479. doi: <http://dx.doi.org/10.1198/016214506000000384>. 924
- Hjort, N. L., Holmes, C., Müller, P., and Walker, S. G. (2010). *Bayesian Nonparametrics*, volume 28. Cambridge University Press. MR2730659. doi: <http://dx.doi.org/10.1017/CB09780511802478.001>. 920, 931, 936
- Holmes, C. C., Caron, F., Griffin, J. E., and Stephens, D. A. (2015). “Two-sample Bayesian nonparametric hypothesis testing.” *Bayesian Analysis*, 10(2): 297–320. MR3420884. doi: <http://dx.doi.org/10.1214/14-BA914>. 920
- Hsu, C.-L., Juan, H.-F., and Huang, H.-C. (2015). “Functional Analysis and Characterization of Differential Coexpression Networks.” *Scientific Reports*, 5. 932
- Kinney, J. B. and Atwal, G. S. (2014). “Equitability, mutual information, and the maximal information coefficient.” *Proceedings of the National Academy of Sciences*, 111(9): 3354–3359. MR3200177. doi: <http://dx.doi.org/10.1073/pnas.1309933111>. 927, 928
- Langfelder, P. and Horvath, S. (2007). “Eigengene networks for studying the relationships between co-expression modules.” *BMC Systems Biology*, 1(1): 54. 934
- Langfelder, P. and Horvath, S. (2008). “WGCNA: an R package for weighted correlation network analysis.” *BMC Bioinformatics*, 9(1): 559. 934
- Lavine, M. (1992). “Some aspects of Polya tree distributions for statistical modelling.” *The Annals of Statistics*, 1222–1235. MR1186248. doi: <http://dx.doi.org/10.1214/aos/1176348767>. 920, 922

- Lavine, M. (1994). “More aspects of Polya tree distributions for statistical modelling.” *The Annals of Statistics*, 1161–1176. MR1311970. doi: <http://dx.doi.org/10.1214/aos/1176325623>. 920, 922
- Leonard, T. (1973). “A Bayesian method for histograms.” *Biometrika*, 60(2): 297–308. MR0326902. 920
- Ma, L. and Wong, W. H. (2011). “Coupling optional Pólya trees and the two sample problem.” *Journal of the American Statistical Association*, 106(496). MR2896856. doi: <http://dx.doi.org/10.1198/jasa.2011.tm10003>. 920
- MacKay, D. J. (2003). *Information Theory, Inference, and Learning Algorithms*, volume 7. Citeseer. 925, 929
- Mauldin, R. D., Sudderth, W. D., and Williams, S. (1992). “Polya trees and random distributions.” *The Annals of Statistics*, 1203–1221. MR1186247. doi: <http://dx.doi.org/10.1214/aos/1176348766>. 920, 922
- Paddock, S. M., Ruggeri, F., Lavine, M., and West, M. (2003). “Randomized Polya tree models for nonparametric Bayesian inference.” *Statistica Sinica*, 13(2): 443–460. MR1977736. 925
- Reshef, D. N., Reshef, Y. A., Finucane, H. K., Grossman, S. R., McVean, G., Turnbaugh, P. J., Lander, E. S., Mitzenmacher, M., and Sabeti, P. C. (2011). “Detecting novel associations in large data sets.” *Science*, 334(6062): 1518–1524. 920, 931, 936
- Shannon, C. E. and Weaver, W. (1949). *The Mathematical Theory of Communication*. The University of Illinois Press. MR0032134. 920, 931, 936
- Trippa, L., Müller, P., and Johnson, W. (2011). “The multivariate beta process and an extension of the Polya tree model.” *Biometrika*, 98(1): 17–34. MR2804207. doi: <http://dx.doi.org/10.1093/biomet/asq072>. 920
- Walker, S. and Mallick, B. K. (1999). “A Bayesian semiparametric accelerated failure time model.” *Biometrics*, 55(2): 477–483. MR1705102. doi: <http://dx.doi.org/10.1111/j.0006-341X.1999.00477.x>. 924
- Wills, Q. F., Livak, K. J., Tipping, A. J., Enver, T., Goldson, A. J., Sexton, D. W., and Holmes, C. (2013). “Single-cell gene expression analysis reveals genetic associations masked in whole-tissue experiments.” *Nature Biotechnology*, 31(8): 748–752. 931
- Wong, W. H., Ma, L., et al. (2010). “Optional Pólya tree and Bayesian inference.” *The Annals of Statistics*, 38(3): 1433–1459. MR2662348. doi: <http://dx.doi.org/10.1214/09-AOS755>. 922

Acknowledgments

We thank Moustafa Abdalla for sharing his insights into gene expression analysis and pathway enrichment, and in his support for the analysis in Section 4.2. Holmes is supported by the Oxford-Man Institute, the Alan Turing Institute, the Engineering and Physical Sciences Research Council programme grant i-like EP/K014463/1, and the Medical Research Council, UK, grant MC_UP_A390_1107.