



Confidence Bands for a Log-Concave Density

Guenther Walther^a, Alnur Ali^{a,b}, Xinyue Shen^{b,c}, and Stephen Boyd^b

^aDepartment of Statistics, Stanford University, Stanford, CA; ^bDepartment of Electrical Engineering, Stanford University, Stanford, CA; ^cShenzhen Research Institute of Big Data and Future Network of Intelligence Institute, Chinese University of Hong Kong, Longgang, Shenzhen

ABSTRACT

We present a new approach for inference about a univariate log-concave distribution: Instead of using the method of maximum likelihood, we propose to incorporate the log-concavity constraint in an appropriate nonparametric confidence set for the cdf F . This approach has the advantage that it automatically provides a measure of statistical uncertainty and it thus, overcomes a marked limitation of the maximum likelihood estimate. In particular, we show how to construct confidence bands for the density that have a finite sample guaranteed confidence level. The nonparametric confidence set for F which we introduce here has attractive computational and statistical properties: It allows to bring modern tools from optimization to bear on this problem via difference of convex programming, and it results in optimal statistical inference. We show that the width of the resulting confidence bands converges at nearly the parametric $n^{-\frac{1}{2}}$ rate when the log density is k -affine. Supplementary materials for this article are available online.

ARTICLE HISTORY

Received December 2020
Accepted May 2022

KEYWORDS

Density estimation;
Difference of convex
programs; Goodness-of-fit
methods; Log-concave
distribution; Nonparametric
methods

1. Introduction

Statistical inference under shape constraints has been the subject of continued considerable research activity. Imposing shape constraints on the inference about a function f , that is, assuming that the function satisfies certain qualitative properties, such as monotonicity or convexity on certain subsets of its domain, has two main motivations: First, such shape constraints may directly arise from the problem under investigation and it is then desirable that the result of the inference reflects this fact. The second reason is that alternative nonparametric estimators typically involve the choice of a tuning parameter, such as the bandwidth in the case of a kernel estimator. A good choice for such a tuning parameter is usually far from trivial. More importantly, selecting a tuning parameter injects a certain amount of subjectivity into the estimator, and the resulting choice may prove quite consequential for relevant aspects of the inference. In contrast, imposing shape constraints often allows to derive an estimator that does not depend on a tuning parameter while at the same time exhibiting a good statistical performance, such as achieving optimal minimax rates of convergence to the underlying function f .

This article is concerned with inference about a univariate log-concave density, that is, a density of the form

$$f(x) = \exp \phi(x),$$

where $\phi : \mathbf{R} \rightarrow [-\infty, \infty)$ is a concave function. It was argued in Walther (2002) that log-concave densities represent an attractive and useful nonparametric surrogate for the class of Gaussian distributions for a range of problems in inference and modeling. The appealing properties of this class are that

it contains most commonly encountered parametric families of unimodal densities with exponentially decaying tails and that the class is closed under convolution, affine transformations, convergence in distribution and marginalization. In a similar vein as a normal distribution can be seen as a prototypical model for a homogenous population, one can use the class of log-concave densities as a more flexible nonparametric model for this task, from which heterogeneous models can then be built for example, via location mixtures. Historically such homogeneous distributions have been modeled with unimodal densities, but it is known that the maximum likelihood estimate (MLE) of a unimodal density does not exist; see, for example, Birgé (1997). In contrast, it was shown in Walther (2002) that the MLE of a log-concave density exists and can be computed with readily available algorithms. Therefore, the class of log-concave densities is a sufficiently rich nonparametric model while at the same time it is small enough to allow nonparametric inference without a tuning parameter.

Due to these attractive properties there has been a considerable research activity in the last 15 years about the statistical properties of the MLE, computational aspects, applications in modeling and inference, and the multivariate setting. Many of the key properties of the MLE are now well understood: Existence of the MLE was shown in Walther (2002), consistency was proved by Pal, Woodroffe, and Meyer (2007), and rates of convergence in certain uniform metrics was established by Dümbgen and Rufibach (2009). Balabdaoui, Rufibach, and Wellner (2009) provided pointwise limit distribution theory, while Doss and Wellner (2016) and Kim and Samworth (2016) gave rates of convergence in the Hellinger metric, and Kim, Guntuboyina, and Samworth (2018) proved adaptation properties.

Accompanying results for the multivariate case are given in, for example, Cule, Samworth, and Stewart (2010), Schuhmacher and Dümbgen (2010), Seregin and Wellner (2010), Kim and Samworth (2016), and Feng et al. (2021).

Computation of the univariate MLE was initially approached with the Iterative Convex Minorant Algorithm, see Walther (2002), Pal, Woodroffe, and Meyer (2007), and Rufibach (2007), but it appears that the fastest algorithms currently available are the active set algorithm given in Dümbgen, Hüsler, and Rufibach (2007) and the constrained Newton method proposed by Liu and Wang (2018).

Overviews of some of these results and other work involving modeling and inference with log-concave distributions are given in the review papers of Walther (2009), Saumard and Wellner (2014), and Samworth (2018).

Notably, the existing methodology for estimating a log-concave density appears to be exclusively focused on the method of maximum likelihood. Here we will employ a different methodology: We will derive a confidence band by intersecting the log-concavity constraint with a goodness-of-fit test. One important advantage of this approach is that such a confidence band satisfies a key principle of statistical inference: an estimate needs to be accompanied by some measure of standard error in order to be useful for inference. There appears to be no known method for obtaining such a confidence band via the maximum likelihood approach. Balabdaoui, Rufibach, and Wellner (2009) construct *pointwise* confidence intervals for a log-concave density based on asymptotic limit theory, which requires to estimate the second derivative of $\log f(x)$. Azadbakhsh, Jankowski, and Gao (2014) compare several methods for estimating this nuisance parameter and they report that this task is quite difficult. An alternative approach is given by Deng, Han, and Sen (2020). In Section 5 we compare the confidence bands introduced here with pointwise confidence intervals obtained via the asymptotic limit theory as well as with the bootstrap. Of course, pointwise confidence intervals have a different goal than confidence bands. The pointwise intervals will be shorter but the confidence level will not hold simultaneously across multiple locations. In contrast, the method we introduce here comes with strong guarantees in terms of finite sample valid coverage levels across locations.

2. Constructing a Confidence Band for a Log-Concave Density

Given data X_1, \dots, X_n from a log-concave density f we want to find functions $\hat{\ell}(x) = \hat{\ell}(x, X_1, \dots, X_n)$ and $\hat{\mu}(x) = \hat{\mu}(x, X_1, \dots, X_n)$ such that

$$\mathbb{P}_f \left\{ \hat{\ell}(x) \leq f(x) \leq \hat{\mu}(x) \text{ for all } x \in \mathbf{R} \right\} \geq 1 - \alpha$$

for a given confidence level $1 - \alpha \in (0, 1)$. It is well known that in the case of a general density f no nontrivial confidence interval $(\hat{\ell}(x), \hat{\mu}(x))$ exists, see, for example, Donoho (1988). However, assuming a shape-constraint for f such as log-concavity allows to construct pointwise and uniform confidence statements as follows:

Let $C_n(\alpha)$ be a $1 - \alpha$ confidence set for the distribution function F of f , that is,

$$\mathbb{P}_F \{F \in C_n(\alpha)\} \geq 1 - \alpha. \tag{1}$$

Such a nonparametric confidence set always exists, for example, the Kolmogorov-Smirnov bands give a confidence set for F (albeit a nonoptimal one). Define

$$\hat{\ell}(x) := \inf_{f \text{ is log-concave and } F \in C_n(\alpha)} f(x) \tag{2}$$

and define $\hat{\mu}(x)$ analogously with sup in place of inf. If f is log-concave then (1) and (2) imply

$$\mathbb{P}_f \left\{ \hat{\ell}(x) \leq f(x) \leq \hat{\mu}(x) \right\} \geq 1 - \alpha,$$

so we obtain a $1 - \alpha$ confidence interval for $f(x)$ by solving the optimization problem (2). Moreover, if we solve (2) at multiple locations x_1, \dots, x_m , then we obtain

$$\mathbb{P}_f \left\{ \hat{\ell}(x_i) \leq f(x_i) \leq \hat{\mu}(x_i) \text{ for all } i = 1, \dots, m \right\} \geq 1 - \alpha. \tag{3}$$

So the coverage probability is automatically simultaneous across multiple locations and comes with a finite sample guarantee, since it is inherited from the confidence set C_n in (1). Likewise, the quality of the confidence band, as measured for example, by the width $\hat{\mu}(x) - \hat{\ell}(x)$, will also derive from C_n , which therefore, plays a central role in this approach. Finally, the log-concavity constraint allows to extend the confidence set (3) to a confidence band on the real line, as we will show in Section 2.4.

Hengartner and Stark (1995), Dümbgen (1998), Dümbgen (2003), and Davies and Kovac (2004) employ the above approach for inference about a unimodal or a k -modal density. Here we introduce a new confidence set $C_n(\alpha)$ for F . This confidence set is adapted from methodology developed in the abstract Gaussian White Noise model by Walther and Perry (2019) for optimal inference in settings related to the one considered here. Therefore, this confidence set should also prove useful for the works about inference in the unimodal and k -modal setting cited above.

The key conceptual problem for solving the optimization problem (2) is that f is infinite dimensional. We will overcome this by using the log-concavity of f to relax $C_n(\alpha)$ to a finite dimensional superset, which makes it possible to compute (2) with fast optimization algorithms. We will address these tasks in turn in the following sections.

2.1. A Confidence Set for F

Given X_1, \dots, X_n iid from the continuous cdf F we set $s_n := \lceil \log_2 \log n \rceil$ and

$$x_i := X_{(1+(i-1)2^{s_n})}, \quad i = 1, \dots, m := \left\lfloor \frac{n-1}{2^{s_n}} \right\rfloor + 1, \tag{4}$$

where $X_{(j)}$ denotes the j th order statistic. Our analysis will use only the subset $\{x_i\}$ of the data, that is, the set containing every $\log n$ th order statistic; see Remark 3 for why this is sufficient.

Translating the methodology of the ‘‘Bonferroni scan’’ developed in Walther and Perry (2019) from the Gaussian White

Noise model to the density setting suggests employing a confidence set of the form

$$C_n(\alpha) := \left\{ F : c_{jkB} \leq F(x_k) - F(x_j) \leq d_{jkB} \text{ for all } (j, k) \in \mathcal{I} = \bigcup_B \mathcal{I}_B \right\}.$$

with $c_{jkB}, d_{jkB}, \mathcal{I}$ given below. The idea is to choose an index set \mathcal{I} that is rich enough to detect relevant deviations from the empirical distribution, but which is also sparse enough so that the $|\mathcal{I}|$ constraints can be combined with a simple weighted Bonferroni adjustment and still result in optimal inference. The second key ingredient of this construction is to let the weights of the Bonferroni adjustment depend on $j - i$ in a certain way. See Walther and Perry (2019) for a comparison of the finite sample and asymptotic performance of this approach with other relevant calibrations, such as the ones used in the works cited above.

Note that the confidence set $C_n(\alpha)$ checks the probability content of random intervals (x_j, x_k) , which automatically adapt to the empirical distribution. This makes it possible to detect relevant deviations from the empirical distribution with a relatively small number of such intervals, which is key for making the Bonferroni adjustment powerful as well as for efficient computation. Moreover, using such random intervals makes the bounds c_{jkB}, d_{jkB} distribution-free since $F(x_k) - F(x_j) \sim \text{Beta}((k - j)2^{s_n}, n + 1 - (k - j)2^{s_n})$, see chap. 3.1 in Shorack and Wellner (1986).

The precise specifications of $c_{jkB}, d_{jkB}, \mathcal{I}$ are as follows:

$$\begin{aligned} \mathcal{I} &:= \bigcup_{B=0}^{B_{\max}} \mathcal{I}_B, \text{ where } B_{\max} := \lfloor \log_2 \frac{n}{8} \rfloor - s_n \\ \mathcal{I}_B &:= \left\{ (j, k) : j = 1 + (i - 1)2^B, k = 1 + i2^B \right. \\ &\quad \left. \text{for } i = 1, \dots, n_B := \lfloor \frac{n - 1}{2^{B+s_n}} \rfloor \right\} \\ c_{jkB} &:= c_B := q\text{Beta} \left(\frac{\alpha}{2(B + 2)n_B t_n}, 2^{B+s_n}, n + 1 - 2^{B+s_n} \right) \\ d_{jkB} &:= d_B := q\text{Beta} \left(1 - \frac{\alpha}{2(B + 2)n_B t_n}, 2^{B+s_n}, n + 1 - 2^{B+s_n} \right) \end{aligned}$$

where $t_n := \sum_{B=0}^{B_{\max}} \frac{1}{B+2}$ and $q\text{Beta}(\alpha, r, s)$ denotes the α -quantile of the beta distribution with parameters r and s . The term $\frac{1}{B+2}$ is a weighting factor in the Bonferroni adjustment which results in an advantageous statistical performance, see Walther and Perry (2019). It follows from the union bound that $\mathbb{P}_F(F \in C_n(\alpha)) \geq 1 - \alpha$ whenever F is continuous.

Remark. 1. An alternative way to control the distribution of $F(x_k) - F(x_j)$ is via a log-likelihood ratio type transformation and Hoeffding’s inequality, see Rivera and Walther (2013) and Li et al. (2020). This results in a loss of power due to the slack in Hoeffding’s inequality and the slack from inverting the log-likelihood ratio transformation with an inequality. Simulations show that the above approach using an exact beta distribution is less conservative despite the use of Bonferroni’s inequality to combine the statistics across \mathcal{I} .

2. The inference is based on the statistic $F((x_j, x_k))$, that is, the unknown F evaluated on the random interval (x_j, x_k) , rather than on the more commonly used statistic $F_n(I)$, which evaluates the empirical measure on deterministic intervals I . The

latter statistic follows a binomial distribution whose discreteness makes it difficult to combine these statistics across \mathcal{I} using Bonferroni’s inequality without incurring substantial conservatism and hence, loss of power. This is another important reason for using random intervals (x_j, x_k) besides the adaptivity property mentioned above. Moreover, a deterministic system of intervals would have to be anchored around the range of the data and this dependence on the data is difficult to account for and is therefore, typically glossed over in the inference.

3. The definition of x_j in (4) means that we do not consider intervals $(X_{(j)}, X_{(k)})$ with $k - j < \log n$. Thus, as opposed to the regression setting in Walther and Perry (2019) we omit the first block¹ of intervals. This derives from the folklore knowledge in density estimation that at least $\log n$ observations are required in order to obtain consistent inference simultaneously across such intervals. Indeed, this choice is sufficient to yield the asymptotic optimality result in Theorem 1.

We further simplify the construction in Walther and Perry (2019) by restricting ourselves to a dyadic spacing of the indices $k - j$ since we already obtain quite satisfactory results with this set of intervals.

2.2. Bounds for $\int_a^b f$ when f is Log-Concave

The confidence set $C_n(\alpha)$ describes a set of plausible distributions in terms of $\int_{x_j}^{x_k} f(t) dt$ for infinite dimensional f . In the special case when f is log-concave it is possible to construct a finite dimensional superset of $C_n(\alpha)$ by deriving bounds for this integral in terms of functions of a finite number of variables:

Lemma 1. Let f be a univariate log-concave function. For given $x_1 < \dots < x_m$ write $\ell_i := \log f(x_i)$, $i = 1, \dots, m$. Then there exist real numbers g_2, \dots, g_{m-1} such that

$$\ell_j \leq \ell_i + g_i (x_j - x_i) \text{ for all } i \in \{2, \dots, m - 1\}, j \in \{i - 1, i + 1\}$$

and

$$\begin{aligned} (x_{i+1} - x_i) \exp(\ell_i) E(\ell_{i+1} - \ell_i) &\leq \int_{x_i}^{x_{i+1}} f(t) dt \\ &\leq \begin{cases} \exp(\ell_i)(x_{i+1} - x_i) E(g_i(x_{i+1} - x_i)), & i \in \{2, \dots, m - 1\} \\ \exp(\ell_{i+1})(x_{i+1} - x_i) E(g_{i+1}(x_i - x_{i+1})), & i \in \{1, \dots, m - 2\} \end{cases} \end{aligned}$$

where

$$E(s) := \int_0^1 \exp(st) dt = \begin{cases} \frac{\exp(s) - 1}{s} & \text{if } s \neq 0 \\ 1 & \text{if } s = 0 \end{cases}$$

is a strictly convex and infinitely often differentiable function.

The proof of Lemma 1 is given in the Appendix, supplementary materials. Importantly, the bounds given in the lemma are convex and smooth functions of the g_i and ℓ_i , despite the fact that these variables appear in the denominator in the formula for E . This makes it possible to bring fast optimization routines to bear on the problem (2).

¹We also shift the index B to let it start at 0 rather than at 2. This results in a simpler notation but does not change the methodology.

2.3. Computing Pointwise Confidence Intervals

We are now in position to define a superset of $\mathcal{C}_n(\alpha)$ by relaxing the inequalities $c_B \leq \int_{x_j}^{x_k} f(t) dt \leq d_B$ in the definition of $\mathcal{C}_n(\alpha)$. To this end define for $i = 1, \dots, m - 1$ the functions

$$L_i(\mathbf{x}, \boldsymbol{\ell}) := (x_{i+1} - x_i) \exp(\ell_i) E(\ell_{i+1} - \ell_i)$$

$$U_i(\mathbf{x}, \boldsymbol{\ell}, \mathbf{g}) := \begin{cases} \exp(\ell_i)(x_{i+1} - x_i) \\ E(g_i(x_{i+1} - x_i)), & i = m - 1 \\ \exp(\ell_{i+1})(x_{i+1} - x_i) \\ E(g_{i+1}(x_i - x_{i+1})), & i \in \{1, \dots, m - 2\} \end{cases}$$

$$V_i(\mathbf{x}, \boldsymbol{\ell}, \mathbf{g}) := \begin{cases} \exp(\ell_i)(x_{i+1} - x_i) \\ E(g_i(x_{i+1} - x_i)), & i \in \{2, \dots, m - 1\} \\ \exp(\ell_{i+1})(x_{i+1} - x_i) \\ E(g_{i+1}(x_i - x_{i+1})), & i = 1 \end{cases}$$

where $\mathbf{x} = (x_1, \dots, x_m)$, $\boldsymbol{\ell} = (\ell_1, \dots, \ell_m)$, $\mathbf{g} = (g_1, \dots, g_{m-1})$ and $E(\cdot)$ is defined in Lemma 1.

Given the subset $x_1 < \dots < x_m$ of the order statistics defined in (4), we define $\tilde{\mathcal{C}}_n(\alpha)$ to be the set of densities f for which there exist real g_1, \dots, g_{m-1} such that $\ell_i := \log f(x_i)$, $i = 1, \dots, m$, satisfy (5)–(8):

$$\ell_j \leq \ell_i + g_i(x_j - x_i) \quad \text{for all } i \in \{2, \dots, m - 1\}, j \in \{i - 1, i + 1\}. \tag{5}$$

For $B = 0, \dots, B_{\max}$:

$$c_B \leq \sum_{i=j}^{k-1} U_i(\mathbf{x}, \boldsymbol{\ell}, \mathbf{g}) \quad \text{for all } (j, k) \in \mathcal{I}_B \tag{6}$$

$$c_B \leq \sum_{i=j}^{k-1} V_i(\mathbf{x}, \boldsymbol{\ell}, \mathbf{g}) \quad \text{for all } (j, k) \in \mathcal{I}_B \tag{7}$$

$$\sum_{i=j}^{k-1} L_i(\mathbf{x}, \boldsymbol{\ell}) \leq d_B \quad \text{for all } (j, k) \in \mathcal{I}_B. \tag{8}$$

Now we can implement a computable version of the confidence bound (2) by optimizing over $\tilde{\mathcal{C}}_n(\alpha)$ rather than over $\mathcal{C}_n(\alpha)$. Note that if f is log-concave then it follows from Lemma 1 that $f \in \mathcal{C}_n(\alpha)$ implies $f \in \tilde{\mathcal{C}}_n(\alpha)$. This proves the following key result:

Proposition 1. If f is log-concave then $\mathbb{P}_f\{f \in \tilde{\mathcal{C}}_n(\alpha)\} \geq 1 - \alpha$. Consequently, if we define pointwise lower and upper confidence bounds at the x_i , $i = 1, \dots, m$, via the optimization problems

$$\hat{\ell}(x_i) := \min \ell(x_i) \tag{9}$$

subject to $f \in \tilde{\mathcal{C}}_n(\alpha)$, that is, subject to (5)–(8)

$$\hat{\mu}(x_i) := \max \ell(x_i)$$

subject to $f \in \tilde{\mathcal{C}}_n(\alpha)$, that is, subject to (5)–(8),

then the following simultaneous confidence statement holds whenever f is log-concave:

$$\mathbb{P}_f \left\{ \exp(\hat{\ell}(x_i)) \leq f(x_i) \leq \exp(\hat{\mu}(x_i)) \quad \text{for all } i = 1, \dots, m \right\} \geq 1 - \alpha.$$

It is an important feature of these confidence bounds that they come with a finite sample guaranteed confidence level $1 - \alpha$. On the other hand, it is desirable that the construction is not overly conservative (i.e., has coverage not much larger than $1 - \alpha$) as otherwise it would result in unnecessarily wide confidence bands. This is the motivation for deriving a statistically optimal confidence set in Section 2.1 and for deriving bounds in Lemma 1 that are sufficiently tight. Indeed, it will be shown in Section 4 that the above construction results in statistically optimal confidence bands.

2.4. Constructing Confidence Bands

The simultaneous pointwise confidence bounds $(\hat{\ell}(x_i), \hat{\mu}(x_i))$, $i = 1, \dots, m$, from the optimization problem (9) imply a confidence band on the real line due to the concavity of $\log f$. In more detail, we can extend the definition of $\hat{\ell}$ to the real line simply by interpolating between the $\hat{\ell}(x_i)$:

$$\hat{\ell}(x) := \begin{cases} \hat{\ell}(x_i) + (x - x_i) \frac{\hat{\ell}(x_{i+1}) - \hat{\ell}(x_i)}{x_{i+1} - x_i} & \text{if } x \in [x_i, x_{i+1}), \\ -\infty & \text{otherwise.} \end{cases} \tag{10}$$

Then $\log f(x_i) \geq \hat{\ell}(x_i)$ for $i = 1, \dots, m$ implies $\log f(x) \geq \hat{\ell}(x)$ for $x \in \mathbf{R}$ since $\log f$ is concave and $\hat{\ell}$ is piecewise linear. (In fact, it follows from (9) that $\hat{\ell}$ is also concave.)

In order to construct an upper confidence bound note that concavity of $\log f$ together with $\hat{\ell}(x_i) \leq \log f(x_i) \leq \hat{\mu}(x_i)$ for all $i \in \{1, \dots, m\}$ implies for $x > x_k$ with $k \in \{2, \dots, m\}$:

$$\begin{aligned} \frac{\log f(x) - \hat{\mu}(x_k)}{x - x_k} &\leq \frac{\log f(x) - \log f(x_k)}{x - x_k} \\ &\leq \min_{j \in \{1, \dots, k-1\}} \frac{\log f(x_k) - \log f(x_j)}{x_k - x_j} \\ &\leq \min_{j \in \{1, \dots, k-1\}} \frac{\hat{\mu}(x_k) - \hat{\ell}(x_j)}{x_k - x_j} =: L_k \end{aligned}$$

and likewise for $x < x_k$ with $k \in \{1, \dots, m - 1\}$:

$$\frac{\hat{\mu}(x_k) - \log f(x)}{x_k - x} \geq \max_{j \in \{k+1, \dots, m\}} \frac{\hat{\ell}(x_j) - \hat{\mu}(x_k)}{x_j - x_k} =: R_k.$$

Hence, $\log f$ is bounded above by

$$\hat{\mu}(x) := \begin{cases} \hat{\mu}(x_{i+1}) + R_{i+1}(x - x_{i+1}) & \text{if } x \in (x_i, x_{i+1}), \\ & i \in \{0, 1\}, \\ & \text{where } x_0 := -\infty \\ M_i(x) & \text{if } x \in [x_i, x_{i+1}), \\ & i \in \{2, \dots, m - 2\} \\ \hat{\mu}(x_i) + L_i(x - x_i) & \text{if } x \in [x_i, x_{i+1}), \\ & i \in \{m - 1, m\}, \\ & \text{where } x_{m+1} := \infty \end{cases}$$

with

$$\begin{aligned}
M_i(x) &:= \min\left(\hat{\mu}(x_i) + L_i(x - x_i), \hat{\mu}(x_{i+1}) + R_{i+1}(x - x_{i+1})\right) \\
&= \left(\hat{\mu}(x_i) + L_i(x - x_i)\right) \mathbf{1}(x \in [x_i, \bar{x}_i]) \\
&\quad + \left(\hat{\mu}(x_{i+1}) + R_{i+1}(x - x_{i+1})\right) \mathbf{1}(x \in [\bar{x}_i, x_{i+1}])
\end{aligned}$$

where $\bar{x}_i := \frac{\hat{\mu}(x_{i+1}) - \hat{\mu}(x_i) + L_i x_i - R_{i+1} x_{i+1}}{L_i - R_{i+1}}$.

Thus, we proved:

Proposition 2. If f is log-concave then

$$\mathbb{P}_f \left\{ \exp(\hat{\ell}(x)) \leq f(x) \leq \exp(\hat{\mu}(x)) \text{ for all } x \in \mathbf{R} \right\} \geq 1 - \alpha.$$

The upper bound $\hat{\mu}(x)$ need not be concave but it is minimal in the sense that it can be shown that for every real x there exist a concave function g with $\hat{\ell}(x_i) \leq g(x_i) \leq \hat{\mu}(x_i)$ for all $i \in \{1, \dots, m\}$ and $g(x) = \hat{\mu}(x)$.

As an alternative to $\exp(\hat{\mu}(x))$ we tried a simple interpolation between the points $(x_i, \exp(\hat{\mu}(x_i)))$. This interpolation will result in a smoother bound than $\exp(\hat{\mu}(x))$, but the coverage guarantee of **Proposition 2** does not apply any longer. However, the difference between $\exp(\hat{\mu}(x))$ and the interpolation will vanish as the sample size increases (or by increasing the number m of design points x_i in (4) for a given sample size n), and the simulations in **Section 5** show that the empirical coverage exceeds the nominal level in all cases considered. Therefore, we also recommend this interpolation as a simple and smoother alternative to $\exp(\hat{\mu}(x))$.

Finally, we point out that the computational effort can be lightened simply by solving the optimization problem (9) for a subset of $\{x_i, 1 \leq i \leq m\}$ and then constructing $\hat{\ell}$ and $\hat{\mu}$ as described above based on that smaller subset of x_i . Such a confidence band will still satisfy **Proposition 2**, but it will be somewhat wider at locations between those design points x_i as it is based on fewer pointwise confidence bounds. Hence, there is a tradeoff between the width of the band and the computational effort required. While a larger subset of the x_i will result in a somewhat reduced width of the band between the x_i , there are diminishing returns as the width at the x_i will not change. It follows from **Theorem 1** that solving the optimization problem (9) for $\{x_i, 1 \leq i \leq m\}$ with m given in (4) is sufficient to produce statistically optimal confidence bands in a representative setting.

3. Solving the Optimization Problem

Next we describe a method for computing the pointwise confidence intervals $(\hat{\ell}(x_i), \hat{\mu}(x_i))$, $i = 1, \dots, m$, from observations X_i , $i = 1, \dots, n$, by efficiently solving the optimization problems (9). Constructing the confidence band is then straightforward with the post-processing steps given in **Section 2.4**.

Inspecting the optimization problems (9), we see that these problems possess some interesting structure: The criterion functions are linear, and the constraints (5) and (8) are convex. However, the constraints (6) and (7) are nonconvex. Finding the global minimizer of a nonconvex optimization problem (even a well-structured one) can be challenging; instead, we focus on a method for finding critical points of the problems (9). Taking a closer look at the nonconvex constraints (6) and (7), we

make the simple observation that they may be expressed as the difference of two convex functions (namely, a constant function minus a convex function). This property puts the problems (9) into the special class of nonconvex problems commonly referred to as *difference of convex programs* (Hartman 1959; Tao 1986; Horst and Thoai 1999; Horst, Pardalos, and Van Thoai 2000), for which a critical point can be efficiently found. The class of difference of convex programs is quite broad, encompassing many problems encountered in practice, with a good amount of research into this area continuing on today. Important references include Hartman (1959), Tao (1986), Horst and Thoai (1999), Horst, Pardalos, and Van Thoai (2000), Yuille and Rangarajan (2003), Smola, Vishwanathan, and Hofmann (2005), and Lipp and Boyd (2016).

A natural approach to finding a critical point of a difference of convex program is to linearize the nonconvex constraints, then solve the convexified problem using any suitable off-the-shelf solver, and repeating these steps as necessary. This strategy underlies the well-known convex-concave procedure in Yuille and Rangarajan (2003), a popular heuristic for difference of convex programs. In more detail, the convex-concave iteration as applied to the problems (9) works as follows: Given feasible initial points, we first replace the (nonconvex) constraints (6) and (7) by their first-order Taylor approximations centered around the initial points. Formally, letting $\ell^{(K)}$ and $\mathbf{g}^{(K)}$ denote the log-densities and subgradients on iteration K , respectively, we form

$$\begin{aligned}
\hat{U}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &:= U_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}) \\
&\quad + \langle \nabla U_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}), (\ell, \mathbf{g}) - (\ell^{(K)}, \mathbf{g}^{(K)}) \rangle
\end{aligned} \tag{13}$$

$$\begin{aligned}
\hat{V}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &:= V_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}) \\
&\quad + \langle \nabla V_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}), (\ell, \mathbf{g}) - (\ell^{(K)}, \mathbf{g}^{(K)}) \rangle,
\end{aligned} \tag{14}$$

for $i = 1, \dots, m - 1$. We then solve the convexified problems (using the constraints (13), (14) instead of (6), (7)) with any off-the-shelf solver. Then we recompute the approximations using the obtained solutions to the convexified problems and repeat these steps until an appropriate stopping criterion has been satisfied (e.g., some predetermined maximum number of iterations has been reached, the change in criterion values are smaller than some specified tolerance, the sum of the slack variables is less than some tolerance, and/or we have that $\tau_K \geq \tau_{\max}$). From this description, it may be apparent to the reader that the convex-concave procedure is actually a generalization of the majorization-minimization class of algorithms (which includes the well-known expectation-maximization algorithm as a special case).

We give a complete description of the convex-concave procedure as applied to the optimization problems (9) in **Algorithm 1** appearing above, along with one important modification that we explain now. In practice it is not easy to obtain feasible initial points for the problems (9). Therefore, the penalty convex-concave procedure, a modification to the basic convex-concave procedure that was introduced by Le Thi and Dinh (2014), Dinh and Le Thi (2014), and Lipp and Boyd (2016), works around this issue by allowing for an (arbitrary) infeasible initial

Algorithm 1 Penalty convex-concave procedure for computing pointwise confidence intervals for a log-concave density

Input: Subset of observations x_t , $t = 1, \dots, m$; collection of interval endpoints $\mathcal{I} = \bigcup_B \mathcal{I}_B$; initial points $\ell^{(0)}, \mathbf{g}^{(0)}$; initial penalty strength $\tau_0 > 0$; penalty growth factor $\kappa > 1$; maximum penalty strength $\tau_{\max} > \tau_0$; maximum number of iterations K_{\max}

For $t = 1, \dots, m, K = 0, \dots, K_{\max}$:

Convexify the constraints (6) and (7), by forming the linearizations, for $i = 1, \dots, m - 1$,

$$\begin{aligned} \hat{U}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &= U_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}) + \langle \nabla U_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}), (\ell, \mathbf{g}) - (\ell^{(K)}, \mathbf{g}^{(K)}) \rangle \\ \hat{V}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &= V_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}) + \langle \nabla V_i(\mathbf{x}, \ell^{(K)}, \mathbf{g}^{(K)}), (\ell, \mathbf{g}) - (\ell^{(K)}, \mathbf{g}^{(K)}) \rangle. \end{aligned}$$

Solve the pair of convexified problems

$$\ell_t^{(K+1)} = \min \ell_t + \tau_K \cdot \sum_{B=0}^{B_{\max}} \sum_{(j,k) \in \mathcal{I}_B} s_{B,j,k} \tag{11}$$

subject to (5), (8), and

For $B = 0, \dots, B_{\max}$:

$$\begin{aligned} c_B - \sum_{i=j}^{k-1} \hat{U}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &\leq s_{B,j,k} \text{ for all } (j, k) \in \mathcal{I}_B \\ c_B - \sum_{i=j}^{k-1} \hat{V}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &\leq s_{B,j,k} \text{ for all } (j, k) \in \mathcal{I}_B \\ s_{B,j,k} &\geq 0 \text{ for all } (j, k) \in \mathcal{I}_B \end{aligned}$$

$$\mu_t^{(K+1)} = \max \ell_t - \tau_K \cdot \sum_{B=0}^{B_{\max}} \sum_{(j,k) \in \mathcal{I}_B} s_{B,j,k} \tag{12}$$

subject to (5), (8), and

For $B = 0, \dots, B_{\max}$:

$$\begin{aligned} c_B - \sum_{i=j}^{k-1} \hat{U}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &\leq s_{B,j,k} \text{ for all } (j, k) \in \mathcal{I}_B \\ c_B - \sum_{i=j}^{k-1} \hat{V}_i(\mathbf{x}, \ell, \mathbf{g}; \ell^{(K)}, \mathbf{g}^{(K)}) &\leq s_{B,j,k} \text{ for all } (j, k) \in \mathcal{I}_B \\ s_{B,j,k} &\geq 0 \text{ for all } (j, k) \in \mathcal{I}_B. \end{aligned}$$

Update the penalty strength, by setting $\tau_{K+1} = \min\{\kappa \cdot \tau_K, \tau_{\max}\}$.

Output: Pointwise confidence intervals $(\ell_t^{(K)}, \mu_t^{(K)})$, $t = 1, \dots, m$.

point and then gradually driving the iterates into feasibility by adding a penalty for constraint violations into the criterion that grows with the number of iterations (explaining the word “penalty” in the name of the procedure), through the use of slack variables.

Standard convergence theory for the (penalty) convex-concave procedure (see, e.g., Section 3.1 in Lipp and Boyd (2016) as well as Theorem 10 in Sriperumbudur and Lanckriet (2009) and Proposition 1 in Khamaru and Wainwright (2018) tells us that the criterion values (11) and (12) generated by Algorithm 1 converge. Moreover, under regularity conditions (see Section 3.1 in Lipp and Boyd 2016), the iterates generated by Algorithm 1 converge to critical points of the problems (9). At convergence, the pointwise confidence intervals generated

by Algorithm 1 can be turned in confidence bands as described in Section 2.4.

Finally, we mention that although not necessary, additionally linearizing the (convex) constraints (8) around the previous iterate, that is, forming

$$\begin{aligned} L_i(\mathbf{x}, \ell^{(K)}) + \langle \nabla L_i(\mathbf{x}, \ell^{(K)}), \ell - \ell^{(K)} \rangle, \\ i = j, \dots, k - 1, (j, k) \in \mathcal{I}_B, B = 0, \dots, B_{\max}, \end{aligned}$$

on iteration K , can help circumvent numerical issues. Furthermore, as all the constraints in the problems (11) and (12) are now evidently affine functions, this relaxation has the added benefit of turning the problems (11) and (12) into linear programs, for which there (of course) exist heavily optimized solvers.

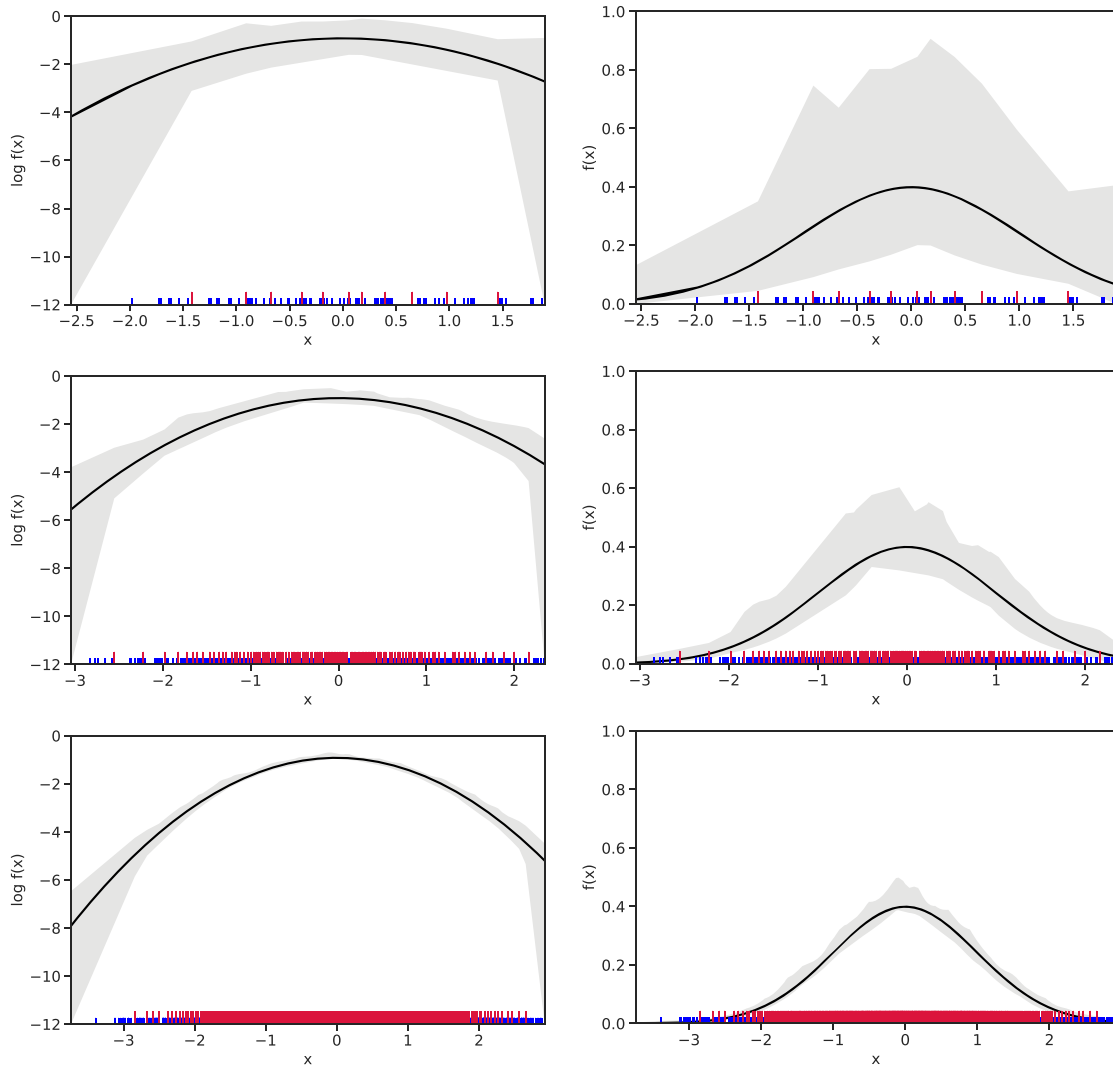


Figure 1. Confidence bands (shaded in gray) generated by Algorithm 1 and linear interpolation as described in Section 2.4, for a Gaussian density. The left column shows the bands for the log density, while the right column shows the bands for the density. The solid black line marks the underlying (log)density. Top, middle and bottom rows show results for sample sizes $n = 100, 1000,$ and $10,000,$ respectively. At the bottom of each plot, the observations X_i are indicated in blue (short lines), while the points $x_i, i = 1, \dots, m$ are marked in red (long lines).

Table 1. Empirical coverages, average widths at the sample quartiles (denoted $Q_1, Q_2,$ and Q_3), and average runtimes in seconds, for the confidence bands generated by Algorithm 1 with linear interpolation and nominal coverage levels 90% and 95%, when applied to $n \in \{100, 1000\}$ observations drawn from four different underlying densities (Gaussian, uniform, Chi-squared, and exponential).

Distribution	n	Nominal coverage 90%			Nominal coverage 95%			Time (secs.)		
		Coverage	Width		Coverage	Width				
			Q_1	Q_2		Q_3	Q_1		Q_2	Q_3
Gaussian	100	0.96	0.58	0.65	0.61	0.98	0.66	0.74	0.70	5.6
	1000	0.95	0.24	0.28	0.24	0.97	0.27	0.31	0.27	151.2
Uniform	100	0.93	0.08	0.07	0.08	0.96	0.09	0.08	0.09	5.1
	1000	0.91	0.03	0.03	0.03	0.97	0.03	0.03	0.03	128.8
Chi-squared	100	0.94	0.36	0.28	0.19	0.98	0.41	0.33	0.23	5.4
	1000	0.94	0.16	0.12	0.07	0.97	0.18	0.13	0.08	132.5
Exponential	100	0.93	0.94	0.69	0.44	0.96	1.06	0.77	0.51	5.7
	1000	0.91	0.38	0.25	0.15	0.97	0.43	0.29	0.17	140.4

NOTE: All results are based on 1000 simulations. The runtimes for the nominal coverage of 90% are similar to the case of 95% and are not displayed.

4. Large Sample Statistical Performance

The large sample performance of the log-concave MLE has been studied intensively; see, for example, Dümbgen and Rufibach (2009), Kim and Samworth (2016), and Doss and Wellner (2016). The main message is that the MLE attains the optimal minimax rate of convergence of $O(n^{-2/5})$ with respect to various global loss functions. Recently, Kim, Guntuboyina, and Samworth (2018) have shown that the MLE can achieve a faster rate of convergence when the log density is k -affine, that is, when $\log f$ consists of k linear pieces. They show that the MLE is able to adapt to this simpler model, where it will converge with nearly the parametric rate, namely with $O(n^{-1/2} \log^{5/8} n)$. Here we show that the construction of our confidence band via the particular confidence set $C_n(\alpha)$ will also result in a nearly parametric rate of convergence for the width of the confidence band in that case. To this end, we first consider the case where some part of f is log-linear:

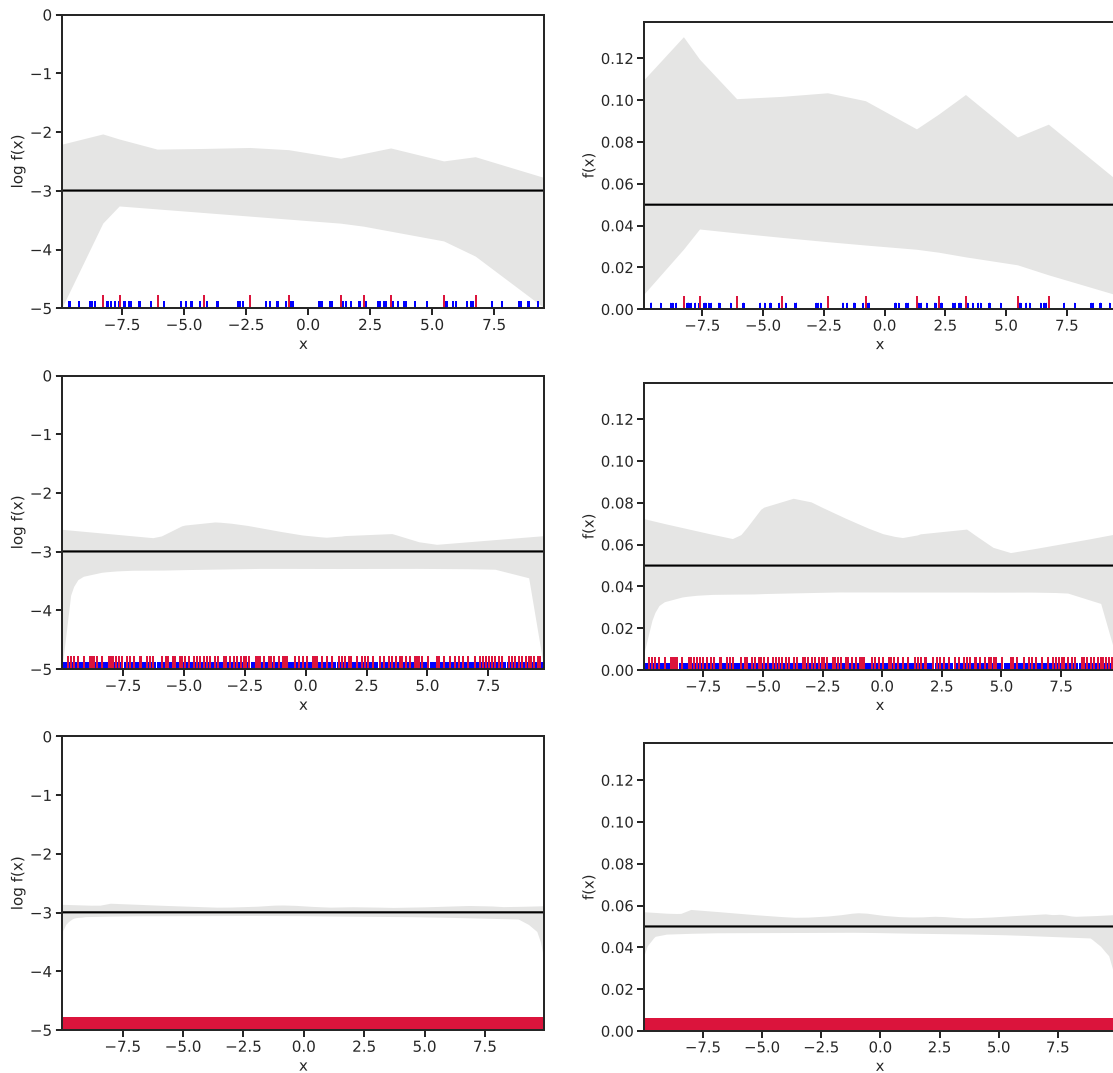


Figure 2. Confidence bands (shaded in gray) generated by Algorithm 1 and linear interpolation as described in Section 2.4, for a uniform density. The left column shows the bands for the log density, while the right column shows the bands for the density. The solid black line marks the underlying (log)density. Top, middle and bottom rows show results for sample sizes $n = 100, 1000,$ and $10,000,$ respectively. At the bottom of each plot, the observations X_i are indicated in blue (short lines), while the points $x_i, i = 1, \dots, m$ are marked in red (long lines).

Theorem 1. Let f be a log-concave density and suppose $\log f$ is linear on some interval J . (So J may be a proper subset of the support of f). Then on every closed interval $I \subset \text{int} J$:

$$\mathbb{P}_f \left\{ \sup_{x \in I} \left| \hat{u}(x) - f(x) \right| \leq C \sqrt{\frac{\log \log n}{n}} \right\} \rightarrow 1$$

for some constant C , and the same statement holds for $\hat{\ell}$ in place of \hat{u} . In particular, the width of the confidence band satisfies $\max_{x \in I} (\hat{u}(x) - \hat{\ell}(x)) \leq 2C \sqrt{\frac{\log \log n}{n}}$ with probability converging to 1.

If there are k such intervals, then the theorem holds for the maximum width over the k intervals. This includes k -affine log densities as a special case. The proof of Theorem 1 is given in the Appendix, supplementary materials.

We conjecture that the width of the confidence band will likewise achieve the optimal minimax rate if $\log f$ is smooth rather than linear.

5. Some Examples

Finally, we present some numerical examples of our methodology, highlighting the empirical coverage and widths of our confidence bands as well as the computational cost of computing the bands, for a number of different distributions.

To calculate coverage, we first simulated $n \in \{100, 1000\}$ observations from a (i) standard normal distribution, (ii) uniform distribution on $[-10, 10]$, (iii) chi-squared distribution with three degrees of freedom, and (iv) exponential distribution with parameter 1. Then, we computed our confidence bands from the data by running the penalty convex-concave procedure described in Algorithm 1 and then computing $\exp(\hat{\ell}(x))$ and $\exp(\hat{u}(x))$ as discussed in Section 2.4, where $\exp(\hat{u}(x))$ was

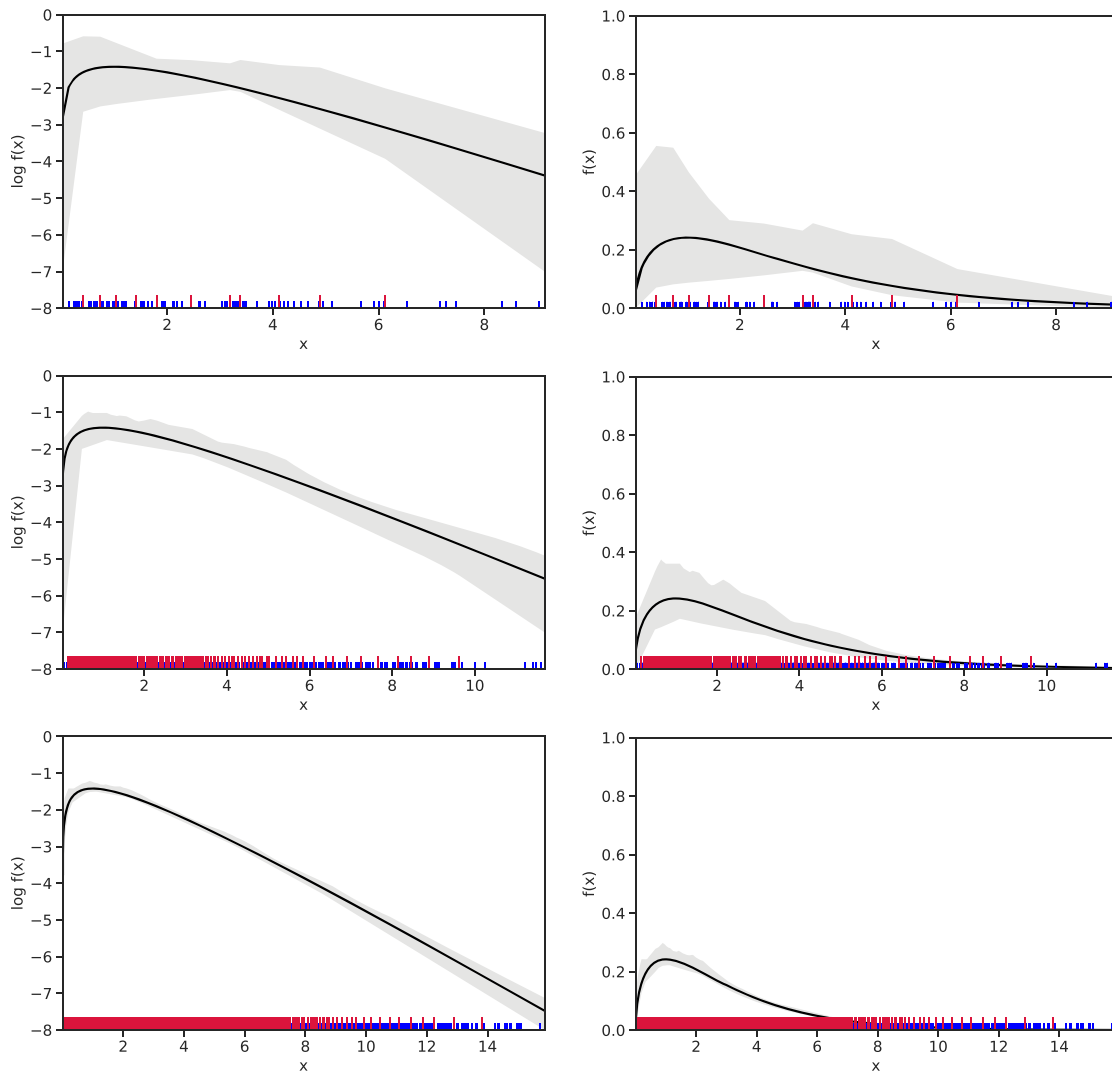


Figure 3. Confidence bands (shaded in gray) generated by Algorithm 1 and linear interpolation as described in Section 2.4, for a chi-squared density. The left column shows the bands for the log density, while the right column shows the bands for the density. The solid black line marks the underlying (log)density. Top, middle, and bottom rows show results for sample sizes $n = 100$, 1000, and 10,000, respectively. At the bottom of each plot, the observations X_i are indicated in blue (short lines), while the points x_i , $i = 1, \dots, m$ are marked in red (long lines).

computed by linearly interpolating between the $(x_i, \exp(\hat{\mu}(x_i)))$. We repeated these two steps (simulating data and computing bands) 1000 times. We calculated the empirical coverage for each density f by evaluating the band at 10,000 points $\{t_j\}$, evenly spaced across the range of the data, to check whether $\exp(\hat{\ell}(t_j)) \leq f(t_j) \leq \exp(\hat{\mu}(t_j))$ for all j , and then computed the empirical frequency of this event across the 1000 repetitions. In order to calculate the widths of the bands, we averaged the widths at the sample quartiles over all of the repetitions. We calculated the computational effort by averaging the runtimes, obtained by running Algorithm 1 on a workstation with four Intel E5-4620 2.20GHz processors and 15 GB of RAM, over all the repetitions. To speed up the computation for $n \geq 1000$, we ran Algorithm 1 on a subset of 30% of the points x_i , $i = 1, \dots, m$, as discussed at the end of Section 2.4; the coverages and widths were virtually indistinguishable from those obtained using the full set of points x_i , $i = 1, \dots, m$.

Algorithm 1 requires a few tuning parameters, which are important for assuring quick convergence. In general, we found

that the initial penalty strength τ_0 , the maximum penalty strength τ_{\max} , and the penalty growth factor κ had the greatest impact on convergence. In our experience, setting τ_0 to a small value and τ_{\max} to a large value worked well; we used $\tau_0 \in \{10^{-5}, 10^{-4}, 10^{-3}\}$ and $\tau_{\max} \in \{10^3, 10^4, 10^5\}$, with the most suitable values depending on the characteristics of the problem. We experimented with various penalty growth factors $\kappa \in \{1, 2, \dots, 10\}$, finding that the best value of κ again varied with the problem setup. We always set the maximum number of iterations $K_{\max} = 50$ as our method usually converged after around 20–30 iterations across all problem settings. We initialized the points $\ell^{(0)}, \mathbf{g}^{(0)}$ randomly. Finally, we set the miscoverage level $\alpha = 0.1$ but we also report results for $\alpha = 0.05$.

Table 1 summarizes the results. The table (reassuringly) shows us that the bands achieve coverage at or above the nominal level. In Figures 1–5 we present a visualization of the bands, from a single repetition chosen at random, for each of the four underlying densities as well as for a density that is




Figure 4. Confidence bands (shaded in gray) generated by Algorithm 1 and linear interpolation as described in Section 2.4, for an exponential density. The left column shows the bands for the log density, while the right column shows the bands for the density. The solid black line marks the underlying (log)density. Top, middle and bottom rows show results for sample sizes $n = 100, 1000,$ and $10,000,$ respectively. At the bottom of each plot, the observations X_i are indicated in blue (short lines), while the points $x_i, i = 1, \dots, m$ are marked in red (long lines).

proportional to $\exp(-x^4)$. In addition, we depict the bands for the case of a larger sample with $n = 10,000$. The figures and Table 1 show that while the bands are naturally wider when the sample size is small ($n = 100$), they quickly tighten as the sample size grows ($n \in \{1000, 10,000\}$).

As for the computational cost, we found that around 20–30 iterations of Algorithm 1 were enough to reach convergence, for each point $x_i, i = 1, \dots, m$. Table 1 shows that this translates into just a few seconds to compute the entire band when $n = 100$, and a couple of minutes when $n = 1000$. We found that Algorithm 1 converged to the exact same solutions even when started from a number of different initial points, suggesting that it is in fact finding the global minimizers of the problems (9). Therefore, these runtimes appear to be reasonable, as it is worth bearing in mind that Algorithm 1 is effectively solving a potentially large number of nonconvex optimization problems (precisely: 13, 39, and 188 such problems, corresponding to $n \in \{100, 1000, 10,000\}$, respectively). Moreover, we point out

that the computation in Algorithm 1 can easily be parallelized, for example, across the points $x_i, i = 1, \dots, m$.

In order to compare the confidence band with pointwise confidence intervals, we performed these experiments also with two methods that compute pointwise confidence intervals for log-concave densities. Azadbakhsh, Jankowski, and Gao (2014) compare several such methods and report that no one approach appears to uniformly dominate the others and that each method works well only in a certain range of the data. The first group of methods examined by Azadbakhsh, Jankowski, and Gao (2014) is based on the pointwise asymptotic theory developed in Balabdaoui, Rufibach, and Wellner (2009) and requires estimating a nuisance parameter, for which Azadbakhsh, Jankowski, and Gao (2014) investigate several options. We picked the option that they report works best, namely their method (iv) in their Section 4. This method is called “Asymptotic theory with approximation” in Table 2. The second group of methods analyzed by Azadbakhsh, Jankowski, and Gao (2014) concerns various

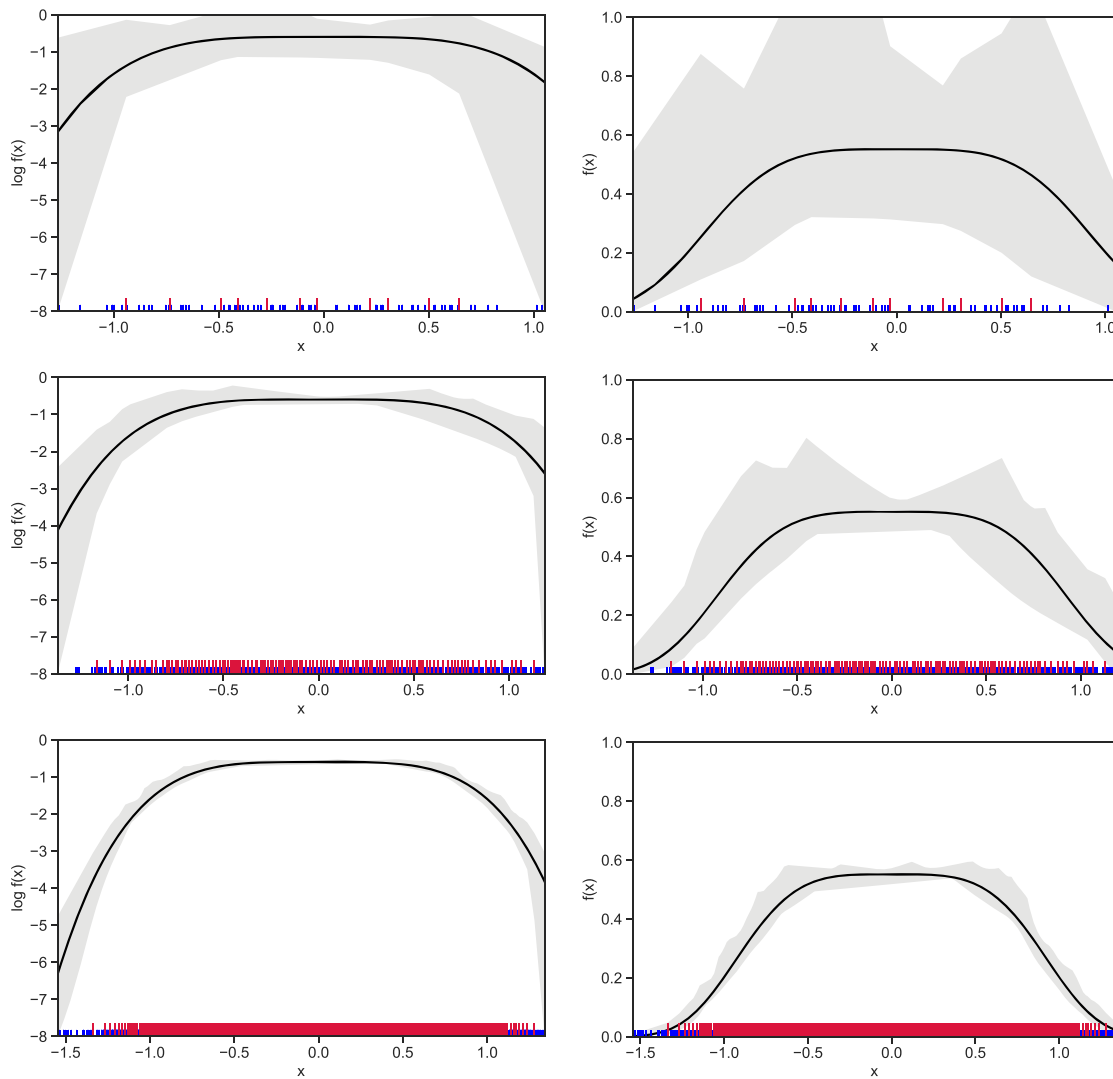


Figure 5. Confidence bands (shaded in gray) generated by Algorithm 1 and linear interpolation as described in Section 2.4, for a density proportional to $\exp(-x^4)$. The left column shows the bands for the log density, while the right column shows the bands for the density. The solid black line marks the underlying (log)density. Top, middle and bottom rows show results for sample sizes $n = 100, 1000,$ and $10,000,$ respectively. At the bottom of each plot, the observations X_i are indicated in blue (short lines), while the points $x_i, i = 1, \dots, m$ are marked in red (long lines).

Table 2. Empirical coverages, average widths at the sample quartiles (denoted $Q_1, Q_2,$ and Q_3), and average runtimes in seconds, for the confidence bands obtained from pointwise asymptotic theory and approximation (iv) in Azadbakhsh, Jankowski, and Gao (2014) and by the bootstrap, with nominal level 90% in the same settings as in Table 1.

Distribution	n	Asymptotic theory with approximation				Time (secs.)	Bootstrap				Time (secs.)
		Coverage	Width				Coverage	Width			
			Q_1	Q_2	Q_3			Q_1	Q_2	Q_3	
Gaussian	100	0.22	0.13	0.20	0.13	1.4	0.37	0.16	0.16	0.12	98.2
	1000	0.08	0.04	0.08	0.04	6.4	0.15	0.04	0.07	0.03	194.8
Uniform	100	0.32	0.03	0.03	0.03	1.4	0.01	0.02	0.01	0.02	97.4
	1000	0.15	0.01	0.01	0.01	6.5	0.11	0.01	0.01	0.01	184.5
Chi-squared	100	0.06	0.08	0.04	0.02	1.4	0.52	0.04	0.03	0.02	97.2
	1000	0.00	0.02	0.01	0.00	6.6	0.31	0.02	0.01	0.01	191.0
Exponential	100	0.37	0.04	0.02	0.01	1.4	0.05	0.07	0.05	0.03	95.7
	1000	0.15	0.01	0.01	0.01	6.5	0.04	0.02	0.01	0.01	180.0

bootstrapping schemes, and we chose the one that they report to have the best performance, namely the ECDF-bootstrap, listed as (v) in their Section 4, which we use with 250 bootstrap repetitions. This method computes the MLE for each bootstrap sample and then computes the bootstrap percentile interval at a

point x_0 based on the 250 bootstrap replicates of the MLE at x_0 . We used the R function `logConCI` produced by Azadbakhsh, Jankowski, and Gao (2014) for implementing both methods. With each of the two methods we compute the pointwise 90% confidence interval for each point x_0 in the grid of points that

we use to evaluate empirical coverage. Since these are *pointwise* 90% confidence intervals, we expect that the coverage for the band (i.e., the simultaneous coverage across all x_0 in the grid) is smaller than 90%, but that the intervals are narrower than those for a simultaneous confidence band. This is confirmed by the results in Table 2, which show that both methods seriously undercover. The bootstrap is also seen to be significantly more computationally intensive than the method based on asymptotic theory as well as Algorithm 1.

6. Discussion

The article shows how to construct confidence bands for a log-concave density by intersecting the log-concavity constraint with an appropriate nonparametric confidence set. This approach has three strong points: First, it produces confidence bands with a finite sample guaranteed confidence level. Our simulations have shown that this guaranteed confidence level is not overly conservative. Second, the approach allows to bring modern tools from optimization to bear on this problem. This aspect is particularly important in a multivariate setting where it is known that computing the MLE is very time consuming. We expect that the key ideas of the univariate construction introduced here can be carried over to the multivariate setting and we are working on implementing this program in the multivariate setting. Third, it was shown that this approach results in attractive statistical properties, namely that the confidence bands converge at nearly the parametric $n^{-\frac{1}{2}}$ rate when the log density is k -affine. We conjecture that the width of these confidence bands will likewise achieve the optimal minimax rate if $\log f$ is smooth rather than piecewise linear, and we leave the proof of such a result as an open problem.

Supplementary Materials

Appendix.pdf contains the mathematical proofs for Lemma 1 and of Theorem 1, and log_ccv_conf_int-master.zip contains the code. The code is also available at github https://github.com/cvxgrp/log_ccv_conf_int.

Funding

Guenter Walther—research supported by NSF grants DMS-1501767 and DMS-1916074. Alnur Ali—research supported by the Intelligence Community Postdoctoral Research Fellowship Program. Xinyue Shen—research supported by the National Key R&D Program of China with grant no. 2018YFB1800800, the Key Area R&D Program of Guangdong Province with grant no. 2018B030338001, the Shenzhen Outstanding Talents Training Fund, and the Guangdong research project no. 2017ZT07X152.

References

Azadbakhsh, M., Jankowski, H., and Gao, X. (2014), “Computing Confidence Intervals for Log-Concave Densities,” *Computational Statistics & Data Analysis*, 75, 248–264. [1427,1435,1436]

Balabdaoui, F., Rufibach, K., and Wellner, J. A. (2009), “Limit Distribution Theory for Maximum Likelihood Estimation of a Log-Concave Density,” *The Annals of Statistics*, 37, 1299–1331. [1426,1427,1435]

Birgé, L. (1997), “Estimation of Unimodal Densities Without Smoothness Assumptions,” *The Annals of Statistics*, 25, 970–981. [1426]

Cule, M., Samworth, R., and Stewart, M. (2010), “Maximum Likelihood Estimation of a Multi-Dimensional Log-Concave Density,” *Journal of the Royal Statistical Society, Series B*, 72, 545–607. [1427]

Davies, P. L., and Kovac, A. (2004), “Densities, Spectral Densities and Modality,” *The Annals of Statistics*, 32, 1093–1136. [1427]

Deng, H., Han, Q., and Sen, B. (2020), “Inference for Local Parameters in Convexity Constrained Models,” arXiv preprint arXiv:2006.10264. [1427]

Dinh, T. P., and Le Thi, H. A. (2014), “Recent Advances in DC Programming and DCA,” in *Transactions on Computational Intelligence XIII*, eds. N.-T. Nguyen, and H. A. Le-Thi, 1–37, Berlin: Springer. [1430]

Donoho, D. L. (1988), “One-Sided Inference for Functionals of a Density,” *The Annals of Statistics*, 16, 1390–1420. [1427]

Doss, C. R., and Wellner, J. A. (2016), “Global Rates of Convergence of the MLEs of Log-Concave and s-concave Densities,” *The Annals of Statistics*, 44, 954–981. [1426,1432]

Dümbgen, L. (1998), “New Goodness-of-Fit Tests and their Application to Nonparametric Confidence Sets,” *The Annals of Statistics*, 26, 288–314. [1427]

——— (2003), “Optimal Confidence Bands for Shape-Restricted Curves,” *Bernoulli*, 9, 423–449. [1427]

Dümbgen, L., Hüsler, A., and Rufibach, K. (2007), “Active Set and EM Algorithms for Log-Concave Densities Based on Complete and Censored Data,” arXiv preprint arXiv:0707.4643v4. [1427]

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. [1426,1432]

Feng, O., Guntuboyina, A., Kim, A. K. H., and Samworth, R. J. (2021), “Adaptation in Multivariate Log-Concave Density Estimation,” *The Annals of Statistics*, 49, 129–153. [1427]

Hartman, P. (1959), “On Functions Representable as a Difference of Convex Functions,” *Pacific Journal of Mathematics*, 9, 707–713. [1430]

Hengartner, N. W., and Stark, P. B. (1995), “Finite-Sample Confidence Envelopes for Shaperestricted Densities,” *The Annals of Statistics*, 23, 525–550. [1427]

Horst, R., Pardalos, P. M., and Van Thoai, N. (2000), *Introduction to Global Optimization*, New York: Springer. [1430]

Horst, R., and Thoai, N. V. (1999), “DC Programming: Overview,” *Journal of Optimization Theory and Applications*, 103, 1–43. [1430]

Khamaru, K., and Wainwright, M. J. (2018), “Convergence Guarantees for a Class of Non-convex and Non-smooth Optimization Problems,” arXiv preprint arXiv:1804.09629. [1431]

Kim, A. K. H., and Samworth, R. J. (2016), “Global Rates of Convergence in Log-Concave Density Estimation,” *The Annals of Statistics*, 44, 2756–2779. [1426,1427,1432]

Kim, A. K. H., Guntuboyina, A., and Samworth, R. J. (2018), “Adaptation in Log-Concave Density Estimation,” *The Annals of Statistics*, 46, 2279–2306. [1426,1432]

Le Thi, H. A., and Dinh, T. P. (2014), “DC Programming and DCA for General DC Programs,” in *Advanced Computational Methods for Knowledge Engineering*, eds. N. T. Nguyen, T. Do, H. A. Thi, 15–35, Cham: Springer. [1430]

Li, H., Munk, A., Sieling, H., and Walther, G. (2020), “The Essential Histogram,” *Biometrika*, 107, 347–364. [1428]

Lipp, T., and Boyd, S. (2016), “Variations and Extension of the Convex-Concave Procedure,” *Optimization and Engineering*, 17, 263–287. [1430,1431]

Liu, Y., and Wang, Y. (2018), “A Fast Algorithm for Univariate Log-Concave Density Estimation,” *Australian and New Zealand Journal of Statistics*, 60, 258–275. [1427]

Pal, J. K., Woodroffe, M., and Meyer, M. (2007), “Estimating a Polya Frequency Function,” in *Complex Datasets and Inverse Problems. IMS Lecture Notes Monograph Series*, 54, 239–249, Beachwood, OH; Institute of Mathematical Statistics. [1426,1427]

Rivera, C., and Walther, G. (2013), “Optimal Detection of a Jump in the Intensity of a Poisson Process or in a Density with Likelihood Ratio Statistics,” *Scandinavian Journal of Statistics*, 40, 752–769. [1428]

Rufibach, K. (2007), “Computing Maximum Likelihood Estimators of a Log-Concave Density Function,” *Journal of Statistical Computation and Simulation*, 77, 561–574. [1427]

Samworth, R. J. (2018), “Recent Progress in Log-Concave Density Estimation,” *Statistical Science*, 33, 493–509. [1427]

- Saumard, A., and Wellner, J. A. (2014), “Log-Concavity and Strong Log-Concavity: A Review,” *Statistics Surveys*, 8, 45–114. [1427]
- Schuhmacher, D., and Dümbgen, L. (2010), “Consistency of Multivariate Log-Concave Density Estimators,” *Statistics & Probability Letters*, 80, 376–380. [1427]
- Seregin, A., and Wellner, J. A. (2010), “Nonparametric Estimation of Multivariate Convex-Transformed Densities,” *The Annals of Statistics*, 38, 3751–3781. [1427]
- Shorack, G. R., and Wellner, J. A. (1986), *Empirical Processes with Applications to Statistics*, New York: Wiley. [1428]
- Smola, A. J., Vishwanathan, S. V. N., and Hofmann, T. (2005), “Kernel Methods for Missing Variables,” in *AISTATS*. [1430]
- Sriperumbudur, B. K., and Lanckriet, G. R. (2009), “On the Convergence of the Concave-Convex Procedure,” in *Proceedings of the 22nd International Conference on Neural Information Processing Systems*, 1759–1767. Curran Associates Inc. [1431]
- Tao, P. D. (1986), “Algorithms for Solving a Class of Nonconvex Optimization Problems. Methods of Subgradients,” *North-Holland Mathematics Studies*, 129, 249–271. [1430]
- Walther, G. (2002), “Detecting the Presence of Mixing with Multiscale Maximum Likelihood,” *Journal of the American Statistical Association*, 97, 508–513. [1426,1427]
- Walther, G. (2009), “Inference and Modeling with Log-Concave Distributions,” *Statistical Science*, 24, 319–327. [1427]
- Walther, G., and Perry, A. (2019), “Calibrating the Scan Statistic: Finite Sample Performance vs. Asymptotics,” arXiv preprint arXiv:2008.06136. [1427,1428]
- Yuille, A. L., and Rangarajan, A. (2003), “The Concave-Convex Procedure,” *Neural Computation*, 15, 915–936. [1430]