

SMOOTHED LOG-CONCAVE MAXIMUM LIKELIHOOD ESTIMATION WITH APPLICATIONS

Yining Chen and Richard J. Samworth

University of Cambridge

Abstract: We study the smoothed log-concave maximum likelihood estimator of a probability distribution on \mathbb{R}^d . This is a fully automatic nonparametric density estimator, obtained as a canonical smoothing of the log-concave maximum likelihood estimator. We demonstrate its attractive features both through an analysis of its theoretical properties and a simulation study. Moreover, we use our methodology to develop a new test of log-concavity, and show how the estimator can be used as an intermediate stage of more involved procedures, such as constructing a classifier or estimating a functional of the density. Here again, the use of these procedures can be justified both on theoretical grounds and through its finite sample performance, and we illustrate its use in a breast cancer diagnosis (classification) problem.

Key words and phrases: Classification, functional estimation, log-concave maximum likelihood estimation, smoothing, testing log-concavity.

1. Introduction

Maximum likelihood estimation of shape-constrained densities has received a great deal of interest recently. The allure is the prospect of obtaining fully automatic nonparametric estimators, with no tuning parameters to choose. The general idea dates back to Grenander (1956), who derived the maximum likelihood estimator of a decreasing density on $[0, \infty)$. A characteristic feature of these shape-constrained maximum likelihood estimators is that they are not smooth. For instance, the Grenander estimator has discontinuities at some of the data points. The maximum likelihood estimator of a multi-dimensional log-concave density is the exponential of what Cule, Samworth, and Stewart (2010) call a *tent function*; it may have several ridges. Moreover, in this (and other) examples, the estimator drops discontinuously to zero outside the convex hull of the data.

In some applications, the lack of smoothness may not be a drawback in itself. However, in other circumstances, a smooth estimate might be preferred, because:

- (a) it has a more attractive visual appearance, without ridges or discontinuities that might be difficult to justify to a practitioner;
- (b) it has the potential to offer substantially improved estimation performance, particularly for small sample sizes, where the convex hull of the data is likely to be rather small;

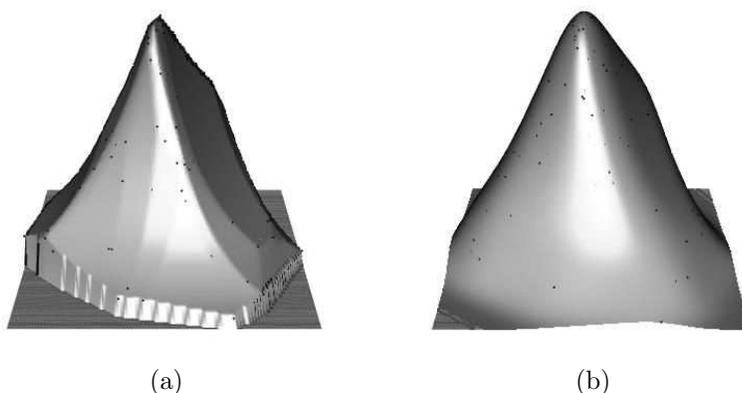


Figure 1. Density estimates based on $n = 200$ observations, plotted as dots, from a standard bivariate normal distribution: (a) log-concave maximum likelihood estimator; (b) smoothed log-concave maximum likelihood estimator.

- (c) for certain applications, e.g., classification, the maximum likelihood estimator being zero outside the convex hull of the data may present problems; see Section 4.1 for further discussion.

For these reasons, we investigate a smoothed version of the d -dimensional log-concave maximum likelihood estimator. The smoothing is achieved by a convolution with a Gaussian density, which preserves the log-concavity shape constraint. To decide how much to smooth, we exploit an interesting property of the log-concave maximum likelihood estimator, that provides a canonical choice of covariance matrix for the Gaussian density, thereby retaining the fully automatic nature of the estimate. The basic idea, introduced by Dümbgen and Rufibach (2009, 2011) for the case $d = 1$ and touched upon in Cule, Samworth, and Stewart (2010), is described in greater detail in Section 2.1.

The challenge of computing the estimator, which involves a d -dimensional convolution integral, is taken up in Section 2.2; see Figure 1 for an illustration of the estimates obtained. The theoretical properties of the smoothed log-concave estimator are studied in Section 2.3. Our framework handles both cases where the log-concavity assumption holds and where it is violated. In Section 2.4, we present new results on the infinite-dimensional projection from a probability distribution on \mathbb{R}^d to its closest log-concave approximation; these give further insight into the misspecified setting. A simulation study follows in Section 2.5, confirming the excellent finite-sample performance.

In Section 3, we introduce a new hypothesis test of log-concavity of multivariate distributions based on our choice of covariance matrix for the Gaussian

density. This test is consistent, easy to implement, and has much improved finite-sample performance compared to existing methods. Section 4 is devoted to applications of the smoothed log-concave maximum likelihood estimator to classification and other functional estimation problems. We provide theory, under both correct and incorrect model specification, for the performance of the resulting procedures in these cases. The classification methodology is applied to the Wisconsin breast cancer data set, where the aim is to aid the diagnosis of future potential breast cancer instances. All proofs are deferred to the Appendix.

Theoretical properties of the unsmoothed log-concave maximum likelihood estimator have been studied in Walther (2002), Pal, Woodroffe, and Meyer (2007), Balabdaoui, Rufibach, and Wellner (2009), and Dümbgen and Rufibach (2009) for the case $d = 1$, and Cule and Samworth (2010), Schuhmacher and Dümbgen (2010), and Dümbgen, Samworth, and Schuhmacher (2011) for the multivariate case. Further properties of log-concave distributions are discussed in Schuhmacher, Hüsler, and Dümbgen (2011), and Walther (2009) provides an overview of the field. Other methods for enforcing various shape constraints have been studied in Braun and Hall (2001), Groeneboom, Jongbloed, and Wellner (2001), Balabdaoui and Wellner (2007), Balabdaoui and Wellner (2010), Pavlides and Wellner (2012), and Carroll, Delaigle, and Hall (2011).

2. The Smoothed Log-concave Maximum Likelihood Estimator

2.1. Definition and basic properties

Let \mathcal{P} denote the set of all probability distributions P on \mathbb{R}^d such that $P(H) < 1$ for all hyperplanes H . In this section, we assume that X_1, X_2, \dots are independent random vectors in \mathbb{R}^d with distribution $P_0 \in \mathcal{P}$. In that case, for sufficiently large n the convex hull of the data, denoted $C_n = \text{conv}(X_1, \dots, X_n)$, is d -dimensional with probability 1. It is then known that there exists a unique log-concave density \hat{f}_n that maximises the likelihood function

$$L(f) = \prod_{i=1}^n f(X_i)$$

over all log-concave densities f . The estimator \hat{f}_n is supported on C_n , and $\log \hat{f}_n$ is piecewise affine on this set. More precisely, there exists an index set J consisting of $(d + 1)$ -tuples $j = (j_0, \dots, j_d)$ of distinct indices in $\{1, \dots, n\}$, such that C_n can be triangulated into simplices $C_{n,j} = \text{conv}(X_{j_0}, \dots, X_{j_d})$ in such a way that

$$\log \hat{f}_n(x) = \begin{cases} b_j^T x - \beta_j & \text{if } x \in C_{n,j}, \\ -\infty & \text{otherwise,} \end{cases}$$

for some vectors $\{b_j : j \in J\}$ in \mathbb{R}^d and real numbers $\{\beta_j : j \in J\}$. Such a function was called a *tent function* in Cule, Samworth, and Stewart (2010) because when $d = 2$ one can think of associating a ‘tent pole’ with each observation, extending vertically out of the plane. For certain tent pole heights, the graph of $\log \hat{f}_n$ is then the roof of a taut tent stretched over the tent poles.

Despite the attractive asymptotic properties of \hat{f}_n derived in the papers cited in the introduction, the simulation results in Cule, Samworth, and Stewart (2010) and Chen (2010) indicate that the finite-sample performance is only strong relative to competitors (e.g., kernel-based methods) for moderate or large sample sizes (say $n > 500$). It appears that for smaller values of n , the convex hull of the data is typically not large enough for good performance.

The idea for fully automatic smoothing of the log-concave maximum likelihood estimator comes from the following observation: Remark 2.3 of Dümbgen, Samworth, and Schuhmacher (2011) (see also Corollary 2.3 of Dümbgen and Rufibach (2009)) shows that while the log-concave maximum likelihood estimator is a good estimator of the first moment of P_0 , it underestimates the covariance matrix. More precisely, we have that

$$\int_{\mathbb{R}^d} x \hat{f}_n(x) dx = \frac{1}{n} \sum_{i=1}^n X_i \equiv \bar{X},$$

say. On the other hand, however,

$$\begin{aligned} \tilde{\Sigma} &\equiv \int_{\mathbb{R}^d} (x - \bar{X})(x - \bar{X})^T \hat{f}_n(x) dx \leq \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})(X_i - \bar{X})^T \\ &< \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})(X_i - \bar{X})^T \equiv \hat{\Sigma}. \end{aligned} \quad (2.1)$$

Here, $A \leq B$ and $A < B$ mean the matrix $B - A$ is non-negative definite and positive definite respectively.

This allows us to define a modified estimator, that we call the *smoothed log-concave maximum likelihood estimator*, denoted \tilde{f}_n . It is given by

$$\tilde{f}_n = \hat{f}_n * \phi_{d, \hat{A}}, \quad (2.2)$$

where $\phi_{d, \hat{A}}$ is the d -variate normal density with zero mean and covariance matrix $\hat{A} = \hat{\Sigma} - \tilde{\Sigma}$. Note that the level of smoothing is automatically determined through the matrix \hat{A} .

The basic properties of \tilde{f}_n are summarised in the proposition below.

Proposition 1. *Let $P_0 \in \mathcal{P}$, and let \tilde{f}_n denote the smoothed log-concave maximum likelihood estimator \tilde{f}_n based on independent observations X_1, \dots, X_n having distribution P_0 . Then*

- (a) \tilde{f}_n is log-concave;
- (b) the support of \tilde{f}_n is \mathbb{R}^d ;
- (c) \tilde{f}_n is a real analytic function on \mathbb{R}^d (in particular, it is infinitely differentiable);
- (d) the mean and covariance matrix corresponding to \tilde{f}_n agree with the sample mean and sample covariance matrix: $\int_{\mathbb{R}^d} x \tilde{f}_n(x) dx = \bar{X}$ and $\int_{\mathbb{R}^d} (x - \bar{X})(x - \bar{X})^T \tilde{f}_n(x) dx = \hat{\Sigma}$.

2.2. Computational issues

The aim of this section is to describe algorithms for computing the smoothed log-concave maximum likelihood estimator \hat{f}_n . As a preliminary step, we need to compute the covariance matrix \hat{A} of the multivariate normal distribution used in the convolution (2.2).

2.2.1. Computation of the covariance matrix \hat{A}

Recall that $\hat{A} = \hat{\Sigma} - \tilde{\Sigma}$, where $\hat{\Sigma}$ is the sample covariance matrix, and

$$\tilde{\Sigma} = \int_{\mathbb{R}^d} xx^T \hat{f}_n(x) dx - \bar{X}\bar{X}^T = \sum_{j \in J} \int_{C_{n,j}} xx^T \exp(b_j^T x - \beta_j) dx - \bar{X}\bar{X}^T. \quad (2.3)$$

We make an affine transformation of each of the regions of integration onto the unit simplex. With $C_{n,j} = \text{conv}(X_{j_0}, \dots, X_{j_d})$, set $D_j = \det[X_{j_1} - X_{j_0}, X_{j_2} - X_{j_0}, \dots, X_{j_d} - X_{j_0}]$, and let $U_d = \{u = (u_1, \dots, u_d) \in [0, \infty)^d : \sum_{l=1}^d u_l \leq 1\}$ be the unit simplex in \mathbb{R}^d . Following Cule and Dürmbgen (2008), we further define the auxiliary functions $J_{\bar{d}} : \mathbb{R}^{\bar{d}+1} \rightarrow \mathbb{R}$ by

$$J_{\bar{d}}(y_0, y_1, \dots, y_{\bar{d}}) = \int_{U_{\bar{d}}} \exp\left(\sum_{l=0}^{\bar{d}} u_l y_l\right) du_1 \dots du_{\bar{d}},$$

where $u_0 = 1 - \sum_{l=1}^{\bar{d}} u_l$. Then, writing $y_{j_l} = \log \hat{f}_n(X_{j_l})$, we have

$$\begin{aligned} & \sum_{j \in J} \int_{C_{n,j}} xx^T \exp(b_j^T x - \beta_j) dx \\ &= \sum_{j \in J} |D_j| \int_{U_d} \left(\sum_{l=0}^d u_l X_{j_l}\right) \left(\sum_{l=0}^d u_l X_{j_l}\right)^T e^{\sum_{l=0}^d u_l y_{j_l}} du \\ &= \sum_{j \in J} |D_j| \left\{ \sum_{l=0}^d \sum_{l'=0}^d X_{j_l} X_{j_{l'}}^T \frac{\partial^2 J_d(y_{j_0}, y_{j_1}, \dots, y_{j_d})}{\partial y_{j_l} \partial y_{j_{l'}}} \right\} \\ &= \sum_{j \in J} |D_j| \left\{ \sum_{l=0}^d \sum_{l'=0}^d X_{j_l} X_{j_{l'}}^T J_{d+2}(y_{j_0}, y_{j_1}, \dots, y_{j_d}, y_{j_l}, y_{j_{l'}}) \right\} \end{aligned}$$

$$+ \sum_{l=0}^d X_{j_l} X_{j_l}^T J_{d+2}(y_{j_0}, y_{j_1}, \dots, y_{j_d}, y_{j_l}, y_{j_l}) \Big\}.$$

We have applied the basic results of Cule and Dümbgen (2008) in the last step. An exact expression for $J_{d+2}(\cdot)$ is given in Appendix B.1 of Cule, Samworth, and Stewart (2010) when its arguments are non-zero and distinct. The Taylor approximation of Cule and Dümbgen (2008) can be used when some of the arguments are small or have similar (or equal) values.

2.2.2. Computation of the smoothed log-concave maximum likelihood estimator

We have

$$\tilde{f}_n(x_0) = \sum_{j \in J} \int_{C_{n,j}} e^{b_j^T x - \beta_j} \frac{1}{(2\pi)^{d/2} (\det \hat{A})^{1/2}} e^{-(x_0 - x)^T \hat{A}^{-1} (x_0 - x)/2} dx.$$

By making an affine transformation of each $C_{n,j}$ onto the unit simplex as in Section 2.2.1, we reduce the problem to integrating the exponential of a quadratic polynomial over the unit simplex. In general, this has no explicit solution, so it has to be evaluated numerically.

Stroud (1971) gives a brief introduction to the problem of evaluating integrals over the unit simplex, while Grundmann and Möller (1978) proposed a combinatorial method. We apply their method, first noting that by integrating out one variable, the dimensionality of the integral can be reduced by one. To see this, consider any $d \times d$ positive definite, symmetric matrix $A \equiv [a_{ll'}]$, any vector $B = (b_1, \dots, b_d)^T \in \mathbb{R}^d$, and any constant $c \in \mathbb{R}$. Writing $\Phi(\cdot)$ for the standard normal distribution function, $u = (u_1, \dots, u_d)^T$, and $u_0 = 1 - \sum_{l=1}^{d-1} u_l$, we have

$$\begin{aligned} & \int_0^1 \int_0^{1-u_1} \dots \int_0^{1-u_1-\dots-u_{d-1}} e^{-u^T A u + B^T u + c} du_d \dots du_2 du_1 \\ &= \int_0^1 \int_0^{1-u_1} \dots \int_0^{1-u_1-\dots-u_{d-1}} e^{-a' u_d^2 + b' u_d + c'} du_d \dots du_2 du_1 \\ &= \int_0^1 \int_0^{1-u_1} \dots \int_0^{1-u_1-\dots-u_{d-2}} e^{c' + b'^2/4a'} \sqrt{\frac{\pi}{a'}} \\ & \quad \times \left\{ \Phi\left(u_0 \sqrt{2a'} - \frac{b'}{\sqrt{2a'}}\right) - \Phi\left(\frac{-b'}{\sqrt{2a'}}\right) \right\} du_{d-1} \dots du_2 du_1. \end{aligned} \quad (2.4)$$

Here, a' , b' and c' are

$$a' = a_{dd}, \quad b' = b_d + 2 \sum_{l=1}^{d-1} a_{dl} u_l \quad \text{and} \quad c' = u_{-d}^T [a_{ll'}]_{1 \leq l, l' \leq d-1} u_{-d} + \sum_{l=1}^{d-1} b_l u_l + c,$$

where $u_{-d} = (u_1, \dots, u_{d-1})^T$. It follows that we can use the combinatorial method to integrate over the $(d-1)$ -dimensional unit simplex. Some special cases are as follows.

- (a) $d = 1$. Here (2.4) is a simple function of $\Phi(\cdot)$ and the smoothed log-concave maximum likelihood estimator can be computed straightforwardly. This method is implemented in the R package `logcondens` (Rufibach and Dümbgen (2006); Dümbgen and Rufibach (2011)).
- (b) $d = 2$. In this case, (2.4) is an integral over $[0, 1]$. Other standard numerical integration methods, such as Gaussian quadrature, can be applied.

The combinatorial method and its variants are implemented in the latest version of the R package `LogConcDEAD` (Cule et al. (2007); Cule, Gramacy, and Samworth (2009)). We found this method to be numerically stable even with several thousand observations, when $\det \hat{A}$ may be rather small (note that in such cases, a' in (2.4) will typically not be close to zero). However, we briefly present two other ways of computing $\tilde{f}_n(x_0)$; while slower in most cases, they do not require the inversion of \hat{A} , so can be used even when $\det \hat{A}$ is small.

(a) **Monte Carlo method.**

- (1) Conditional on X_1, \dots, X_n , generate independent random vectors X_1^*, \dots, X_B^* as $N_d(x_0, \hat{A})$.
- (2) Approximate $\tilde{f}_n(x_0)$ by $\frac{1}{B} \sum_{b=1}^B \hat{f}_n(X_b^*)$.

The validity of this approximation follows from the Strong Law of Large Numbers, applied conditional on X_1, \dots, X_n .

- (b) **Fourier transform.** We can take advantage of the convolution property of the Fourier transform \mathcal{F} as follows. First note that

$$\mathcal{F}(\hat{f}_n)(\xi) = \int_{\mathbb{R}^d} \hat{f}_n(x) e^{-i\xi^T x} dx = \sum_{j \in J} \int_{C_{n,j}} e^{(b_j - i\xi)^T x - \beta_j} dx,$$

which can be evaluated by extending the auxiliary functions J_d to the complex plane. Since $\mathcal{F}(\tilde{f}_n)(\xi) = e^{-ix_0^T \xi - \xi^T \hat{A} \xi / 2} \mathcal{F}(\hat{f}_n)(\xi)$, we can invert $\mathcal{F}(\tilde{f}_n)$ on a fine grid using the fast Fourier transform.

2.2.3. Sampling from the fitted density estimate

Since \tilde{f}_n is the convolution of \hat{f}_n and a multivariate normal density, conditional on X_1, \dots, X_n , it is straightforward to draw an observation X^{**} from \tilde{f}_n , as follows.

- (a) Draw X^* from \hat{f}_n using the algorithm described in Appendix B.3 of Cule, Samworth, and Stewart (2010) or the algorithm of Gopal and Casella (2010).
- (b) Draw $u \sim N_d(0, \hat{A})$, independent of X^* .
- (c) Return $X^{**} = X^* + u$.

2.3. Theoretical performance

It is convenient to define, for $r = 1, 2$, the classes of probability distributions on \mathbb{R}^d given by

$$\mathcal{P}_r = \left\{ P \in \mathcal{P} : \int_{\mathbb{R}^d} \|x\|^r dP(x) < \infty \right\}.$$

The condition $P_0 \in \mathcal{P}_1$ is necessary and sufficient for the existence of a unique upper semi-continuous log-concave density f^* that maximises $\int \log f dP_0$ over all log-concave densities f (Dümbgen, Samworth, and Schuhmacher, 2011, Theorem 2.2). In fact, if P_0 has a density f_0 , and provided that $\int f_0 \log f_0 < \infty$ (which is certainly the case if f_0 is bounded), f^* minimises the Kullback–Leibler divergence $d_{KL}(f, f_0) = \int f_0 \log(f_0/f)$ over all log-concave densities f . In this sense, f^* is the closest log-concave density to P_0 .

The density f^* plays an important role in the following theorem, which describes the asymptotic behaviour of the smoothed log-concave maximum likelihood estimator \tilde{f}_n .

Theorem 1. *Suppose that $P_0 \in \mathcal{P}_2$, and write $\mu = \int_{\mathbb{R}^d} x dP_0(x)$ and $\Sigma = \int_{\mathbb{R}^d} (x - \mu)(x - \mu)^T dP_0(x)$. Let $f^{**} = f^* * N_d(0, A^*)$, where $A^* = \Sigma - \Sigma^*$ with $\Sigma^* = \int_{\mathbb{R}^d} (x - \mu)(x - \mu)^T f^*(x) dx$. Taking $a_0 > 0$ and $b_0 \in \mathbb{R}$ such that $f^{**}(x) \leq e^{-a_0 \|x\| + b_0}$, we have for all $a < a_0$ that*

$$\int_{\mathbb{R}^d} e^{a\|x\|} |\tilde{f}_n(x) - f^{**}(x)| \xrightarrow{a.s.} 0$$

and, if f^{**} is continuous, $\sup_{x \in \mathbb{R}^d} e^{a\|x\|} |\tilde{f}_n(x) - f^{**}(x)| \xrightarrow{a.s.} 0$.

The condition that $P_0 \in \mathcal{P}_2$ imposed in Theorem 1 ensures the finiteness of A^* . We see that in general, \tilde{f}_n converges to a slightly smoothed version of the closest log-concave density to P_0 . However, if P_0 has a log-concave density f_0 , then $f_0 = f^* = f^{**}$, so \tilde{f}_n is strongly consistent in these exponentially weighted total variation and supremum norms. In fact, suppose that $a : \mathbb{R}^d \rightarrow \mathbb{R}$ is a sublinear function, i.e. $a(x+y) \leq a(x) + a(y)$ and $a(rx) = ra(x)$ for all $x, y \in \mathbb{R}^d$ and $r \geq 0$, satisfying $e^{a(x)} f(x) \rightarrow 0$ as $\|x\| \rightarrow \infty$. It can be shown that under the conditions of Theorem 1,

$$\int_{\mathbb{R}^d} e^{a(x)} |\tilde{f}_n(x) - f^{**}(x)| \xrightarrow{a.s.} 0$$

(Schuhmacher, Hüsler, and Dümbgen (2011)).

Despite being smooth and having full support, it turns out that \tilde{f}_n is rather close to \hat{f}_n . This is quantified in the finite-sample bound below.

Proposition 2. *If $x \in C_{n,j}$, and $\hat{f}_n(x) = \exp(b_j^T x - \beta_j)$, then*

$$\frac{\tilde{f}_n(x) - \hat{f}_n(x)}{\hat{f}_n(x)} \leq e^{(1/2)b_j^T \hat{A}b_j} - 1.$$

Moreover,
$$\int_{\mathbb{R}^d} |\tilde{f}_n - \hat{f}_n| \leq 2(e^{(1/2)\lambda_{\max}} - 1 + \delta_n)$$

where $\lambda_{\max} = \max_{j \in J} b_j^T \hat{A}b_j$, and $\delta_n = \int_{C_n^c} \tilde{f}_n$.

2.4. Properties of (smoothed) log-concave approximations

In this subsection, we give new insights into the maps from a probability distribution P to its log-concave approximation f^* , and its smoothed version f^{**} . Results such as these enhance our understanding of the behaviour of maximum likelihood estimators in non-convex, misspecified models, where existing results are limited. Theorem 2 below shows that log-concave approximations and their smoothed analogues preserve independence of components. As well as being of use in our simulation studies, this is the key result that underpins a new approach to fitting independent component analysis models using nonparametric maximum likelihood (Samworth and Yuan (2012)).

Theorem 2. *Suppose that $P \in \mathcal{P}_1$ is the product measure $P = P_1 \otimes P_2$ on \mathbb{R}^d , where P_1 and P_2 are probability measures on \mathbb{R}^{d_1} and \mathbb{R}^{d_2} respectively, with $d_2 = d - d_1$. Let f^* be the log-concave approximation to P , and let f_ℓ^* be the log-concave approximation to P_ℓ , for $\ell = 1, 2$. Then writing $x = (x_1^T, x_2^T)^T$, where $x_1 \in \mathbb{R}^{d_1}$ and $x_2 \in \mathbb{R}^{d_2}$, we have*

$$f^*(x) = f_1^*(x_1)f_2^*(x_2).$$

*If, further, $P \in \mathcal{P}_2$, f^{**} is the smoothed log-concave approximation to P , and f_ℓ^{**} is the smoothed log-concave approximation to P_ℓ for $\ell = 1, 2$, then for all $x = (x_1^T, x_2^T)^T$,*

$$f^{**}(x) = f_1^{**}(x_1)f_2^{**}(x_2).$$

Our next theorem characterises the log-concavity constraint through the trace of the non-negative definite matrix A^* of Theorem 1.

Theorem 3. *Suppose that $P \in \mathcal{P}_1$. Then $\text{tr}(A^*) = 0$ if and only if P has a log-concave density.*

The ‘if’ part of this statement is well-known, but the ‘only if’ part is new. The two parts together motivate our testing procedure for log-concavity, that is developed in Section 3.

In most cases, it is difficult to explicitly find the log-concave approximation f^* to a given distribution $P \in \mathcal{P}_1$. The final result here is straightforward to prove, but is of interest because it shows that some log-concave densities can have a large ‘domain of attraction’.

Proposition 3. *Let f^* be an upper semi-continuous, log-concave density on \mathbb{R}^d . Then the class of distributions $P \in \mathcal{P}_1$ with log-concave approximation f^* is convex.*

For instance, if $f(x; \alpha, \sigma) = \alpha\sigma^\alpha/[2(|x| + \sigma)^{\alpha+1}]$ is a symmetrised Pareto density with $\alpha > 1$ and $\sigma > 0$, then it can be shown that its log-concave projection is $f^*(x; \alpha, \sigma) = [(\alpha - 1)/2\sigma] \exp\{-(\alpha - 1)|x|/\sigma\}$. Thus the class of distributions whose log-concave projection is the standard Laplace density is infinite-dimensional.

2.5. Finite sample performance

Our simulation study considered the normal location mixture density $f(\cdot) = 0.4\phi_d(\cdot) + 0.6\phi_d(\cdot - \mu)$ for $\|\mu\| = 1, 2$ and 3 , where $\phi_d = \phi_{d,I}$. This mixture density is log-concave if and only if $\|\mu\| \leq 2$. For each density, for $d = 2$ and $d = 3$ and sample sizes $n = 100$ and $n = 1,000$, we computed the Integrated Squared Error (ISE) of the smoothed log-concave maximum likelihood estimator for each of 50 replications. We also computed the ISE of the log-concave maximum likelihood estimator and that of a kernel density estimator with a Gaussian kernel and the optimal ISE bandwidth for each individual data set, which is unknown in practice. The boxplots of the ISEs for the different methods are given in Figure 2 for $d = 3$. The analogous plots for the case $d = 2$ can be found in Chen and Samworth (2011).

We see that when the true density is log-concave, the smoothed log-concave estimator offers substantial ISE improvements over its unsmoothed analogue for both sample sizes, particularly at the smaller sample size $n = 100$. It also outperforms by a considerable margin the kernel density estimator with the optimal ISE bandwidth. When the log-concavity assumption is violated, the smoothed log-concave estimator is still competitive with the optimal-ISE kernel estimator at the smaller sample size $n = 100$, and also improves on its unsmoothed analogue. However, when $n = 1,000$, the bias caused by the fact that $\int_{\mathbb{R}^d} (f^* - f)^2 > 0$ dominates the contribution from the variance of the estimator, and the kernel estimator is an improvement. These results suggest that the smoothed log-concave estimator has excellent performance when the true density is log-concave, and remains competitive in situations where the log-concavity assumption is violated,

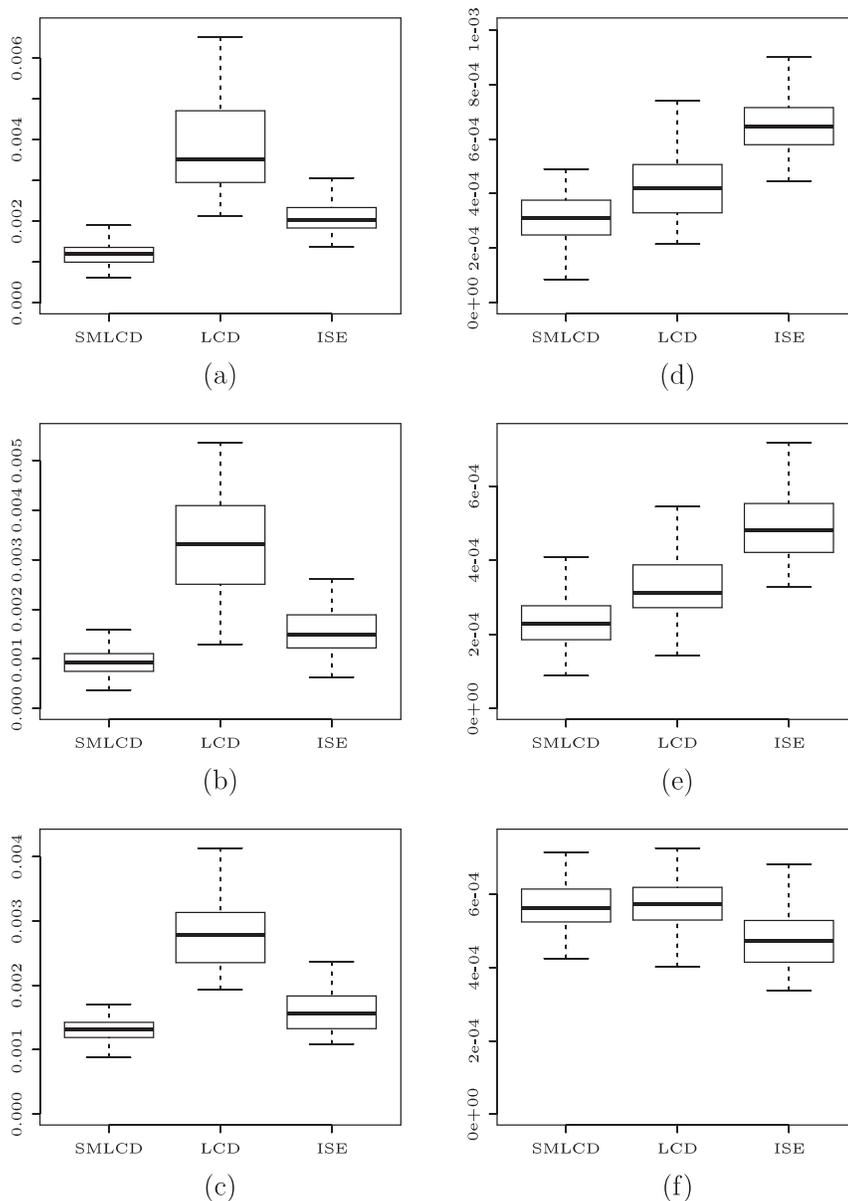


Figure 2. Boxplots of ISEs for $d = 3$ with the Gaussian location mixture true density for the smoothed log-concave maximum likelihood estimator SMLCD, log-concave maximum likelihood estimator LCD and kernel density estimator with the 'oracle' optimal ISE bandwidth: (a) $n = 100$, $\|\mu\| = 1$; (b) $n = 100$, $\|\mu\| = 2$; (c) $n = 100$, $\|\mu\| = 3$; (d) $n = 1,000$, $\|\mu\| = 1$; (e) $n = 1,000$, $\|\mu\| = 2$; (f) $n = 1,000$, $\|\mu\| = 3$.

provided that the modelling bias caused by this misspecification is not too large relative to the sampling variability of the estimator.

3. A New Test of Log-concavity

Several tests of log-concavity have been proposed in the literature. An (1995) and Walther (2002) discuss various tests for univariate data, while Cule, Samworth, and Stewart (2010) presented two tests of log-concavity for multivariate data. Hazelton (2011) proposed another multivariate test based on kernel density estimates that had improved finite-sample performance on his simulated examples. However, none of these multivariate tests has theoretical support.

Suppose $X_1, \dots, X_n \stackrel{iid}{\sim} P_0 \in \mathcal{P}_1$, and we seek a size $\alpha \in (0, 1)$ test of $H_0 : P_0$ has a log-concave density against $H_1 : P_0$ does not have a log-concave density. Motivated by Theorem 3, we propose the following.

- (a) Compute the log-concave maximum likelihood density estimate \hat{f}_n .
- (b) Compute the test statistic $\text{tr}(\hat{A})$, where $\hat{A} = \hat{\Sigma} - \tilde{\Sigma}$, as in (2.1).
- (c) Generate a reference distribution as follows: for $b = 1, \dots, B$, draw conditionally independent samples $X_{1b}^*, \dots, X_{nb}^*$ from \hat{f}_n . For each bootstrap sample, compute the log-concave maximum likelihood estimator \hat{f}_{nb} , then compute $\text{tr}(\hat{A}_{nb})$, where

$$\begin{aligned} \hat{A}_{nb} &\equiv \hat{\Sigma}_b - \tilde{\Sigma}_b \\ &\equiv \frac{1}{n-1} \sum_{i=1}^n (X_{ib}^* - \bar{X}_b^*)(X_{ib}^* - \bar{X}_b^*)^T - \int_{\mathbb{R}^d} (x - \bar{X}_b^*)(x - \bar{X}_b^*)^T \hat{f}_{nb}(x) dx, \end{aligned}$$

$$\text{and } \bar{X}_b^* = n^{-1} \sum_{i=1}^n X_{ib}^*.$$

- (d) Reject H_0 if $(B+1)^{-1} \sum_{b=1}^{B+1} \mathbb{1}_{\{\text{tr}(\hat{A}) > \text{tr}(\hat{A}_{nb})\}} > 1 - \alpha$.

We call this procedure a *trace* test. Its consistency is justified as follows.

Theorem 4. *Suppose that $P_0 \in \mathcal{P}_1$. If P_0 is not log-concave, then for each $B \in \mathbb{N}$ the power of the test converges to one as $n \rightarrow \infty$.*

We remark that if $P_0 \in \mathcal{P}_2$, one can also draw bootstrap samples from \tilde{f}_n instead of \hat{f}_n in Step (c). To illustrate the performance of the test, we ran two small simulation studies. In the first, we simulated from the bivariate mixture $f(x) = (1/2)\phi_{2,I}(x) + (1/2)\phi_{2,I}(x - \mu)$, with $\|\mu\| = 0, 2, 4$ (which we recall is log-concave if and only if $\|\mu\| \leq 2$). For each simulation setup, we performed 200 hypothesis tests with $B = 99$. The proportion of times that the null hypothesis was rejected in a size $\alpha = 0.05$ test is reported in Table 1. For comparison, we also report the results from the critical bandwidth test proposed by Hazelton (2011).

Table 1. Proportion of times out of 200 repetitions that the null hypothesis was rejected with $\alpha = 0.05$.

n	Method	$\ \mu\ = 0$	$\ \mu\ = 2$	$\ \mu\ = 4$
200	critical bandwidth	0.065	0.015	0.985
	trace	0.045	0.045	1.000
500	critical bandwidth	0.045	0.005	1.000
	trace	0.045	0.055	1.000

Table 2. Proportion of times out of 200 repetitions that the null hypothesis was rejected with $\alpha = 0.05$.

n	Method	Cases			
		(a)	(b)	(c)	(d)
200	critical bandwidth	0.520	0.195	0.395	0.295
	trace	1.000	0.960	1.000	1.000
500	critical bandwidth	0.760	0.340	0.710	0.505
	trace	1.000	1.000	1.000	1.000

The permutation test studied by Cule, Samworth, and Stewart (2010) did not perform as well as the critical bandwidth test, so we omitted its results here. For the second study, we replicated the settings considered in Hazelton (2011), where four different types of bivariate densities of independent components were chosen. The marginal distributions were:

- (a) A $N(0, 1/4)/2 + N(0, 4)/2$ distribution and a $N(0, 1/4)/2 + N(2, 4)/2$ distribution;
- (b) A t_4 distribution in both cases;
- (c) A $N(0, 1/4)/2 + N(2, 4)/2$ and a t_4 distribution;
- (d) A $N(0, 1/4)/2 + N(2, 5)/2$ density, and a $\Gamma(2, 1)$ distribution.

Note that all of these densities are unimodal but not log-concave. The corresponding estimates of the power of the tests are presented in Table 2. The first study shows that the trace test controls the Type I error satisfactorily (and appears to be less conservative than the critical bandwidth test when $\|\mu\| = 2$). The results of the second study, though, are quite striking, and suggest that our test for log-concavity has considerably improved finite-sample power compared to the critical bandwidth test. Hazelton (2011) noted that the critical bandwidth test can have reduced power due to the boundary bias of the kernel estimators and is quite sensitive to the outliers (in fact, one also needs to pick a compact region containing the majority of the data, and this choice is somewhat arbitrary). Our test avoids these issues and performs well even in the presence of outliers or when the true density has bounded support.

4. Other Applications

4.1. Classification problems

Changing notation slightly from the previous section, we assume that (X, Y) , $(X_1, Y_1), \dots, (X_n, Y_n)$ are independent and identically distributed pairs taking values in $\mathbb{R}^d \times \{1, \dots, K\}$. Let $\mathbb{P}(Y = k) = \pi_k$ for $k = 1, \dots, K$, and suppose that conditional on $Y = k$, the random vector X has distribution P_k .

A *classifier* is a measurable function $C : \mathbb{R}^d \rightarrow \{1, \dots, K\}$, with the interpretation that the classifier assigns the point $x \in \mathbb{R}^d$ to class $C(x)$. The *misclassification error rate*, or *risk*, of C is

$$\text{Risk}(C) = \mathbb{P}\{C(X) \neq Y\}.$$

In the case where each distribution P_k has a density f_k , the classifier that minimises the risk is the *Bayes classifier*

$$C^{\text{Bayes}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} \pi_k f_k(x).$$

(For all classifiers defined by an argmax as above, for the sake of definiteness we split ties by taking the smallest element of the argmax .) We are interested in the *log-concave Bayes classifier* and *smoothed log-concave Bayes classifier*, defined respectively by

$$C^{\text{LCBayes}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} \pi_k f_k^*(x) \quad \text{and} \quad C^{\text{SLCBayes}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} \pi_k f_k^{**}(x).$$

Here, f_k^* and f_k^{**} are the log-concave approximation to P_k and its smoothed analogue, defined in Theorem 1. In particular, both classifiers coincide with the Bayes classifier when $\{P_k : k = 1, \dots, K\}$ have log-concave densities. Empirical analogues of these theoretical classifiers are

$$\hat{C}_n^{\text{LC}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} N_k \hat{f}_{n,k}(x) \quad \text{and} \quad \hat{C}_n^{\text{SLC}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} N_k \tilde{f}_{n,k}(x).$$

Here, $N_k = \sum_{i=1}^n \mathbb{1}_{\{Y_i=k\}}$ is the number of observations from the k th class, and $\hat{f}_{n,k}$ and $\tilde{f}_{n,k}$ are, respectively, the log-concave maximum likelihood estimator of f_k and its smoothed analogue, based on $\{X_i : Y_i = k\}$.

The next theorem describes the asymptotic behaviour of these classifiers. It reveals that the risks of \hat{C}_n^{LC} and \hat{C}_n^{SLC} converge not (in general) to the Bayes risk, but instead to the risks of C^{LCBayes} and C^{SLCBayes} , respectively. This is a similar situation to that encountered when a parametric classifier, such as linear or quadratic discriminant analysis, is used but the relevant parametric modelling assumptions fail to hold. It suggests that the classifiers \hat{C}_n^{LC} and \hat{C}_n^{SLC} should only be used when the hypothesis of log-concavity can be expected to hold, at least approximately.

Theorem 5. (a) Assume $P_k \in \mathcal{P}_1$ for $k = 1, \dots, K$. Let $\mathcal{X}^* = \{x \in \mathbb{R}^d : |\operatorname{argmax}_k \pi_k f_k^*(x)| = 1\}$. Then $\hat{C}_n^{\text{LC}}(x) \xrightarrow{a.s.} C^{\text{LCBayes}}(x)$ for almost all $x \in \mathcal{X}^*$, and

$$\text{Risk}(\hat{C}_n^{\text{LC}}) \rightarrow \text{Risk}(C^{\text{LCBayes}}).$$

(b) If $P_k \in \mathcal{P}_2$ for $k = 1, \dots, K$, and $\mathcal{X}^{**} = \{x \in \mathbb{R}^d : |\operatorname{argmax}_k \pi_k f_k^{**}(x)| = 1\}$, then $\hat{C}_n^{\text{SLC}}(x) \xrightarrow{a.s.} C^{\text{SLCBayes}}(x)$ for almost all $x \in \mathcal{X}^{**}$, and

$$\text{Risk}(\hat{C}_n^{\text{SLC}}) \rightarrow \text{Risk}(C^{\text{SLCBayes}}).$$

In fact, the smoothed log-concave classifier is somewhat easier to apply in practical classification problems than its unsmoothed analogue. This is because if $x_0 \in \mathbb{R}^d$ is outside the convex hull of the training data for each of the K classes (an event of positive probability), then the log-concave maximum likelihood estimates of the densities at x_0 are all zero. Thus all such points would be assigned by \hat{C}_n^{LC} to Class 1. On the other hand, \hat{C}_n^{SLC} avoids this problem altogether. For these reasons, we considered only \hat{C}_n^{SLC} in our simulation study (Chen and Samworth (2011)) and below.

We remark that the direct use of \hat{C}_n^{SLC} (or any other classifier based on non-parametric density estimation) is not recommended when $d > 4$, due to the curse of dimensionality. In such circumstances there are two options: dimension reduction (cf. Section 4.2), or such modelling assumptions as independent component analysis models (Samworth and Yuan (2012)). In either case, the methodology we develop remains applicable, but now as part of a more involved procedure.

4.2. Breast cancer example

In the Wisconsin breast cancer data set (Street, Wolberg, and Mangasarian (1993)), 30 measurements were taken from a digitised image of a fine needle aspirate of different breast masses. There were 357 benign and 212 malignant instances, and we aim to construct a classifier based on this training data set to aid future diagnoses. Only the first two principal components of the training data were considered, and these capture 63% of the total variability; cf. Figure 3(a). This was done to make our procedure computationally feasible, to reduce the effect of the curse of dimensionality, and to facilitate plots such as Figure 3 below.

In Figure 3(b), we show the smoothed log-concave density estimates of both the benign and malignant classes. Figure 3(c) plots the decision boundaries of the smoothed log-concave classifier, where we treat benign cases and malignant cases equally. However in practice, misdiagnosing a malignant tumour as benign is much more serious than misidentifying a benign one as malignant. One may therefore seek to incorporate different losses into the classifier. For $k = 1, 2$, let L_k denote the cost of failure to recognise the class k (this notion can easily be

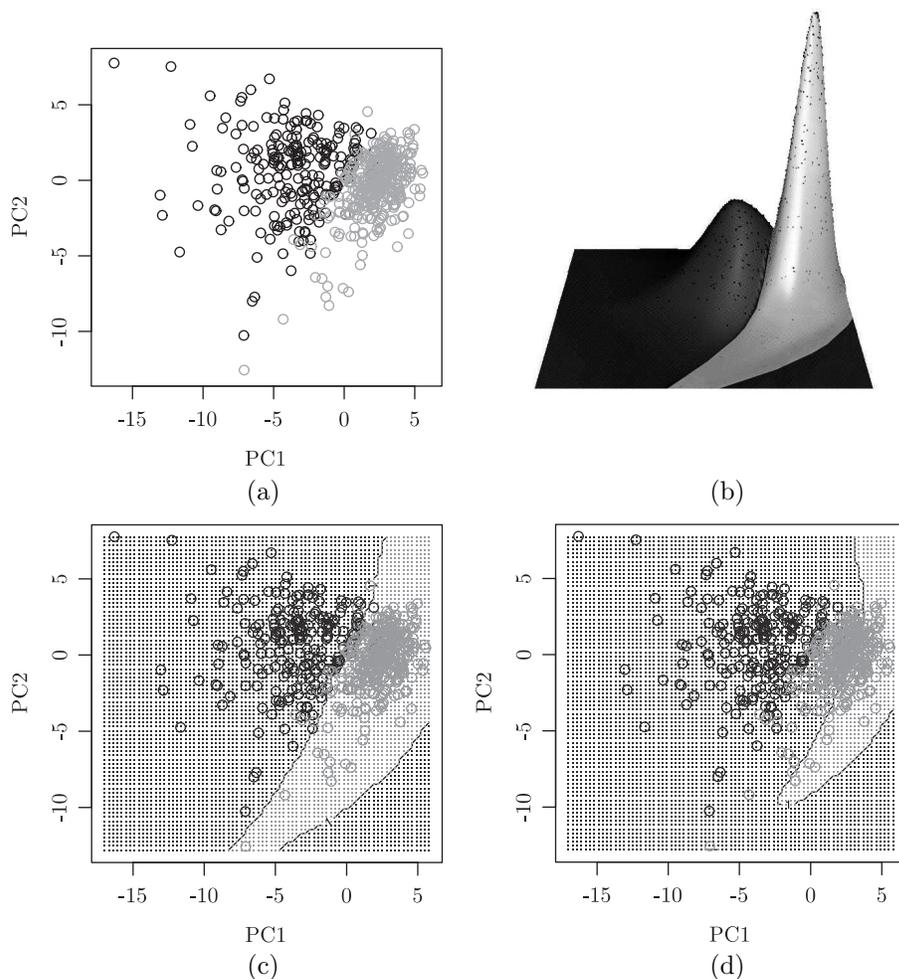


Figure 3. (a) Wisconsin breast cancer data (benign cases are light; malignant cases are dark); (b) smoothed log-concave maximum likelihood density estimates; (c) and (d) plot the decision boundaries of the smoothed log-concave classifier, where the loss $L_2 = 1$ and $L_2 = 100$, respectively.

generalised to multicategory situations with $L_{kk'}$ the loss incurred in assigning the pair (X, Y) to class k' when $Y = k$. Redefining the risk as

$$\text{Risk}(C) = L_1 \mathbb{P}\{C(X) = 2 \cap Y = 1\} + L_2 \mathbb{P}\{C(X) = 1 \cap Y = 2\},$$

the same asymptotic properties continue to hold, mutatis mutandis, for the classifier

$$\hat{C}_n^{\text{SLC}^*}(x) = \operatorname{argmax}_{k \in \{1,2\}} N_k L_k \tilde{f}_{n,k}(x).$$

We observe that this modification requires no recalculation of the smoothed log-concave density estimates and there is no loss of generality in taking $L_1 = 1$. A GUI with slider is implemented in the R package `LogConcDEAD`; it provides a way of demonstrating how the decision boundaries change as L_2 varies. For the purpose of illustration, Figure 3(d) plots the decision boundaries of $\hat{C}_n^{\text{SLC}*}$ when the cost L_2 of misidentifying a malignant tumour is 100. Compared with Figure 3(c), observations are considerably more likely to be classified as malignant under this setting.

4.3. Functional estimation problems

Classification problems are an important example of a situation in which one is more interested in a functional of one or more density estimates than in the density estimate itself. For simplicity, we return in this section to the situation in which we have a single independent sample X_1, \dots, X_n distributed according to a distribution P_0 .

In general, we can estimate a functional $\theta_0 = \theta(P_0)$ using the plug-in smoothed log-concave estimate $\tilde{\theta}_n = \theta(\tilde{P}_n)$, where \tilde{P}_n is the distribution with density \tilde{f}_n . Note that even if this functional cannot be computed directly, it is usually straightforward to construct a Monte Carlo approximation to $\tilde{\theta}_n$ by applying the algorithm for sampling from \tilde{f}_n outlined in Section 2.2.3. To describe the theoretical properties of these functional estimates, for $a > 0$, let \mathcal{B}_a denote the set of signed measures P on \mathbb{R}^d with $\int_{\mathbb{R}^d} e^{a\|x\|} d|P|(x) < \infty$. Equip \mathcal{B}_a with the norm

$$\|P\|_a = \int_{\mathbb{R}^d} e^{a\|x\|} d|P|(x).$$

We can then consider θ as a measurable function on $(\mathcal{B}_a, \|\cdot\|_a)$ taking values in some other normed space $(\mathcal{B}, \|\cdot\|)$.

Proposition 4. *Let $P_0 \in \mathcal{P}_2$, and let P_0^{**} denote the probability distribution whose density is the smoothed version of the log-concave approximation to P_0 . Suppose that $\theta : \mathcal{B}_a \rightarrow \mathcal{B}$ is continuous and let $\theta^{**} = \theta(P_0^{**})$. Then $\|\tilde{\theta}_n - \theta^{**}\| \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$.*

Once again, we remark that if P_0 has a log-concave density, then $P_0 = P_0^{**}$. The fact that the topology on \mathcal{B}_a is rather strong means that the continuity requirement on θ is relatively weak. This is illustrated in the following consideration of *linear* functionals.

Corollary 1. *Let $P_0 \in \mathcal{P}_2$, and let $a_0 > 0$ and $b_0 \in \mathbb{R}$ be such that $f^{**}(x) \leq e^{-a_0\|x\|+b_0}$, where f^{**} is the smoothed log-concave approximation to P_0 . Let $\theta(P) = \int_{\mathbb{R}^d} g dP$ for some measurable function $g : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying*

$$\sup_{x \in \mathbb{R}^d} e^{-a\|x\|} |g(x)| < \infty \tag{4.1}$$

*for some $a < a_0$. Then $\tilde{\theta}_n \xrightarrow{a.s.} \theta^{**}$.*

Acknowledgement

We thank an associate editor and two referees for their helpful comments. The second author is grateful for the support of a Leverhulme Research Fellowship and an EPSRC Early Career Fellowship.

Appendix

Proof of Proposition 1.

- (a) This follows immediately from Theorems 2.8 and 2.18 of Dharmadhikari and Joag-Dev (1988).
- (b) Note that for any non-empty open set $B \subseteq \mathbb{R}^d$,

$$\int_B \tilde{f}_n(x) dx = \int_B \int_{C_n} \hat{f}_n(y) \phi_{d, \hat{A}}(x-y) dy dx,$$

which is positive, since the integrand is positive and continuous on the region of integration.

- (c) The fact that \tilde{f}_n is infinitely differentiable follows from Proposition 8.10 of Folland (1999). In fact, using standard multi-index notation with $\alpha = (\alpha_1, \dots, \alpha_d)$ and $\partial^\alpha = \left(\frac{\partial}{\partial x_1}\right)^{\alpha_1} \dots \left(\frac{\partial}{\partial x_d}\right)^{\alpha_d}$, we have $\partial^\alpha \tilde{f}_n = \hat{f}_n * \partial^\alpha \phi_{d, \hat{A}}$. Writing $|\alpha| = \sum_{l=1}^d \alpha_l$ and $\alpha! = \prod_{l=1}^d \alpha_l!$, it follows that for any $x_0 \in \mathbb{R}^d$ and $k \in \mathbb{N}$,

$$\begin{aligned} & \left| \tilde{f}_n(x) - \sum_{|\alpha| \leq k} \frac{(\partial^\alpha \tilde{f}_n)(x_0)}{\alpha!} (x-x_0)^\alpha \right| \\ & \leq \int_{C_n} \hat{f}_n(y) \left| \phi_{d, \hat{A}}(x-y) - \sum_{|\alpha| \leq k} \frac{(\partial^\alpha \phi_{d, \hat{A}})(x_0)}{\alpha!} (x-y-x_0)^\alpha \right| dy \rightarrow 0 \end{aligned}$$

as $k \rightarrow \infty$, by the Dominated Convergence Theorem and Lemma 1 of Cule and Samworth (2010).

- (d) Conditional on X_1, \dots, X_n , let X^* and Y^* be independent, with X^* having density \hat{f}_n and Y^* having density $\phi_{d, \hat{A}}$, so that $X^* + Y^*$ has conditional density \tilde{f}_n . Then

$$\begin{aligned} \mathbb{E}(X^* + Y^* | X_1, \dots, X_n) &= \mathbb{E}(X^* | X_1, \dots, X_n) = \int_{\mathbb{R}^d} x \hat{f}_n(x) dx = \bar{X}, \\ \text{Cov}(X^* + Y^* | X_1, \dots, X_n) &= \text{Cov}(X^* | X_1, \dots, X_n) + \text{Cov}(Y^* | X_1, \dots, X_n) \\ &= \tilde{\Sigma} + \hat{A} = \hat{\Sigma}. \end{aligned}$$

Proof of Theorem 1. Let d_P and d_{TV} denote the Prohorov and total variation metrics on the space of probability measures on \mathbb{R}^d . Recall that d_P metrises weak convergence, and that $d_P \leq d_{TV}$. Let $\hat{\mu}_n$ denote the probability measure corresponding to the density \hat{f}_n , let $\hat{\nu}_n$ denote the probability measure corresponding to the convolution of f^* with the measure $N_d(0, \hat{A})$, and let ν denote the probability measure corresponding to f^{**} . Then

$$\begin{aligned} d_P(\hat{\mu}_n, \nu) &\leq d_P(\hat{\mu}_n, \hat{\nu}_n) + d_P(\hat{\nu}_n, \nu) \\ &\leq d_{TV}(\hat{\mu}_n, \hat{\nu}_n) + d_P(\hat{\nu}_n, \nu) \\ &= \frac{1}{2} \int_{\mathbb{R}^d} |\hat{f}_n * N_d(0, \hat{A}) - f^* * N_d(0, \hat{A})| + d_P(\hat{\nu}_n, \nu) \\ &\leq \frac{1}{2} \int_{\mathbb{R}^d} |\hat{f}_n - f^*| + d_P(\hat{\nu}_n, \nu). \end{aligned} \quad (\text{A.1})$$

The first term of (A.1) converges almost surely to zero, by Theorem 2.15 of Dümbgen, Samworth, and Schuhmacher (2011). The second term also converges almost surely to zero, using the fact that $\hat{A} \xrightarrow{a.s.} A^*$ as $n \rightarrow \infty$. Proposition 2 of Cule and Samworth (2010) strengthens the mode of convergence and yields the result.

Proof of Proposition 2. If $x \in C_{n,j}$, and $\hat{f}_n(x) = \exp(b_j^T x - \beta_j)$, then $\hat{f}_n(x - y) \leq \exp\{b_j^T(x - y) - \beta_j\}$ for all $y \in \mathbb{R}^d$. It follows that

$$\frac{\tilde{f}_n(x) - \hat{f}_n(x)}{\hat{f}_n(x)} \leq \int_{\mathbb{R}^d} e^{-b_j^T y} \phi_{d, \hat{A}}(y) dy - 1 = e^{b_j^T \hat{A} b_j / 2} - 1. \quad (\text{A.2})$$

Now

$$\begin{aligned} \int_{\mathbb{R}^d} |\tilde{f}_n - \hat{f}_n| &= \int_{C_n} |\tilde{f}_n - \hat{f}_n| + \delta_n \\ &= \int_{C_n} (\tilde{f}_n - \hat{f}_n)_+ + \int_{C_n} (\hat{f}_n - \tilde{f}_n)_+ + \delta_n. \end{aligned}$$

But

$$\int_{C_n} (\hat{f}_n - \tilde{f}_n)_+ = \int_{C_n} (\tilde{f}_n - \hat{f}_n)_+ - \int_{C_n} (\tilde{f}_n - \hat{f}_n) = \int_{C_n} (\tilde{f}_n - \hat{f}_n)_+ + \delta_n.$$

It therefore follows from this and (A.2) that

$$\int_{\mathbb{R}^d} |\tilde{f}_n - \hat{f}_n| \leq 2 \sum_{j \in J} \int_{C_{n,j}} \hat{f}_n(x) (e^{b_j^T \hat{A} b_j / 2} - 1) dx + 2\delta_n \leq 2(e^{\lambda_{\max} / 2} - 1 + \delta_n),$$

as required.

Proof of Theorem 2. (a) Let f be an arbitrary log-concave density on \mathbb{R}^d , and let X be a random vector with density f . Letting $X = (X_1^T, X_2^T)^T$, where X_1

and X_2 take values in \mathbb{R}^{d_1} and \mathbb{R}^{d_2} , respectively, we write f_{X_1} for the marginal density of X_1 and $f_{X_2|X_1}(\cdot|x_1)$ for the conditional density of X_2 given $X_1 = x_1$. By Theorem 6 of Prékopa (1973), f_{X_1} is log-concave and by Proposition 1 of Cule, Samworth, and Stewart (2010), $f_{X_2|X_1}(\cdot|x_1)$ is log-concave for each x_1 .

There is also no loss of generality in assuming f upper semi-continuous. Since $P \in \mathcal{P}_1$, we can assume without loss of generality that $\int_{\mathbb{R}^d} |\log f| dP < \infty$. We can then apply Fubini's Theorem and seek to maximise over all upper semi-continuous log-concave densities the quantity

$$\begin{aligned} \int_{\mathbb{R}^d} \log f dP &= \int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} \{\log f_{X_1}(x_1) + \log f_{X_2|X_1}(x_2|x_1)\} dP_2(x_2) dP_1(x_1) \\ &= \int_{\mathbb{R}^{d_1}} \log f_{X_1}(x_1) P_1(dx_1) \\ &\quad + \int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} \log f_{X_2|X_1}(x_2|x_1) dP_2(x_2) dP_1(x_1). \end{aligned} \tag{A.3}$$

The first term on the right-hand side of (A.3) is maximised uniquely over all upper semi-continuous log-concave densities by setting $f_{X_1} = f_1^*$. Moreover, for any fixed x_1 , the quantity $\int_{\mathbb{R}^{d_2}} \log f_{X_2|X_1}(x_2|x_1) dP_2(x_2)$ is maximised uniquely over upper semi-continuous log-concave densities by setting $f_{X_2|X_1}(\cdot|x_1) = f_2^*$. Since this choice does not depend on x_1 , it maximises the second term on the right-hand side of (A.3). Because both terms can be maximised simultaneously, it follows that $f^* = f_1^* f_2^*$, as desired.

(b) Write Σ and Σ^* for the covariance matrices corresponding to the probability distribution P and the density f^* , respectively. The independence structure of P_0 and f^* gives that $\Sigma = \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix}$ and $\Sigma^* = \begin{bmatrix} \Sigma_1^* & 0 \\ 0 & \Sigma_2^* \end{bmatrix}$. Here, Σ_1 and Σ_1^* are $d_1 \times d_1$ submatrices, while Σ_2 and Σ_2^* are $d_2 \times d_2$ submatrices. Therefore, $A^* = \Sigma - \Sigma^*$ is of the form $A^* = \begin{bmatrix} A_1^* & 0 \\ 0 & A_2^* \end{bmatrix}$. Writing $x, y \in \mathbb{R}^d$ as $(x_1^T, x_2^T)^T$ and $(y_1^T, y_2^T)^T$, respectively, where $x_1, y_1 \in \mathbb{R}^{d_1}$ and $x_2, y_2 \in \mathbb{R}^{d_2}$, it follows again by Fubini's Theorem that

$$\begin{aligned} f^{**}(x) &= (f^* * N_d(0, A^*))(x) \\ &= \int_{\mathbb{R}^{d_1}} \int_{\mathbb{R}^{d_2}} f_1^*(y_1) f_2^*(y_2) dN_{d_2}(0, A_2^*)(x_2 - y_2) dN_{d_1}(0, A_1^*)(x_1 - y_1) \\ &= \left\{ \int_{\mathbb{R}^{d_1}} f_1^*(y_1) dN_{d_1}(0, A_1^*)(x_1 - y_1) \right\} \left\{ \int_{\mathbb{R}^{d_2}} f_2^*(y_2) dN_{d_2}(0, A_2^*)(x_2 - y_2) \right\} \\ &= f_1^{**}(x_1) f_2^{**}(x_2). \end{aligned}$$

Proof of Theorem 3. Let $P \in \mathcal{P}_1$, and let f^* denote its log-concave approximation. Without loss of generality, we assume $\int_{\mathbb{R}^d} x dP(x) = 0$, so it suffices to

show that if $A^* := \int_{\mathbb{R}^d} xx^T dP(x) - \int_{\mathbb{R}^d} xx^T f^*(x) dx$ is the zero matrix, then P has a log-concave density.

Let P^* denote the distribution corresponding to f^* , let $X \sim P$, and let $X^* \sim P^*$. For an arbitrary $u \in \mathbb{R}^d$, let F_u and F_u^* denote the distribution functions of $u^T X$ and $u^T X^*$, respectively, and let

$$G_u(s) = \int_{-\infty}^s F_u(t) dt \quad \text{and} \quad G_u^*(s) = \int_{-\infty}^s F_u^*(t) dt.$$

Fix $s \in \mathbb{R}$. By applying Remark 2.3 of Dümbgen, Samworth, and Schuhmacher (2011) to the convex function $x \mapsto (s - u^T x)_+$ and Fubini's Theorem, we have that

$$\begin{aligned} 0 &\leq \int_{\mathbb{R}^d} (s - u^T x)_+ d(P - P^*)(x) = \int_{\mathbb{R}^d} \int_{-\infty}^{\infty} \mathbb{1}_{\{u^T x \leq t < s\}} dt d(P - P^*)(x) \\ &= \int_{-\infty}^s (F_u - F_u^*)(t) dt = G_u(s) - G_u^*(s). \end{aligned} \quad (\text{A.4})$$

Since all moments of log-concave densities are finite, we have $\int_{\mathbb{R}^d} xx^T f^*(x) dx < \infty$. So, since $A^* = 0$, we must have $P \in \mathcal{P}_2$. We can therefore integrate by parts as follows:

$$\begin{aligned} 0 &= \int_{\mathbb{R}^d} (u^T x)^2 d(P - P^*)(x) = \int_{-\infty}^{\infty} t^2 d(F_u - F_u^*)(t) \\ &= -2 \int_{-\infty}^{\infty} t(F_u - F_u^*)(t) dt = 2 \int_{-\infty}^{\infty} (G_u - G_u^*)(t) dt. \end{aligned} \quad (\text{A.5})$$

Combining (A.4), (A.5), and the fact that $G_u - G_u^*$ is continuous, we deduce that $G_u = G_u^*$. Thus $F_u = F_u^*$, by the fundamental theorem of calculus and the fact that F_u and F_u^* are both right-continuous. It follows that

$$\mathbb{E}(e^{iu^T X}) = \int_{-\infty}^{\infty} e^{it} dF_u(t) = \int_{-\infty}^{\infty} e^{it} dF_u^*(t) = \mathbb{E}(e^{iu^T X^*}).$$

Since $u \in \mathbb{R}^d$ was arbitrary we deduce that $P = P^*$, so P has a log-concave density.

Proof of Proposition 3. Suppose that the upper semi-continuous log-concave density f^* is the log-concave approximation to $P_1, P_2 \in \mathcal{P}_1$. Then for each $t \in (0, 1)$, we see that f^* also maximises

$$\int_{\mathbb{R}^d} \log f d(tP_1 + (1-t)P_2) = t \int_{\mathbb{R}^d} \log f dP_1 + (1-t) \int_{\mathbb{R}^d} \log f dP_2$$

over all upper semi-continuous log-concave densities f on \mathbb{R}^d .

Proof of Theorem 4. Let $d_2(P, Q) = \inf_{(X, Y) \sim (P, Q)} \{\mathbb{E}\|X - Y\|^2\}^{1/2}$, where the infimum is taken over all pairs (X, Y) of random vectors $X \sim P$ and $Y \sim Q$ on a common probability space. Recall that the infimum here is attained, and that if $P, P_1, P_2, \dots \in \mathcal{P}_2$, then $d_2(P_n, P) \rightarrow 0$ if and only if both $P_n \xrightarrow{d} P$ and $\int_{\mathbb{R}^d} \|x\|^2 dP_n(x) \rightarrow \int_{\mathbb{R}^d} \|x\|^2 dP(x)$. Let P^* denote the distribution corresponding to the log-concave approximation to P_0 , and for $\delta > 0$ to be chosen later, let $\mathcal{Q}_{2, \delta}$ denote the subset of \mathcal{P}_2 consisting of those distributions Q with $d_2(Q, P^*) \leq \delta$ that have a log-concave density. Fix $\epsilon > 0$ and let $Q \in \mathcal{Q}_{2, \delta}$. Let \mathbb{P}_n and \mathbb{Q}_n denote the empirical distribution of an independent sample of size n from P^* and an independent sample from Q , respectively. We require a bound for $\mathbb{P}\{d_2(\mathbb{Q}_n, \mathbb{P}_n) \geq \epsilon/4\}$ that holds uniformly over $\mathcal{Q}_{2, \delta}$ using the following coupling argument. Suppose that $(X_1, Y_1), \dots, (X_n, Y_n)$ are independent and identically distributed pairs with $X_i \sim P^*$ and $Y_i \sim Q$ and that \mathbb{P}_n and \mathbb{Q}_n are obtained as the empirical distribution of X_1, \dots, X_n and Y_1, \dots, Y_n , respectively. We may further suppose that $\mathbb{E}\|X_i - Y_i\|^2 = d_2^2(P^*, Q)$. Using standard results (e.g., Equation (8.2) and Lemma 8.7 of Bickel and Freedman (1981)), we deduce that for $\delta \leq \epsilon^{3/2}/32$,

$$\begin{aligned} \sup_{Q \in \mathcal{Q}_{2, \delta}} \mathbb{P}\left\{d_2(\mathbb{Q}_n, \mathbb{P}_n) > \frac{\epsilon}{4}\right\} &\leq \sup_{Q \in \mathcal{Q}_{2, \delta}} \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^n \|X_i - Y_i\|^2 > \frac{\epsilon^2}{16}\right) \\ &\leq \frac{16}{\epsilon^2} \sup_{Q \in \mathcal{Q}_{2, \delta}} \mathbb{E}(\|X_1 - Y_1\|^2) \leq \frac{16\delta^2}{\epsilon^2} \leq \frac{\epsilon}{2}. \end{aligned}$$

Now let \hat{Q}_n denote the distribution corresponding to the log-concave maximum likelihood estimator constructed from X_1, \dots, X_n , and let $\hat{\mathbb{Q}}_n$ denote the empirical distribution of a sample of size n which, conditional on X_1, \dots, X_n , is drawn independently from \hat{Q}_n . By reducing $\delta > 0$ if necessary, we can assume $\delta \leq \epsilon/4$. It follows that

$$\begin{aligned} \mathbb{P}\{d_2(\hat{\mathbb{Q}}_n, P^*) > \epsilon\} &\leq \sup_{Q \in \mathcal{Q}_{2, \delta}} \mathbb{P}\left\{d_2(\mathbb{Q}_n, Q) > \frac{3\epsilon}{4}\right\} + \mathbb{P}\{d_2(\hat{Q}_n, P^*) > \delta\} \\ &\leq \sup_{Q \in \mathcal{Q}_{2, \delta}} \mathbb{P}\left\{d_2(\mathbb{Q}_n, \mathbb{P}_n) > \frac{\epsilon}{4}\right\} + \mathbb{P}\left\{d_2(\mathbb{P}_n, P^*) > \frac{\epsilon}{4}\right\} \\ &\quad + \mathbb{P}\{d_2(\hat{Q}_n, P^*) > \delta\} \\ &\leq \frac{\epsilon}{2} + \mathbb{P}\left\{d_2(\mathbb{P}_n, P^*) > \frac{\epsilon}{4}\right\} + \mathbb{P}\{d_2(\hat{Q}_n, P^*) > \delta\} \leq \epsilon \quad (\text{A.6}) \end{aligned}$$

for sufficiently large n . The final convergence of the second term here follows from the Weak Law of Large Numbers, while for the third term it follows from Proposition 2(c) of Cule and Samworth (2010) and the Dominated Convergence Theorem.

Let \hat{Q}_{nb} and \tilde{Q}_{nb} denote, respectively, the empirical distribution and the distribution corresponding to the log-concave maximum likelihood estimator of the b th bootstrap sample $X_{1b}^*, \dots, X_{nb}^*$ drawn from \hat{Q}_n . We deduce from (A.6), Theorem 2.15 of Dümbgen, Samworth, and Schuhmacher (2011) and another application of Proposition 2(c) of Cule and Samworth (2010) that there exists $a > 0$ such that

$$\int_{\mathbb{R}^d} e^{a\|x\|} d(\hat{Q}_{nb} - P^*)(x) \xrightarrow{P} 0. \quad (\text{A.7})$$

Now let

$$\begin{aligned} \hat{A}_{nb} &\equiv \hat{\Sigma}_b - \tilde{\Sigma}_b \\ &\equiv \frac{n}{n-1} \int_{\mathbb{R}^d} (x - \bar{X}_b^*)(x - \bar{X}_b^*)^T d\hat{Q}_{nb}(x) - \int_{\mathbb{R}^d} (x - \bar{X}_b^*)(x - \bar{X}_b^*)^T d\tilde{Q}_{nb}(x), \end{aligned}$$

where $\bar{X}_b^* = n^{-1} \sum_{i=1}^n X_{ib}^*$. From (A.6), (A.7), the Dominated Convergence Theorem and the Continuous Mapping Theorem, we have that $\text{tr}(\hat{A}_{nb}) \xrightarrow{P} 0$ as $n \rightarrow \infty$. On the other hand, in the notation of Theorem 1,

$$\text{tr}(\hat{A}) = \text{tr}(\hat{\Sigma}) - \text{tr}(\tilde{\Sigma}) \xrightarrow{P} \text{tr}(\Sigma) - \text{tr}(\Sigma^*) = \text{tr}(A^*) > 0,$$

where the final claim follows from Theorem 3 and the fact that P_0 does not have a log-concave density. Note that this claim holds even if $P_0 \in \mathcal{P}_1 \setminus \mathcal{P}_2$, in which case $\text{tr}(\Sigma) = \infty$.

Write $Z_{nb} = \mathbb{1}_{\{\text{tr}(\hat{A}_{nb}) > \text{tr}(A^*)/2\}}$, and note that Z_{n1}, \dots, Z_{nB} are exchangeable (so in particular, identically distributed). Thus, for any $\alpha \in (0, 1)$,

$$\begin{aligned} \mathbb{P}(\text{Do not reject } H_0) &= \mathbb{P}\left(\frac{1}{B+1} \sum_{b=1}^{B+1} \mathbb{1}_{\{\text{tr}(\hat{A}) > \text{tr}(\hat{A}_{nb})\}} \leq 1 - \alpha\right) \\ &\leq \mathbb{P}\left\{\text{tr}(\hat{A}) \leq \frac{\text{tr}(A^*)}{2}\right\} + \mathbb{P}\left(\frac{1}{B+1} \sum_{b=1}^{B+1} Z_{nb} \geq 1 - \alpha\right) \\ &\leq \mathbb{P}\left\{\text{tr}(\hat{A}) \leq \frac{\text{tr}(A^*)}{2}\right\} + \frac{\mathbb{E}(Z_{n1})}{1 - \alpha} \rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. We deduce that for any given size of test $\alpha \in (0, 1)$, the power at any alternative converges to 1.

Proof of Theorem 5. (a) Note that

$$\hat{C}_n^{\text{LC}}(x) = \operatorname{argmax}_{k \in \{1, \dots, K\}} \frac{N_k}{n} \hat{f}_{n,k}(x).$$

We have that $\int_{\mathbb{R}^d} |\hat{f}_{n,k} - f_k^*| \xrightarrow{a.s.} 0$ as $n \rightarrow \infty$ for every k . In fact, by Theorem 10.8 of Rockafellar (1997), it is almost surely the case that $\hat{f}_{n,k}$ converges to f_k^* uniformly on compact sets in the interior of the support of f_k^* . By the Strong Law of Large Numbers and the fact that the boundary of the support of f_k^* has zero d -dimensional Lebesgue measure, it therefore follows that

$$\hat{C}_n^{\text{LC}}(x) \xrightarrow{a.s.} C^{\text{LCBayes}}(x)$$

for almost all $x \in \mathcal{X}^*$.

In fact, with probability one, $\frac{N_k}{n} \hat{f}_{n,k}$ converges to $\pi_k f_k^*$ uniformly on compact sets in the interior of the support of f_k^* . It follows immediately from this and the Dominated Convergence Theorem that

$$\text{Risk}(\hat{C}_n^{\text{LC}}) \rightarrow \text{Risk}(C^{\text{LCBayes}}).$$

(b) The proof is virtually identical to that of Part (a), so is omitted.

Proof of Proposition 4. The conclusion of Theorem 1 can be stated in the notation of Section 4.3 as

$$\|\tilde{P}_n - P_0^{**}\|_a \xrightarrow{a.s.} 0.$$

The result therefore follows immediately by the Continuous Mapping Theorem.

Proof of Corollary 1. It suffices to show that under (4.1), the functional $\theta(P) = \int_{\mathbb{R}^d} g dP$ is continuous. Fix $a < a_0$ such that $\sup_{x \in \mathbb{R}^d} e^{-a\|x\|} |g(x)| < \infty$, and choose a sequence (P_n) such that $\|P_n - P\|_a \rightarrow 0$. Then

$$\begin{aligned} |\theta(P_n) - \theta(P)| &\leq \int_{\mathbb{R}^d} |g| d|P_n - P| \\ &\leq \sup_{x \in \mathbb{R}^d} e^{-a\|x\|} |g(x)| \int_{\mathbb{R}^d} e^{a\|x\|} d|P_n - P| \\ &= \sup_{x \in \mathbb{R}^d} e^{-a\|x\|} |g(x)| \|P_n - P\|_a \rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. Thus θ is continuous, as required.

References

- An, M. Y. (1995). Log-concave probability distributions: theory and statistical testing. Technical report, Economics Department, Duke University, Durham, NC.
- Balabdaoui, F., Rufibach, K. and Wellner, J. A. (2009). Limit distribution theory for maximum likelihood estimation of a log-concave density. *Ann. Statist.* **37**, 1299-1331.
- Balabdaoui, F. and Wellner, J. A. (2007). Estimation of a k -monotone density: limit distribution theory and the spline connection. *Ann. Statist.* **35**, 2536-2564.
- Balabdaoui, F. and Wellner, J. A. (2010). Estimation of a k -monotone density: characterizations, consistency and minimax lower bounds. *Statistica Neerlandica* **64**, 45-70.

- Bickel, P. J. and Freedman, D. A. (1981). Some asymptotic theory for the bootstrap. *Ann. Statist.* **9**, 1196-1217.
- Braun, W. J. and Hall, P. (2001). Data sharpening for nonparametric inference subject to constraints. *J. Comput. Graph. Statist.* **10**, 786-806.
- Carroll, R. J., Delaigle, A. and Hall, P. (2011). Testing and estimating shape-constrained non-parametric density and regression in the presence of measurement error. *J. Amer. Statist. Assoc.* **106**, 191-202.
- Chen, Y. (2010). Discussion of *Maximum likelihood estimation of a multi-dimensional log-concave density* by M. Cule, R. Samworth and M. Stewart. *J. Roy. Statist. Soc. B* **72**, 590-593.
- Chen, Y. and Samworth, R. J. (2011). Smoothed log-concave maximum likelihood estimation with applications. <http://arxiv.org/abs/1102.1191v3>.
- Cule, M. L. and Dümbgen, L. (2008). On an auxiliary function for log-density estimation. Tech. rep. 71, Universität Bern. <http://arxiv.org/abs/0807.4719>.
- Cule, M. L., Gramacy, R. B. and Samworth, R. J. (2009). LogConcDEAD: an R package for maximum likelihood estimation of a multivariate log-concave density. *J. Statist. Software* **29**, Issue 2.
- Cule, M. L., Gramacy, R. B., Samworth, R. J. and Chen, Y. (2007). LogConcDEAD: maximum likelihood estimation. R package version 1.5-4 <http://CRAN.R-project.org/package=LogConcDEAD>.
- Cule, M. and Samworth, R. (2010). Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electron. J. Statist.* **4**, 254-270.
- Cule, M., Samworth, R. and Stewart, M. (2010). Maximum likelihood estimation of a multi-dimensional log-concave density (with discussion). *J. Roy. Statist. Soc. B* **72**, 545-607.
- Dharmadhikari, S. and Joag-Dev, K. (1988). *Unimodality, Convexity and Applications*. Academic Press, San Diego.
- Dümbgen, L. and Rufibach, K. (2009). Maximum likelihood estimation of a log-concave density and its distribution function: Basic properties and uniform consistency. *Bernoulli* **15**, 40-68.
- Dümbgen, L. and Rufibach, K. (2011). `logcondens`: Computations Related to Univariate Log-Concave Density Estimation. *J. Statist. Software* **39**, 1-28.
- Dümbgen, L., Samworth, R. and Schuhmacher, D. (2011). Approximation by log-concave distributions with applications to regression. *Ann. Statist.* **39**, 702-730.
- Folland, G. B. (1999). *Real Analysis*. Wiley, New York.
- Gopal, V. and Casella, G. (2010). Discussion of *Maximum likelihood estimation of a multi-dimensional log-concave density* by M. Cule, R. Samworth and M. Stewart. *J. Roy. Statist. Soc. B* **72**, 580-582.
- Grenander, U. (1956). On the theory of mortality measurement II. *Skandinavisk Aktuarietidskrift* **39**, 125-153.
- Groeneboom, P., Jongbloed, G. and Wellner, J. A. (2001). Estimation of a convex function: Characterizations and asymptotic theory. *Ann. Statist.* **29**, 1653-1698.
- Grundmann, A. and Möller, M. (1978). Invariant Integration Formulas for the N -Simplex by Combinatorial Methods. *SIAM J. Numer. Anal.* **15**, 282-290.
- Hazelton, M. L. (2011). Assessing log-concavity of multivariate log-concave densities. *Statist. Probab. Lett.* **81**, 121-125.

- Pal, J. K., Woodroffe, M. and Meyer, M. (2007). Estimating a Polya frequency function. In *Complex Datasets and Inverse Problems: Tomography, Networks and Beyond*, 239-249, Institute of Mathematical Statistics, Ohio.
- Pavlidis, M. G. and Wellner, J. A. (2012). Nonparametric estimation of multivariate scale mixtures of uniform densities. *J. Mult. Anal.* **107**, 71-89.
- Prékopa, A. (1973). On logarithmically concave measures and functions. *Acta Sci. Math. (Szeged)* **34**, 335-343.
- Rockafellar, R. T. (1997). *Convex Analysis*. Princeton University Press, Princeton.
- Rufibach, K. and Dümbgen, L. (2006). **logcondens**: Estimate a log-concave probability density from *i.i.d.* Observations R package version 2.01 <http://cran.r-project.org/web/packages/logcondens/>.
- Samworth, R. J. and Yuan, M. (2012). Independent component analysis via nonparametric maximum likelihood estimation. *Ann. Statist.* , **40**, 2973-3002.
- Schuhmacher, D. and Dümbgen, L. (2010). Consistency of multivariate log-concave density estimators. *Statist. Probab. Lett.* **80**, 376-380.
- Schuhmacher, D., Hüsler, A. and Dümbgen, L. (2011). Multivariate log-concave distributions as a nearly parametric model. *Statistics and Risk Modeling* **28**, 277-295.
- Street, W. N. Wolberg, W. H., and Mangasarian, O. L. (1993). Nuclear feature extraction for breast tumor diagnosis. In *Proc. Int. Symp. Electronic Imaging: Science and Technology*, **1905**, 861-870.
- Stroud, A. H. (1971). *Approximate Calculation of Multiple Integrals*. Prentice-Hall, Englewood Cliffs, NJ.
- Walther, G. (2002). Detecting the presence of mixing with multiscale maximum likelihood. *J. Amer. Statist. Assoc.* **97**, 508-513.
- Walther, G. (2009). Inference and modeling with log-concave distributions. *Statist. Sci.* **24**, 319-327.

Statistical Laboratory, University of Cambridge, UK.

E-mail: y.chen@statslab.cam.ac.uk

Statistical Laboratory, University of Cambridge, UK.

E-mail: r.samworth@statslab.cam.ac.uk

(Received September 2011; accepted October 2012)