Published in final edited form as:

Biom J. 2015 May ; 57(3): 468–484. doi:10.1002/bimj.201300130.

Bayesian semiparametric copula estimation with application to psychiatric genetics

Ori Rosen^{*,1} and Wesley K. Thompson²

¹Department of Mathematical Sciences, University of Texas at El Paso, El Paso, Texas 79968, U.S.A

²Department of Psychiatry, University of California at San Diego, La Jolla, California 92093, U.S.A

Abstract

This paper proposes a semiparametric methodology for modeling multivariate and conditional distributions. We first build a multivariate distribution whose dependence structure is induced by a Gaussian copula and whose marginal distributions are estimated nonparametrically via mixtures of B-spline densities. The conditional distribution of a given variable is obtained in closed form from this multivariate distribution. We take a Bayesian approach, using Markov chain Monte Carlo methods for inference. We study the frequentist properties of the proposed methodology via simulation and apply the method to estimation of conditional densities of summary statistics, used for computing conditional local false discovery rates, from genetic association studies of schizophrenia and cardiovascular disease risk factors.

Keywords

B-spline densities; Cardiovascular Disease Risk Factors; Gaussian copula; Schizophrenia

1 Introduction

This paper proposes a semiparametric methodology for modeling multivariate and conditional densities. We first build a multivariate distribution whose dependence structure is induced by a Gaussian copula and whose marginal distributions are estimated nonparametrically via mixtures of B-spline densities. The conditional distribution of a given variable as a function of the remaining variables is then obtained from this multivariate distribution. We take a Bayesian approach, using Markov chain Monte Carlo (MCMC) methods for inference.

Our research is motivated by an application in psychiatric genetics. Individuals with schizophrenia have significantly higher mortality rates compared with the general population, corresponding to a 10–20 year reduction in life expectancy (Colton and Manderscheid, 2006; Laursen et al., 2012; Saha et al., 2007). Although the mortality rate

*Corresponding author: ori@math.utep.edu, Phone: +1-915-747-6843, Fax: +1-915-747-6502.

Conflict of Interest

The authors have declared no conflict of interest.

from suicide is high, lifestyle and cardiovascular disease (CVD) risk factors contribute substantially to life expectancy reduction in schizophrenia (Marder et al., 2004; Mitchell et al., 2011). Epidemiological research has shown increased rates of dyslipidemia, type 2 diabetes, obesity, and a high prevalence of metabolic syndrome among people with schizophrenia (Mitchell et al., 2011; De Hert et al., 2006). This increase in CVD risk factors has been primarily attributed to lifestyle factors such as unhealthy diets, sedentary habits, excessive smoking, and to antipsychotic medication side-effects (Laursen et al., 2012; De Hert et al., 2006; Kaddurah-Daouk et al., 2007). However, as suggested by studies (i) predating introduction of antipsychotics (Raphael and Parsons, 1921), (ii) on untreated first episode individuals as well as their healthy relatives (Ryan et al., 2003) and (iii) on overlapping candidate genes (Hansen et al., 2011), shared genetics between schizophrenia and CVD risk factors may also be of importance.

Both schizophrenia and CVD risk factors are heritable, and large genome-wide association studies (GWAS) have reported single nucleotide polymorphisms (SNPs) associated with schizophrenia (Ripke et al., 2011) as well as CVD risk factors, including systolic and diastolic blood pressure (Ehret et al., 2011), low- and high-density lipoprotein cholesterol and triglycerides (Teslovich et al., 2010), body mass index (Speliotes et al., 2010), and waist to hip ratio (Heid et al., 2010). Complex traits such as schizophrenia and CVD risk factors are also highly polygenic (Glazier et al., 2002; Hirschhorn and Daly, 2005; Hindorf et al., 2009; Manolio et al., 2009; Yang et al., 2010), and there is evidence that many of these genetic risk factors overlap across traits (Sivakumaran et al., 2011; Andreassen et al., 2013), i.e., are *pleiotropic*.

We investigate the shared genetic mechanisms of schizophrenia and CVD risk factors by analyzing summary statistics from large independent GWAS of schizophrenia (SCZ), triglycerides (TG), and systolic blood pressure (SBP). The goal of these analyses is to estimate the conditional probability that a given SNP is non-null for schizophrenia given its observed association with all three phenotypes simultaneously. This conditional version of the local false discovery rate (local fdr; Efron, 2007) is biologically informative about the shared genetic architecture of schizophrenia and CVD risk factors, as well as potentially increasing power to discover new genetic loci related to schizophrenia (Andreassen et al., 2013). Estimation of the conditional local fdr requires estimation of the conditional density of the schizophrenia summary statistics given values of the summary statistics of the pleiotropic phenotypes. As noted above, we accomplish this by modeling the marginal densities as a mixture of B-spline densities and linking the marginal densities via a Gaussian copula.

Copulas are defined as follows. Let $Y_1, \dots, Y_p$ be random variables with cumulative distribution functions (cdf) $F_1(y_1), \dots, F_p(y_p)$, respectively. Sklar (1959) showed that there always exists a copula C , such that

$$F(y_1, \dots, y_p) = C(F_1(y_1), \dots, F_p(y_p)). \quad (1)$$

If $F_1, \dots, F_p$ are continuous, then C is unique. Note that C is a cdf for p uniform random variables, and as such is a function from $[0, 1]^p$ onto $[0, 1]$. A copula models the dependence structure among the $Y_j, j = 1, \dots, p$, but does not determine the marginal distributions. It can therefore be used as a tool for devising new multivariate distributions.

Genest and Favre (2007) discuss estimation methods for copula parameters, primarily in the bivariate case. Given pairs $(y_{11}, y_{12}, \dots, (y_{n1}, y_{n2})$, they describe estimation of copula parameters using rank-based estimators. The justification for using rank-based methods is that the dependence structure captured by a copula is not related to the marginal distributions and is invariant under strictly monotonic transformations of the variables. This is a two-step approach, where the marginal cdfs are first estimated by their empirical counterparts, and the copula parameters are then estimated by maximizing the rank-based log likelihood. Unlike the method of moments, maximum pseudolikelihood extends naturally to the multivariate case. As Genest and Favre (2007) acknowledge, there is no consensus in the literature on using only rank-based methods for estimating copula parameters. They also give references for kernel-based nonparametric estimation of the copula density. Chen et al. (2006) consider simultaneous estimation of both the marginal cdfs and the copula parameters. They argue that in general, empirical cdfs are inefficient estimators of the marginal cdfs. Furthermore, except for a few special cases, the two-step approach mentioned above is in general inefficient. Chen et al. (2006) propose a general sieve maximum likelihood estimation procedure for all the unknown parameters, where the marginal densities are approximated by linear combinations of finite-dimensional known basis functions with increasing complexity (sieves). Hoff (2007) proposes the extended-rank likelihood. For continuous data, this likelihood is equivalent to the distribution of the multivariate ranks. As a function of the parameters, the extended likelihood depends only on the copula parameters, not on the parameters of the marginal distributions. Pitt et al. (2006) use a Gaussian copula to handle multivariate dependence, assuming that the marginal distributions are specified. Their Bayesian formulation uses latent variables to transform each of the marginals to a standard normal distribution, and a multivariate normal distribution is assumed for these latent variables. The response variables can be either discrete or continuous or a combination of both. Pitt et al. (2006) also extend the idea of covariance selection to Gaussian copula models. Song et al. (2009) utilize Gaussian copulas to combine separate univariate generalized linear models into a joint regression model accommodating continuous, discrete or mixed outcomes. They specifically consider multidimensional logistic regression and a joint model for mixed normal and binary outcomes. Inference is based on maximum likelihood. Kolev and Paiva (2009) give a survey of copula-based regression models which includes the Gaussian copula regression model of Pitt et al. (2006), transition regression models and longitudinal models. Smith et al. (2010) model the dependence structure of continuous time series data using a sequence of bivariate copulas termed pair-copulas and take a Bayesian approach to inference.

Our proposed method is closest in spirit to that of Chen et al. (2006), described above, in that both the copula parameters and the parameters of the marginal distributions are estimated simultaneously. While Chen et al. (2006) use maximum likelihood estimation, our approach is Bayesian. Chen et al. (2006) first approximate the unknown marginal density functions by appropriate sieves, and then the sieve MLEs are obtained by maximization over

a finite-dimensional parameter space. In their method, the smoothing parameters which control the smoothness of the estimated densities are selected by cross validation. Our method allows for joint estimation of all the parameters, including the smoothing parameters. Another emphasis of our proposed method is on conditional density estimation, i.e., the estimation of $f(Y_1/Y_2 = y_2, \dots, Y_p = y_p)$. We derive an explicit formula for the conditional distribution corresponding to the Gaussian copula and use it to obtain the conditional density in closed form.

The remainder of the paper is organized as follows. Section 2 provides a brief background on Gaussian copulas. Section 3 presents our proposed model, the priors and an outline of the sampling scheme. Section 4 studies the frequentist properties of the model via simulation, and Section 5 applies our model to data obtained from genome-wide association studies of schizophrenia and cardiovascular risk factors. We conclude with some remarks and future directions in Section 6. Technical details are given in the appendices.

2 The Gaussian Copula

This paper focuses on the Gaussian copula (Song (2000)), which belongs to the family of elliptical copulas (Fang et al. (2002)). Let $\mathbf{u} = (u_1, \dots, u_p)'$ such that $\mathbf{u} \in [0, 1]^p$, then the cdf of the Gaussian copula is

$$C(u_1, \dots, u_p) = \Phi_p(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_p); \Omega), \quad (2)$$

where $\Phi_p(\cdot; \Omega)$ is the cdf of a p -variate normal distribution with mean $\mathbf{0}$ and correlation matrix Ω , and Φ^{-1} is the inverse cdf of the standard normal distribution. Taking the derivative $C(\mathbf{u})/\mathbf{u}$ and evaluating it at $u_1 = F_1(y_1), \dots, u_p = F_p(y_p)$ results in the Gaussian copula pdf

$$c(\mathbf{u}) = |\Omega|^{-1/2} \exp \left\{ -\frac{1}{2} \mathbf{q}' (\Omega^{-1} - I_p) \mathbf{q} \right\}, \quad (3)$$

where $\mathbf{q} = (q_1, \dots, q_p)'$, $q_j = \Phi^{-1}(F_j(y_j))$, $j = 1, \dots, p$, and I_p is the $p \times p$ identity matrix.

From Equation (1), it follows that the joint probability density function of $Y_1, \dots, Y_p$ is

$$f(y_1, \dots, y_p) = c(F_1(y_1), \dots, F_p(y_p)) \prod_{i=1}^p f_i(y_i). \quad (4)$$

Partitioning Ω as

$$\Omega = \begin{pmatrix} 1 & \omega' \\ \omega & \Omega_{11} \end{pmatrix},$$

where Ω_{11} is the correlation matrix of $Y_2, \dots, Y_p$ and $\omega = (\omega_{12}, \dots, \omega_{1p})'$, the conditional density of Y_1 given $Y_2, \dots, Y_p$ is

$$f(y_1|y_2, \dots, y_p) = \frac{1}{\sigma} f_1(y_1) \exp \left\{ -\frac{1}{2} \left[\frac{(q_1 - \omega' \Omega_{11}^{-1} \mathbf{q} - 1)^2}{\sigma^2} - q_1^2 \right] \right\}. \quad (5)$$

In (5), $f_1(y_1)$ is the marginal density of Y_1 , $\mathbf{q}_{-1} = (q_2, \dots, q_p)'$ and $\sigma^2 = 1 - \omega' \Omega_{11}^{-1} \omega$. Käärik and Käärik (2009) give without proof a less explicit version of the conditional density (5), while Crane and Van Der Hoek (2008) provide general conditional expectation formulas for continuous copulas. For a proof of (5), see Appendix A.

3 The model, priors and sampling scheme

3.1 The model

Similar to Chen et al. (2006), we propose simultaneous estimation of the marginal pdfs along with the copula parameters. We take a Bayesian approach and focus on the Gaussian copula. Given n observations on p continuous margins, $y = (y_1', \dots, y_n')'$, it follows from equations (3) and (4) that the likelihood is

$$L(\Theta, \Omega; y) = |\Omega|^{-n/2} \exp \left\{ -\frac{1}{2} \sum_{i=1}^n \mathbf{q}_i' (\Omega^{-1} - I) \mathbf{q}_i \right\} \prod_{i=1}^n f_1(y_{i1}) \dots f_p(y_{ip}), \quad (6)$$

where $\mathbf{q}_i$ denotes the vector $\mathbf{q}$ evaluated at the i th observation. The marginal densities are expressed nonparametrically as mixtures of quadratic B-spline densities, which are regular B-splines that integrate to 1 over their support. This idea was first proposed by Ghidey et al. (2004) for modeling the distribution of the random effects in linear mixed models. In fact, Ghidey et al. (2004) used Gaussian densities as the basis functions, citing a result on the convergence of a B-spline of degree q to a normal density, as $q \rightarrow \infty$. Staudenmayer et al. (2008) used the same idea with B-spline densities for density estimation in the presence of heteroscedastic measurement error. Pitt et al. (2006) use a Bayesian approach to copula modeling but assume known parametric margins.

More specifically, to facilitate nonparametric estimation of the marginal pdfs, we express $f_j(y_{ij})$, $i = 1, \dots, n$, $j = 1, \dots, p$, as

$$f_j(y_{ij}) = \sum_{k=1}^{K_j} c_{jk} g_k(y_{ij}), \quad (7)$$

where $c_{jk} = \exp(\gamma_{jk}) / \sum_{l=1}^{K_j} \exp(\gamma_{jl})$ and $g_k(\cdot)$ are the B-spline densities. The number K_j of B-spline densities, specific to f_j is fixed throughout the estimation process. In our

experience, K_j of up to 30 is usually adequate for estimating the marginal densities. Schellhase and Kauermann (2012) cite asymptotic results but recommend that K_j be large enough but usually small compared to the sample size, which is the rule of thumb advocated by Ruppert (2002). This is illustrated in Section 4 via simulation.

The corresponding knots are equally spaced along $y_{1j}, \dots, y_{nj}$, $j=1, \dots, p$. The γ_{jk} , $j=1, \dots, p$, $k=1, \dots, K_j$ are unknown parameters, and for identifiability, γ_{j1} , $j=1, \dots, p$, are set to zero.

Although the γ_{jk} s are unconstrained parameters, the c_{jk} s satisfy $c_{jk} \in (0, 1)$ and

$\sum_{k=1}^{K_j} c_{jk} = 1$, $j=1, \dots, p$. The marginal cdfs are

$$F_j(y_{ij}) = \sum_{k=1}^{K_j} c_{jk} G_k(y_{ij}), \quad (8)$$

where $G_k(\cdot)$, $k=1, \dots, K_j$ are the B-spline cdfs corresponding to the $g_k(\cdot)$ s.

3.2 Priors

Priors on the γ_j s—Let $\gamma_j = (\gamma_{j2}, \dots, \gamma_{jK_j})'$. We assume *a priori* independent γ_j , $j=1, \dots, p$. To obtain smooth fits for the marginal pdfs, the c_{jk} s corresponding to neighboring B-splines must be close to one another (Eilers and Marx (1996)). This can be achieved by constraining the corresponding γ_{jk} s to be close to one another. By analogy to Eilers and Marx (1996)'s idea, Lang and Brezger (2004) require the γ_{jk} s to satisfy $\gamma_{j,\rho} = 2\gamma_{j,\rho-1} - \gamma_{j,\rho-2} + u_\rho$, where $u_\rho \sim N(0, \tau_j^2)$, and $p(\gamma_{j2}) \propto 1$, $p(\gamma_{j3}) \propto 1$. However, this results in an improper prior distribution on γ_j similar to the intrinsic Gaussian Markov random field prior used in spatial statistics. Chib and Jeliazkov (2006) place a joint normal prior on γ_{j2} and γ_{j3} which fixes the impropriety of the prior on γ_j . For example, if $(\gamma_{j2}, \gamma_{j3})' \sim N_2(0, c \tau_j^2 \mathbf{I}_2)$, where c is a fixed constant, the prior on γ_j becomes

$$p(\gamma_j | \tau_j^2) \propto (\tau_j^2)^{-\frac{1}{2}(K_j-1)} \exp\left\{-\frac{1}{2\tau_j^2} \left[\sum_{\rho=4}^{K_j} (\gamma_{j,\rho} - 2\gamma_{j,\rho-1} + \gamma_{j,\rho-2})^2 + c^{-1} \gamma_{j,2:3}' \gamma_{j,2:3} \right]\right\}, \quad (9)$$

where $\gamma_{j,2:3}$ is a vector consisting of the first two entries of γ_j . The summation in the exponent on the right-hand side of (9) can be expressed as $\gamma_j' P_j \gamma_j$, where $P_j = D_j' D_j$ and D_j is the $(K_j - 3) \times (K_j - 1)$ matrix

$$\begin{pmatrix} 1 & -2 & 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & -2 & 1 & 0 & \dots & 0 \\ \vdots & \dots & & & & & \vdots \\ 0 & 0 & \dots & 1 & -2 & 1 \end{pmatrix}$$

Equation (9) can now be re-expressed as

$$p(\gamma_j | \tau_j^2) \propto (\tau_j^2)^{-\frac{1}{2}(K_j-1)} \exp \left\{ -\frac{1}{2\tau_j^2} \gamma_j' P_j^* \gamma_j \right\},$$

where $P_{j,l}^* = P_{j,l} + c^{-1}$, for $l=1, 2$.

Priors on the τ_j s—The τ_j s are assumed independent Half- $t(\nu_j, A_j)$ (Gelman, 2006) with pdf $p(x) \propto [1 + (x/A_j)^2/\nu_j]^{-(\nu_j+1)/2}$, $x > 0$, where the hyperparameters ν_j and A_j are assumed known. The larger the value of A_j , the less informative the prior is. In our computations, $A_j = 10$ and $A_j = 10^5$ gave indistinguishable results. Computationally, it is convenient to utilize the following scale mixture representation (Wand et al., 2012) $(\tau_j^2 | a_j) \sim IG(\nu_j/2, \nu_j/a_j)$, $a_j \sim IG(1/2, 1/A_j^2)$, where $IG(a, b)$, is the inverse Gamma distribution with pdf $p(x) \propto x^{-(a+1)} \exp(-b/x)$, $x > 0$.

Prior on Ω —Following Danaher and Smith (2011), we express Ω as

$$\Omega = \text{diag} \left(\sum \right)^{-1/2} \sum \text{diag} \left(\sum \right)^{-1/2}, \quad (10)$$

where Σ is a non-unique positive definite matrix, and $\text{diag}(\Sigma)$ is a diagonal matrix consisting of the leading diagonal of Σ . The matrix Σ^{-1} is decomposed as $\Sigma^{-1} = LL'$, where L is a lower triangular Cholesky factor with ones on its diagonal. The priors on the elements of L are independent $N(0, \sigma_L^2)$, where σ_L^2 is a fixed large number, reflecting vague knowledge on the elements of L . Alternative approaches include Chan and Jeliazkov (2009) who use the modified Cholesky decomposition and impose constraints on elements of the matrices making up this decomposition. Daniels and Pourahmadi (2009) parametrize correlation matrices in terms of partial autocorrelations as proposed by Joe (2006). These partial autocorrelations take values in the interval $[-1, 1]$ which makes them easy to work with. When the number of variables is large, it may be more appropriate to parameterize the correlation matrix parsimoniously via variable selection methods as in Pitt et al. (2006), for example.

3.3 The Sampling Scheme

This section gives an outline of the MCMC sampling scheme. More details are given in Appendix B. To simplify the sampling scheme, we introduce unobservable mixture indicators, z_{ijk} , such that $z_{ijk} = 1$ if y_{ij} came from the k th spline density. The sampling scheme consists of the following steps

1. Sample γ_j from $p(\gamma_j | z_j)$, $j = 1, \dots, p$, where $z_j = \{z_{ijk}, i = 1, \dots, n, k = 1, \dots, K_j\}$, using a Metropolis-Hastings step.
2. Sample τ_j^2 from $p(\tau_j^2 | \gamma_j, a_j)$, $j = 1, \dots, p$.
3. Sample a_j from $p(a_j | \tau_j^2)$, $j = 1, \dots, p$.

4. Sample Ω from $p(\Omega|\{\gamma_j\}_{j=1}^p, \text{data})$.
5. Sample the indicators, z_{ijk} , from
 $p(z_{ijk}|\{\gamma_j\}_{j=1}^p, \text{data}), i=1, \dots, n, j=1, \dots, p, k=1, \dots, K_j$.

4 Simulation Study

In this section we study the frequentist properties of our model by Monte Carlo simulation studies. Two simulation settings are considered: (i) samples are generated from a trivariate normal distribution with mean vector $(3, 1, 2)'$ and correlation matrix

$$\Omega = \begin{pmatrix} 1 & 0 & .5774 \\ 0 & 1 & -.5774 \\ .5774 & -.5774 & 1 \end{pmatrix}, \quad (11)$$

and (ii) samples are generated from a trivariate Gaussian copula with marginal distributions

$$\begin{aligned} F_1(y_1) &= 0.3\Phi(y_1) + 0.7\Phi(y_1 - 5), \quad -\infty < y_1 < \infty, \\ F_2(y_2) &= 1 - \exp(-y_2/3), \quad y_2 > 0, \\ F_3(y_3) &= \Phi(y_3), \quad -\infty < y_3 < \infty \end{aligned} \quad (12)$$

and the correlation matrix (11). From each setting, 50 samples of size 10^3 and 50 samples of size 10^4 are generated. For each sample, a total of 10,000 MCMC iterations are run with a burn-in period of 2,000. To evaluate the results, we utilize the Kullback-Leibler (KL) divergence

$$KL = - \int h(t) \log \frac{\hat{h}(t)}{h(x)} dt,$$

where $h(\cdot)$ is a probability density function and $\hat{h}(\cdot)$ is its estimate. In our case, $h(\cdot)$ is either a marginal density, $f(y_j)$, $j = 1, 2, 3$ or a conditional density, $f(y_1/y_2, y_3)$. In addition, the distance between the true correlation matrix Ω and its estimate is assessed using Stein's loss

$$\text{trace}(\hat{\Omega}\Omega^{-1}) - \log(\det(\hat{\Omega}\Omega^{-1})) - p,$$

see, for example, Daniels and Kass (2001). Each column in Figure 1 presents, for each of the three marginal distributions, boxplots of the KL divergence between the true and the estimated marginal densities. The first two columns of boxplots from the left correspond, respectively, to samples of size $n = 10^3$ and $n = 10^4$ from setting (i), while the last two columns correspond to setting (ii). The KL values reflect good estimation of the marginal densities in both settings. For $n = 10^4$, these values are less variable.

Figure 2 demonstrates that the fitted marginal densities of setting (ii) are not sensitive to the number of B-spline densities used, provided this number, K_j , is large enough, as discussed in

Section 3.1. This figure consists of boxplots of the KL divergence between the true and estimated marginal densities for $K_j = 5, 10, 20, 30, 60, j = 1, 2, 3$, based on 50 samples of size $n = 10^3$. As can be seen in these boxplots, $K_j = 5$ is somewhat too small, especially for estimating the mixture in (12), but $K_j = 20$ already seems sufficient. Moreover, adding more basis functions does not result in overfitting.

Figure 3 displays boxplots of the KL divergence between the true and estimated conditional densities, organized as in Figure 1. Each column includes nine boxplots corresponding to conditioning on quantiles of Y_2 and Y_3 according to

	1	2	3	4	5	6	7	8	9
Quantile of Y_2	25	50	75	25	50	75	25	50	75
Quantile of Y_3	25	25	25	50	50	50	75	75	75

For example, the fourth boxplot in each panel corresponds to conditioning on the 25th quantile of Y_2 and the 50th quantile of Y_3 . It is seen in Figure 3 that the KL divergences corresponding to $n = 10^4$ are less variable than the ones corresponding to $n = 10^3$. Also, the KL divergences are less variable for setting (i) compared to setting (ii).

Figure 4 displays boxplots of the Stein loss values for setting (i) (first two boxplots from the left) and for setting (ii). From these plots, it is evident that the methodology provides good estimates of the underlying correlation matrix. The Stein loss values are less variable for $n = 10^4$.

Source code to reproduce the results of this section and of Section 5 is available as Supporting Information on the journal's web page (<http://onlinelibrary.wiley.com/doi/xxx/supinfo>).

5 Application

We investigate whether there are shared genetic mechanisms between schizophrenia (SCZ) and CVD risk factors by analyzing summary statistics from independent studies. We obtained two-tailed p -values from a schizophrenia GWAS ($n=21,856$ subjects; Ripke et al. 2011) and from independent GWAS on TG ($n=96,568$ subjects; Teslovich et al. 2010) and SBP ($n=203,056$ subjects; Ehret et al. 2011). We selected all assayed SNPs in linkage disequilibrium with one or more 5' untranslated region loci, a category shown to be highly enriched for non-null effects across many phenotypes (Schork et al., 2013), resulting in 132,765 SNPs assessed in all three phenotypes simultaneously.

Since two-tailed p -values were made publicly available, we were able to recover the absolute z -scores of the test statistics but not their sign. For ease of notation, we denote absolute z -scores by $z = (z_1, z_2, z_3)'$, where $z_j = \Phi^{-1}(1 - .5p_j)$, $j = 1, 2, 3$ and 1=SCZ, 2=SBP and 3=TG. Here, p_j is the two-tailed p -value for the j th phenotype and Φ is the standard normal cdf. Pairwise scatterplots of the data are shown in Figure 5. Histograms of absolute z -scores for each phenotype are displayed in the top panel of Figure 6, along with the fitted marginal densities from the proposed model estimated using 20 equally spaced knots for each phenotype.

The estimates (posterior means) of the correlations making up the matrix Ω are $\hat{\omega}_{12} = .0463$, $\hat{\omega}_{13} = .0659$ and $\hat{\omega}_{23} = .0499$. The 95% credible intervals for ω_{12} , ω_{13} and ω_{23} are [.0455, .0463], [.0654, .0659] and [.0389, .0499], respectively. The lower and upper limits for these credible intervals are obtained as the 0.025 and 0.975 percentiles of the MCMC iterates of the correlations. Thus, quantiles of TG and SBP z -scores are positively correlated with quantiles of SCZ z -scores, demonstrating that SCZ has a partially overlapping genetic architecture with these two phenotypes.

Of primary interest is whether the probability that a SNP has a non-null association with SCZ depends on the strength of its relationship with TG and/or SBP. To investigate this question further, we utilize the local false discovery framework of Efron (2007). The local false discovery rate (*local fdr*) is defined as the posterior probability that a SNP is null for SCZ given its observed absolute z -score. By Bayes' rule this is given by

$$\text{fdr}(z_1) = \frac{\pi_0 \phi(z_1)}{f_1(z_1)} \quad (13)$$

where ϕ is the folded standard normal density (i.e., the distribution of the absolute value of a standard normal random variable), π_0 is the proportion of SNPs that are null for SCZ, and f_1 is the density of z_1 . Use of the folded standard normal density is equivalent to the *theoretical null* of Efron (2007). This is justified in the current example since the summary statistics were made publicly available after performing usual genomic control procedures (Devlin and Roeder, 1999). A conservative estimate of the *fdr* is produced by assuming $\pi_0 = 1$ and utilizing the nonparametric estimate $\hat{f}_1$ of f_1 produced from our proposed model. The assumption that $\pi_0 = 1$ in Equation (13) is reasonable, since even for highly polygenic traits such as SCZ, the proportion of null SNPs is considerably greater than 0.9 (Andreassen et al., 2013). Larger z -scores will have a higher probability of association with the phenotypes of interest. Here, a z -score cut-off of $z_1 = 3.3$ produces estimated $\hat{\text{fdr}}(z_1) \leq .2$. The bottom panel of Figure 6 shows similar histograms to those in the top panel but for absolute z -scores ≥ 3.3 .

The *conditional local fdr* is given by

$$\text{fdr}(z_1|z_2, z_3) = \frac{\pi_0(z_2, z_3) \phi(z_1)}{f(z_1|z_2, z_3)}, \quad (14)$$

where $\pi_0(z_2, z_3)$ is the conditional probability that z_1 is null given (z_2, z_3) and $f(z_1|z_2, z_3)$ is the conditional density of z_1 given (z_2, z_3) . By setting $\pi_0(z_2, z_3) = 1$ in Equation (14) and using an estimate $\hat{f}(z_1|z_2, z_3)$ based on Equation (5), we obtain a conservative estimate $\hat{\text{fdr}}(z_1|z_2, z_3) = \phi(z_1) / \hat{f}(z_1|z_2, z_3)$. Conditional density estimates for $z_1 \geq 3.3$ given quantiles of z_2 and z_3 are displayed in Figure 7, and conditional local *fdr* estimates are given in Figure 8. For a given z -score, the local *fdr* is lower when conditioning on larger quantiles of z_2 and z_3 . For example, for $z_1 = 3.5$, the unconditional *fdr* estimate is $\hat{\text{fdr}}(3.5) = .1364$,

whereas $\widehat{\text{fdr}}(3.5|z_2 = 3.03, z_3 = 3.63) = .08$, where $z_2 = 3.03$ and $z_3 = 3.63$ are the 99th quantiles of SBP and TG absolute z-scores, respectively. Using the commonly-used threshold of .2, there were 1,594 significant SNPs using local fdr, whereas there were 1,715 SNPs using conditional local fdr, for an increase of 7.6%.

To evaluate the goodness of fit of the Gaussian copula, we use the $S_n^{(B)}$ statistic given in Genest et al. (2009). This is a Cramér-von Mises-type statistic which measures the distance between the multivariate empirical distribution function and the independence copula, obtained by using Rosenblatt's probability integral transform (Rosenblatt, 1952). For more details, see Genest et al. (2009). The calculation of $S_n^{(B)}$ is computer intensive, and so to facilitate its computation we analyze a subset of the data obtained by sampling randomly 1000 observations from the full data set. At each MCMC iteration $S_n^{(B)}$ is computed, thus yielding an approximation to the posterior distribution of $S_n^{(B)}$. Figure 9 displays a density histogram of the $S_n^{(B)}$ iterates. The 99%th empirical quantile of these values is equal to 0.36 which indicates good fit of the Gaussian copula to the data.

6 Discussion

In this paper we propose a Bayesian methodology for estimating multivariate distributions wherein the marginal densities are estimated nonparametrically and linked via a Gaussian copula.

Various extensions and applications of our model are possible. For example, the use of Gaussian copulas assumes no tail dependence between pairs of variables (Embrechts et al. (2003)). Upper tail dependence between random variables Y_1 and Y_2 , is defined as

$$\lim_{u \uparrow 1} \Pr(Y_2 > F_2^{-1}(u) | Y_1 > F_1^{-1}(u)). \quad (15)$$

Upper tail dependence is a measure of association in the upper-right quadrant, whereas lower tail dependence is defined analogously for the lower-left quadrant. Both upper and lower tail dependence assess the strength of relationships between extreme events of Y_1 and Y_2 . While Gaussian copulas are upper- and lower-tail independent, other families of elliptical copulas (Fang et al. (2002)) are not. For example, t -copulas are a family of elliptical copulas allowing for tail dependence. Implementing our model using other families of elliptical copulas may provide better fits for random variables exhibiting dependence in extreme events.

Other extensions are possible. For example, the proposed model currently handles only continuous margins. Recent work has proposed Bayesian estimation of copula models with discrete margins by augmentation with uniform latent variables (Smith and Khaled (2012)). Additionally, Pitt et al. (2006) developed a graphical model selection prior on the Gaussian copula correlation matrix.

Finally, it is possible to adapt the copula model to estimate the multivariate local false discovery rate (Andreassen et al., 2013). This involves an extension of Efron's (2007) two-group mixture model to a 2^P -group mixture model, where P is the number of phenotypes of interest. Andreassen et al. (2013) have shown that bivariate local false discovery rate is potentially much more powerful than univariate local false discovery rate, and highly informative biologically.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

O. Rosen was supported in part by NSF grant DMS-0804140 and by the National Security Agency under Grant Number H98230-12-1-0246. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein.

W. Thompson was supported by NIH grants R01HD061414, 5R01AG022381, R01MH092793 and R01GM104400-01A. This project was also supported by grant 2G12MD007592 from the National Institute on Minority Health and Health Disparities (NIMHD). We are also grateful to the referees for their suggestions that helped improve this paper.

Appendix A

Conditional Density Formula for Gaussian Copula

Using Equation (4), it follows that

$$f(y_1|y_2, \dots, y_p) = f_1(y_1) \frac{c(F_1(y_1), \dots, F_p(y_p))}{c_{-1}(F_2(y_2), \dots, F_p(y_p))}, \quad (16)$$

where c is the pdf (3), and c_{-1} is the copula density with correlation matrix Ω_{11} , defined in Section 2. Differentiating (2), substituting $q_j = \Phi^{-1}(F_j(y_j))$ and plugging into (16) results in

$$f(y_1|y_2, \dots, y_p) = f_1(y_1) \frac{\phi_p(q_1, \dots, q_p)}{\phi(q_1)\phi_{p-1}(q_2, \dots, q_p)}, \quad (17)$$

where ϕ_p and ϕ_{p-1} are the densities corresponding to $N(\mathbf{0}_p, \Omega)$ and $N(\mathbf{0}_{p-1}, \Omega_{11})$, respectively, and ϕ is the density of the univariate standard normal. The notation $\mathbf{0}_p$ means a vector of p zeros. Plugging the normal densities into (17) leads to

$$f(y_1|y_2, \dots, y_p) = f_1(y_1) \frac{|\Omega|^{-1/2}}{|\Omega_{11}|^{-1/2}} \exp \left\{ -\frac{1}{2} \mathbf{q}' \Omega^{-1} \mathbf{q} + \frac{q_1^2}{2} + \frac{1}{2} \mathbf{q}'_{-1} \Omega_{11}^{-1} \mathbf{q}_{-1} \right\}. \quad (18)$$

Applying Schur decomposition to Ω^{-1} gives

$$\Omega^{-1} = \begin{pmatrix} 1 & \mathbf{0}'_{p-1} \\ -\Omega_{11}^{-1}\omega & I_{p-1} \end{pmatrix} \begin{pmatrix} (1 - \omega' \Omega_{11}^{-1} \omega)^{-1} & \mathbf{0}'_{p-1} \\ \mathbf{0}_{p-1} & \Omega_{11}^{-1} \end{pmatrix} \begin{pmatrix} 1 & -\omega' \Omega_{11}^{-1} \\ \mathbf{0}_{p-1} & I_{p-1} \end{pmatrix}, \quad (19)$$

where I_{p-1} is the $(p-1) \times (p-1)$ identity matrix. Using (19) implies

$$\mathbf{q}' \Omega^{-1} \mathbf{q} = (q_1 - \omega' \Omega_{11}^{-1} \mathbf{q}_{-1})^2 (1 - \omega' \Omega_{11}^{-1} \omega)^{-1} + \mathbf{q}'_{-1} \Omega_{11}^{-1} \mathbf{q}_{-1} \quad (20)$$

and

$$|\Omega|^{-1} = (1 - \omega' \Omega_{11}^{-1} \omega)^{-1} |\Omega_{11}|^{-1}. \quad (21)$$

Plugging expressions (20) and (21) into (18) gives the conditional density (5).

Appendix B

Details of the Sampling Scheme

Section 3.3 describes the outline of the sampling scheme. This appendix provides some more details.

1. Sampling the γ_j

The augmented conditional posterior of γ_j is

$$p(\gamma_j | \{\gamma_l\}_{l \neq j}, \Omega, \text{data}, z_j) \propto \exp \left\{ -\frac{1}{2} \sum_{i=1}^n q_i' (\Omega^{-1} - I) q_i \right\} \prod_{i=1}^n \prod_k^{K_j} c_{jk}^{z_{ijk}} \\ \times \exp \left\{ -\frac{1}{2\tau_j^2} \gamma_j' P_j^* \gamma_j \right\},$$

where $z_j = \{z_{ijk}, i = 1, \dots, n, k = 1, \dots, K_j\}$. The marginal pdfs and cdfs $f_j(y_{ij})$ and $F_j(y_{ij})$ are modeled as in (7) and (8), respectively. With the latent indicators,

$z_j, q_{ij} = \Phi^{-1} \left(\prod_{k=1}^{K_j} (G_k(y_{ij}))^{z_{ijk}} \right)$. Since this is not a function of γ_j , the log conditional posterior of γ_j reduces to

$$\log p(\gamma_j | z_j) \stackrel{c}{=} \sum_k^{K_j} n_{jk} \log c_{jk} - \frac{1}{2\tau_j^2} \gamma_j' P_j^* \gamma_j, \quad (22)$$

where $\stackrel{c}{=}$ denotes equality up to a constant, and $n_{jk} = \sum_{i=1}^n z_{ijk}$. To sample γ_j , (22) is maximized at each iteration with respect to γ_j . The maximizer and the negative inverse Hessian of (22), evaluated at the maximizer, are used in a

Metropolis-Hastings step with a multivariate t distribution with small degrees of freedom (e.g., 4) as the proposal distribution.

2. Sampling the τ_j^2

The $\tau_j^2, j=1, \dots, p$, are sampled from $IG((K_j + \nu_j - 1)/2, \nu_j/a_j + \frac{1}{2}\gamma_j' P_j^* \gamma_j)$.

3. Sampling the a_j

The $a_j, j=1, \dots, p$, are sampled from $IG((\nu_j + 1)/2, \nu_j/\tau_j^2 + 1/A_j^2)$.

4. Sampling Ω

$$\begin{aligned} p(\Omega | \{\gamma_j\}_{j=1}^p, \text{data}) &\propto |\Omega|^{-n/2} \exp \left\{ -\frac{1}{2} \sum_{i=1}^n \mathbf{q}_i' (\Omega^{-1} - I) \mathbf{q}_i \right\} p(\Omega) \\ &\propto |\Omega|^{-n/2} \exp \left\{ -\frac{1}{2} \text{tr}(S\Omega^{-1}) \right\} p(\Omega), \end{aligned}$$

where $p(\Omega)$ is the prior on Ω , $S = \sum_{i=1}^n \mathbf{q}_i \mathbf{q}_i'$ and $q_{ij} = \Phi^{-1}(\sum_{k=1}^K c_{jk} G_k(y_{ij}))$.

The matrix Ω is sampled in a random-walk Metropolis step via the decomposition described in Section 3.2. The elements of L are sampled individually using a normal proposal centered at the current value.

5. Sampling the indicators

The indicators are sampled as Bernoulli random variables with probability

$$p(z_{ijk}=1 | \text{data}, \{\gamma_j\}_{j=1}^p) \propto \exp \left\{ -\frac{1}{2} \mathbf{q}_i' (\Omega^{-1} - I) \mathbf{q}_i \right\} c_{jk} g_k(y_{ij}),$$

where $q_{ij} = \Phi^{-1}(\prod_{k=1}^K (G_k(y_{ij}))^{z_{ijk}})$.

References

- Andreassen O, Djurovic S, Thompson W, Schork A, Kendler K, O'Donovan M, Rujescu D, Werge T, van de Bunt M, Morris A, McCarthy M, Roddey J, McEvoy L, Desikan R, Dale A. Improved detection of common variants associated with schizophrenia by leveraging pleiotropy with cardiovascular disease risk factors. *American Journal of Human Genetics*. 2013; 92:197–209. [PubMed: 23375658]
- Chan JC, Jeliazkov I. MCMC estimation of restricted covariance matrices. *Journal of Computational and Graphical Statistics*. 2009; 18:457–480.
- Chen X, Fan Y, Tsyrennikov V. Efficient estimation of semiparametric multivariate copula models. *Journal of the American Statistical Association*. 2006; 101:1228–1240.
- Chib S, Jeliazkov I. Inference in semiparametric dynamic models for binary longitudinal data. *Journal of the American Statistical Association*. 2006; 101:685–700.
- Colton C, Manderscheid R. Congruencies in increased mortality rates, years of potential life lost, and causes of death among public mental health clients in eight states. *Preventing Chronic Diseases*. 2006; 3:A42.
- Crane G, Van Der Hoek J. Conditional expectation formulae for copulas. *Australian and New Zealand Journal of Statistics*. 2008; 50:53–67.

- Danaher P, Smith M. Modeling multivariate distributions using copulas: applications in marketing. *Marketing Science*. 2011; 30:4–21.
- Daniels M, Kass R. Shrinkage estimators for covariance matrices. *Biometrics*. 2001; 57:1173–1184. [PubMed: 11764258]
- Daniels M, Pourahmadi M. Modeling covariance matrices via partial autocorrelations. *Journal of Multivariate Analysis*. 2009; 100:2352–2363. [PubMed: 20161018]
- De Hert M, van Winkel R, Van Eyck D, Hanssens L, Wampers M, Scheen A, Peuskens J. Prevalence of the metabolic syndrome in patients with schizophrenia treated with antipsychotic medication. *Schizophrenia Research*. 2006; 83:87–93. [PubMed: 16481149]
- Devlin D, Roeder K. Genomic control for association studies. *Biometrics*. 1999; 55:997–1004. [PubMed: 11315092]
- Efron B. Size, power and false discovery rates. *The Annals of Statistics*. 2007; 35:1351–1377.
- Ehret G, Munroe P, Rice K, Bochud M, Johnson A, Chasman D, Smith A, Tobin M, Verwoert G, Hwang S, et al. Genetic variants in novel pathways influence blood pressure and cardiovascular disease risk. *Nature*. 2011; 478:103–109. [PubMed: 21909115]
- Eilers P, Marx B. Flexible smoothing with B-splines and penalties. *Statistical Science*. 1996; 11:89–121.
- Embrechts, P., Lindskog, F., McNeil, A. Modeling dependence with copulas and applications to risk management. In: Rachev, S., editor. *Handbook of Heavy Tailed Distributions in Finance*. Elsevier; 2003. p. 329-384. chapter 8
- Fang HB, Fang KT, Kotz S. The meta-elliptical distributions with given marginals. *Journal of Multivariate Analysis*. 2002; 82:1–16. Corrigendum: *Journal of Multivariate Analysis* 94, 222–223, 2005.
- Gelman A. Prior distributions for variance parameters (comment on article by Browne and Draper). *Bayesian Analysis*. 2006; 1:515–534.
- Genest C, Favre AC. Everything you always wanted to know about copula modeling but were afraid to ask. *Journal of Hydrologic Engineering*. 2007; 12:347–368.
- Genest C, Rémillard B, Beaudoin D. Goodness-of-fit tests for copulas: a review and a power study. *Insurance: Mathematics and Economics*. 2009; 44:199–213.
- Ghidey W, Lesaffre E, Eilers P. Smooth random effects distribution in a linear mixed modeling. *Biometrics*. 2004; 60:945–953. [PubMed: 15606415]
- Glazier A, Nadeau J, Aitman T. Finding genes that underlie complex traits. *Science*. 2002; 298:2345–2349. [PubMed: 12493905]
- Hansen T, Ingason A, Djurovic S, Melle I, Fenger M, Gustafsson O, Jakobsen K, Rasmussen H, Tosato S, Rietschel M, et al. At-risk variant in TCF7L2 for type II diabetes increases risk of schizophrenia. *Biological Psychiatry*. 2011; 70:59–63. [PubMed: 21414605]
- Heid I, Jackson A, Randall J, Winkler T, Qi L, Steinthorsdottir V, Thorleifsson G, Zillikens M, Speliotes E, Magi R, et al. Meta-analysis identifies 13 new loci associated with waist-hip ratio and reveals sexual dimorphism in the genetic basis of fat distribution. *Nature Genetics*. 2010; 42:949–960. [PubMed: 20935629]
- Hindorf L, Sethupathy P, Junkins H, Ramos E, Mehta J, Collins F, Manolio T. Potential etiologic and functional implications of genome-wide association loci for human diseases and traits. *Proceedings of the National Academy of Sciences*. 2009; 106:9362–9367.
- Hirschhorn J, Daly M. Genome-wide association studies for common diseases and complex traits. *Nature Review Genetics*. 2005; 6:95–108.
- Hoff P. Extending the rank likelihood for semiparametric copula estimation. *The Annals of Applied Statistics*. 2007; 1:265–283.
- Joe H. Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis*. 2006; 97:2177–2189.
- Käärik E, Käärik M. Modeling dropouts by conditional distribution, a copula-based approach. *Journal of Statistical Planning and Inference*. 2009; 139:3830–3835.

- Kaddurah-Daouk R, McEvoy J, Baillie R, Lee D, Yao J, Doraiswamy P, Krishnan K. Metabolomic mapping of atypical antipsychotic effects in schizophrenia. *Molecular Psychiatry*. 2007; 12:934–945. [PubMed: 17440431]
- Kolev N, Paiva D. Copula-based regression models: a survey. *Journal of Statistical Planning and Inference*. 2009; 139:3847–3856.
- Lang S, Brezger A. Bayesian P-splines. *Journal of Computational and Graphical Statistics*. 2004; 13:183–212.
- Laursen T, Munk-Olsen T, Vestergaard M. Life expectancy and cardiovascular mortality in persons with schizophrenia. *Current Opinion Psychiatry*. 2012; 25:83–88.
- Manolio T, Collins F, Cox N, Goldstein D, Hindorf L, Hunter D, McCarthy M, Ramos E, Cardon L, Chakravarti A, et al. Finding the missing heritability of complex diseases. *Nature*. 2009; 461:747–753. [PubMed: 19812666]
- Marder S, Essock S, Miller A, Buchanan R, Casey D, Davis J, Kane J, Lieberman J, Schooler N, Covell N, et al. Physical health monitoring of patients with schizophrenia. *American Journal of Psychiatry*. 2004; 161:1334–1349. [PubMed: 15285957]
- Mitchell A, Vancampfort D, Sweers K, van Winkel R, Yu W, De Hert M. Prevalence of metabolic syndrome and metabolic abnormalities in schizophrenia and related disorders—a systematic review and meta-analysis. *Schizophrenia Bulletin*. 2011
- Pitt M, Chan D, Kohn R. Efficient Bayesian inference for Gaussian copula regression models. *Biometrika*. 2006; 93:537–554.
- Raphael T, Parsons J. Blood sugar studies in dementia praecox and manic depressive insanity. *Archives of Neurology and Psychiatry*. 1921; 5:687–709.
- Ripke S, Sanders A, Kendler K, Levinson D, Sklar P, Holmans P, Lin D, Duan J, Ophoff R, Andreassen O, et al. Genome-wide association study identifies five new schizophrenia loci. *Nature Genetics*. 2011; 43:969–976. [PubMed: 21926974]
- Rosenblatt M. Remarks on a multivariate transformation. *The Annals of Mathematical Statistics*. 1952; 23:470–472.
- Ruppert D. Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics*. 2002; 11:735–757.
- Ryan M, Collins P, Thakore J. Impaired fasting glucose tolerance in first-episode, drug-naïve patients with schizophrenia. *American Journal of Psychiatry*. 2003; 160:284–289. [PubMed: 12562574]
- Saha S, Chant D, McGrath J. A systematic review of mortality in schizophrenia: is the differential mortality gap worsening over time? *Archives of General Psychiatry*. 2007; 64:1123–1131. [PubMed: 17909124]
- Schellhase C, Kauermann G. Density estimation and comparison with a penalized mixture approach. *Computational Statistics*. 2012; 27:757–777.
- Schork A, Thompson W, Pham P, Torkamani A, Roddey J, Sullivan P, Kelsoe J, Purcell S, O'Donovan M, Schork N, Andreassen O, Dale A. All SNPs are not created equal: Genome-wide association studies reveal a consistent pattern of enrichment among functionally annotated SNPs. *PLoS Genetics*. 2013; 9doi: 10.1371/journal.pgen.1003449
- Sivakumaran S, Agakov F, Theodoratou E, Prendergast J, Zgaga L, Manolio T, Rudan I, McKeigue P, Wilson J, Campbell H. Abundant pleiotropy in human complex diseases and traits. *American Journal of Human Genetics*. 2011; 89:607–618. [PubMed: 22077970]
- Sklar A. Fonctions de répartition à n dimensions et leur marges. *Publ Inst Statist Univ Paris*. 1959; 8:229–231.
- Smith M, Khaled M. Estimation of copula models with discrete margins via Bayesian data augmentation. *Journal of the American Statistical Association*. 2012; 107:290–303.
- Smith M, Min A, Almeida C, Czado C. Modeling longitudinal data using a pair-copula decomposition of serial dependence. *Journal of the American Statistical Association*. 2010; 105:1467–1479.
- Song P. Multivariate dispersion models generated from Gaussian copula. *Scandinavian Journal of Statistics*. 2000; 27:305–320.
- Song P, Li M, Yuan Y. Joint regression analysis of correlated data using Gaussian copulas. *Biometrics*. 2009; 65:60–68. [PubMed: 18510653]

- Speliotes E, Willer C, Berndt S, Monda K, Thorleifsson G, Jackson A, Allen H, Lindgren C, Luan J, Magi R, et al. Association analyses of 249,796 individuals reveal 18 new loci associated with body mass index. *Nature Genetics*. 2010; 42:937–948. [PubMed: 20935630]
- Staudenmayer J, Ruppert D, Buonaccorsi J. Density estimation in the presence of heteroscedastic measurement error. *Journal of the American Statistical Association*. 2008; 103:726–736.
- Teslovich T, Musunuru K, Smith A, et al. Biological, clinical, and population relevance of 95 loci for blood lipids. *Nature*. 2010; 466:707–713. [PubMed: 20686565]
- Wand M, Ormerod J, Padoan S, Frühworth R. Mean field variational Bayes for elaborate distributions. *Bayesian Analysis*. 2012; 7:847–900.
- Yang J, Benyamin B, McEvoy B, Gordon S, Henders A, Nyholt D, Madden P, Heath A, Martin N, Montgomery GW, et al. Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics*. 2010; 42:565–569. [PubMed: 20562875]

Figure 1. Boxplots of the Kullback-Leibler divergence between the true and the estimated marginal densities.

Figure 2. Boxplots of the Kullback-Leibler divergence between the true and the estimated marginal densities of setting (ii) with $n = 10^3$ as a function of the number, K_j , of B-spline densities.

Figure 3. Boxplots of the Kullback-Leibler divergence between the true and the estimated conditional densities.

Figure 4. Boxplots of Stein's loss based on the true and estimated correlation matrices.

Figure 5.
Pairwise scatterplots of the data.

Figure 6.

Top panel from left to right: marginal density histograms and density estimates of SCZ, SBP and TG. The solid lines are density estimates based on the method of Section 3. Bottom panel: similar plots for absolute z -scores 3.3.

Figure 7.
Conditional density estimates showing the tails only ($Z_1 \approx 3.3$).

Figure 8.

Conditional local fdr estimates $\hat{f}_1(z_1|z_2, z_3)$ for different quantiles of z_2 and z_3 .

Figure 9.
Density histogram of $S_n^{(B)}$.