

Inference and Modeling with Log-concave Distributions

Guenther Walther

Abstract. Log-concave distributions are an attractive choice for modeling and inference, for several reasons: The class of log-concave distributions contains most of the commonly used parametric distributions and thus is a rich and flexible nonparametric class of distributions. Further, the MLE exists and can be computed with readily available algorithms. Thus, no tuning parameter, such as a bandwidth, is necessary for estimation. Due to these attractive properties, there has been considerable recent research activity concerning the theory and applications of log-concave distributions. This article gives a review of these results.

Key words and phrases: Nonparametric density estimation, shape constraint, log-concave density, Polya frequency function, strongly unimodal, iterative convex minorant algorithm, active set algorithm.

1. INTRODUCTION

There has been considerable recent activity in the area of inference under shape constraints, that is, inference about a (say) function f under the constraint that f satisfies certain qualitative properties, such as monotonicity or convexity on certain subsets of its domain. This approach is appealing for two main reasons: First, such shape constraints are sometimes direct consequences of the problem under investigation (see, e.g., Hampel, 1987, or Wang et al., 2005), or they are at least plausible in many problems. It is then desirable that the result of the inference reflect this fact. There is also the hope that imposing these constraints will improve the quality of the resulting estimator in some sense. The second reason is that alternative nonparametric estimators such as, for example, kernel estimators, typically require the choice of a tuning parameter such as a bandwidth. A good choice for such a tuning parameter is usually far from trivial and injects a certain amount of subjectivity into the estimator. In contrast, inference under shape constraints often results in an explicit solution that does not depend on a tuning parameter.

In the context of density estimation, Grenander (1956) derived the nonparametric maximum likelihood

estimator of a density function that is nonincreasing on a half-line. This estimator is given explicitly by the left derivative of the least concave majorant of the empirical distribution function. However, this result does not carry over to the problem of estimating a unimodal density with unknown mode, as then the nonparametric MLE does not exist; see, for example, Birgé (1997). Even if the mode is known, the estimator suffers from inconsistency near the mode, the so-called spiking problem; see, for example, Woodroofe and Sun (1993). These results are unfortunate since the constraint of unimodality is cited as a reasonable assumption in many problems.

It was argued in Walther (2002) that log-concave densities are an attractive and natural alternative choice to the class of unimodal densities: The class of log-concave densities is a subset of the class of the unimodal densities, but it contains most of the commonly used parametric distributions and is thus a rich and useful nonparametric model. Moreover, it was shown in Walther (2002) that the nonparametric MLE of a univariate log-concave density exists and can be computed with readily available algorithms.

Due to these attractive properties, there has been considerable recent research activity about the statistical properties of the MLE, computational aspects, applications in modeling and inference, as well as about the multivariate case. As an example, Figure 1 shows a scatterplot of measurements on 569 individuals from

Guenther Walther is Associate Professor, Department of Statistics, Stanford University, 390 Serra Mall, Stanford, California 94305, USA (e-mail: gwalther@stanford.edu).

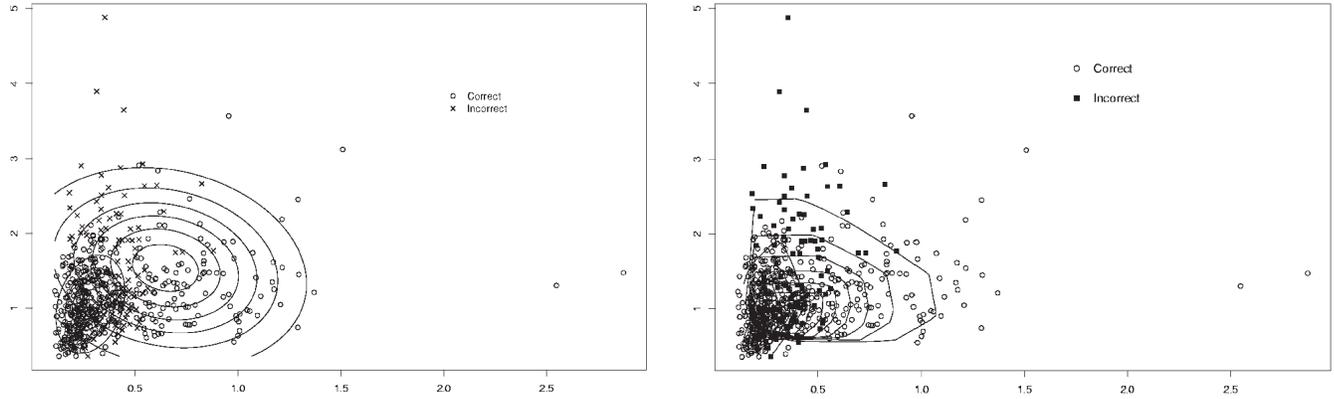


FIG. 1. Contour plot and misclassified instances from the Gaussian EM algorithm (left) and the log-concave EM algorithm (right). The plots are from Cule, Samworth and Stewart (2008).

the Wisconsin breast cancer data set; see Section 6 for a more detailed description. The data were clustered using a two-component normal mixture model fitted with the EM-algorithm; see, for example, Fraley and Raftery (2002). The contour lines of the fitted normal components are shown in the left plot, while the right plot shows the contour lines that obtain when the normal MLE is replaced by the log-concave MLE in the EM algorithm. The log-concave MLE automatically adapts to the multivariate skewness of the data and results in a superior clustering: Each observation is either a benign or a malignant instance. These labels were not used for the fitting but can be employed to assess the quality of the clustering. The EM algorithm with the log-concave MLE resulted in 121 misclassified instances versus 144 for the Gaussian MLE.

This article gives an overview of recent results about inference and modeling with the log-concave MLE. Section 2 gives some basic properties and applications of log-concave distributions. Section 3 addresses the MLE and its statistical properties. Computational aspects are surveyed in Section 4, while Section 5 describes recent advances in the multivariate setting. Section 6 reviews applications of the log-concave MLE for various modeling and inference problems. Section 7 lists some open problems for future work.

2. BASIC PROPERTIES AND APPLICATIONS OF LOG-CONCAVE FUNCTIONS

A function f on \mathbf{R}^d is log-concave if it is of the form

$$(1) \quad f(x) = \exp \phi(x),$$

for some concave function $\phi: \mathbf{R}^d \rightarrow [-\infty, \infty)$. A prime example is the normal density, where $\phi(x)$ is a

quadratic in x . Further, most common univariate parametric densities are log-concave, such as the normal family, all gamma densities with shape parameter ≥ 1 , all Weibull densities with exponent ≥ 1 , all beta densities with both parameters ≥ 1 , the generalized Pareto and the logistic density; see, for example, Marshall and Olkin (1979).

Log-concave functions have a number of properties that are desirable for modeling: Marginal distributions, convolutions and product measures of log-concave distributions are again log-concave; see, for example, Dharmadhikari and Joag-Dev (1988). Notably, the first two properties are not true for the class of unimodal densities.¹ Log-concave distributions may be skewed, and this flexibility is relevant in a number of applications; see, for example, Section 6. On the other hand, log-concave distributions necessarily have subexponential tails and nondecreasing hazard rates; see, for example, Karlin (1968) and Barlow and Proschan (1975).

There are several alternative characterizations and designations for the class of univariate log-concave distributions: Ibragimov (1956) proved that these are precisely the distributions whose convolution with a unimodal distribution is always unimodal; thus, log-concave distributions are sometimes referred to as strongly unimodal. Log-concave densities are also precisely the Polya frequency functions of order 2, as well as precisely those densities f for which the location family $f_\theta(x) := f(x - \theta)$ has monotone likelihood ratio in x ; see Karlin (1968).

Log-concave distribution models have been found useful in economics (see, e.g., An, 1995, 1998; Bagnoli

¹Counterexamples are available from the author upon request.

and Bergstrom, 2005 and Caplin and Nalebuff, 1991), in reliability theory (see, e.g., Barlow and Proschan, 1975) and in sampling and nonparametric Bayesian analysis (see, e.g., Gilks and Wild, 1992; Dellaportas and Smith, 1993 and Brooks, 1998). Recent advances in inference have led to fruitful applications of log-concave distributions in other areas such as clustering, some of which will be discussed in Section 6.

3. PROPERTIES OF THE NONPARAMETRIC MLE

If X_1, \dots, X_n are i.i.d. observations from a univariate log-concave density (1), then the nonparametric MLE exists, is unique, and is of the form $\hat{f}_n = \exp \hat{\phi}_n$, where $\hat{\phi}_n$ is continuous and piecewise linear on $[X_{(1)}, X_{(n)}]$ with the set of knots contained in $\{X_1, \dots, X_n\}$, and $\hat{\phi}_n = -\infty$ on $\mathbf{R} \setminus [X_{(1)}, X_{(n)}]$; see

Walther (2002), Rufibach (2006) or Pal, Woodroffe and Meyer (2007). An example is plotted in Figure 2.

Consistency of \hat{f}_n with respect to the Hellinger metric was established in Pal, Woodroffe and Meyer (2007), while Dümbgen and Rufibach (2009) provide results on the uniform consistency on compact subsets of the interior of the support: If ϕ belongs to a Hölder class with exponent $\beta \in [1, 2]$, then $\hat{\phi}_n$ and \hat{f}_n are uniformly consistent with rate $O_p((\log n/n)^{\beta/(2\beta+1)})$. Thus, in the typical case $\beta = 2$, \hat{f}_n converges uniformly with rate $O_p((\log n/n)^{2/5})$. It is known that these rates are optimal even if β were known. This establishes that the nonparametric MLE adapts to the unknown local smoothness of f , at least for $\beta \in [1, 2]$. Further, under some regularity conditions, the c.d.f. \hat{F}_n of \hat{f}_n is asymptotically equivalent to the empirical c.d.f. \mathbb{F}_n : If $\beta > 1$, then $|\mathbb{F}_n - \hat{F}_n|$ is of order $o_p(n^{-1/2})$ uniformly

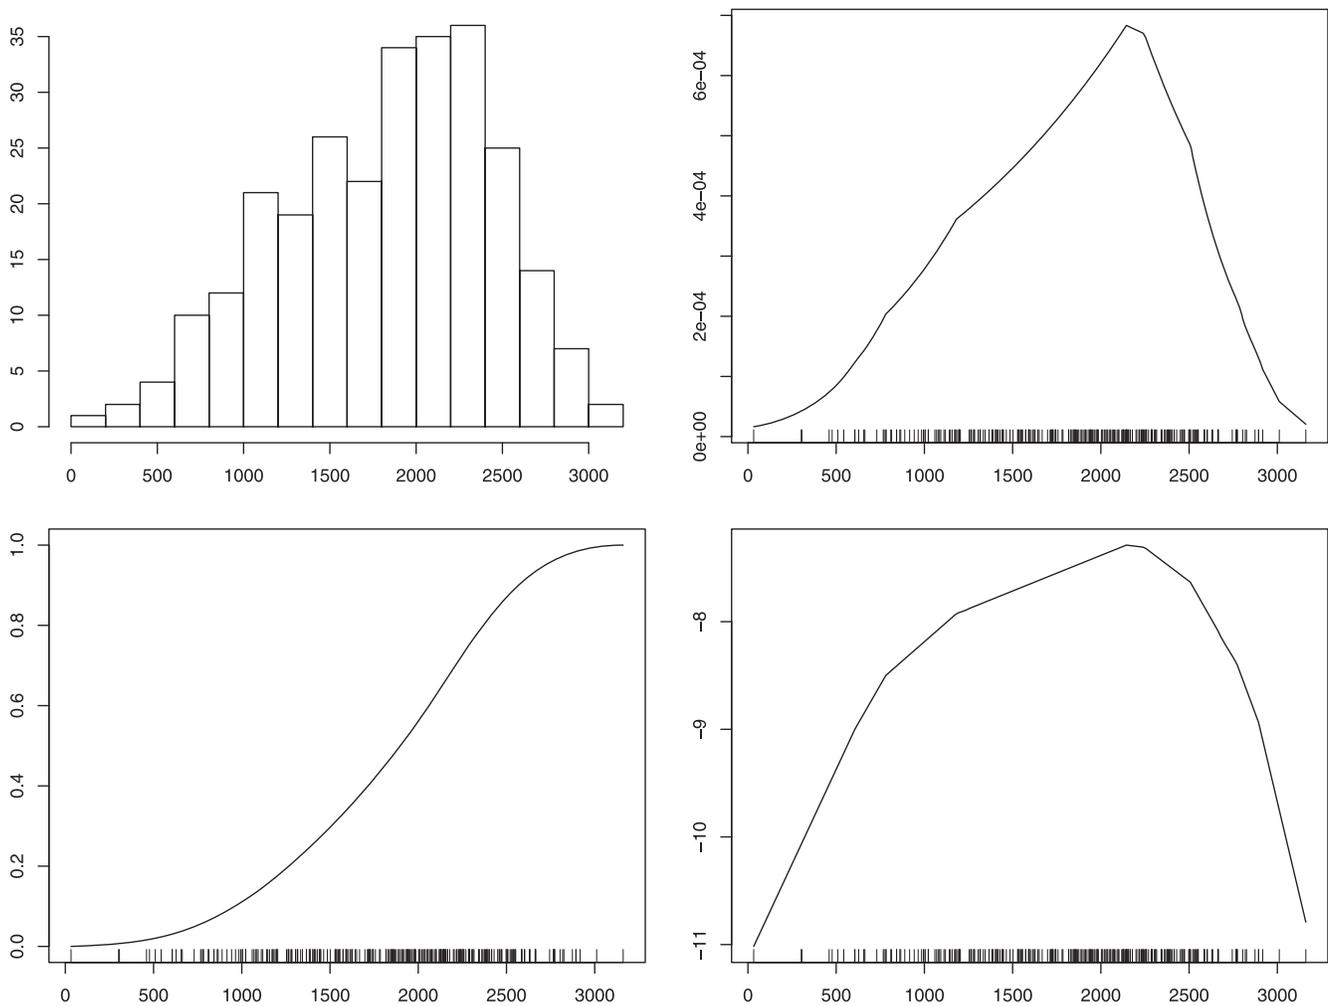


FIG. 2. The histogram of $n = 270$ flow cytometry data (top left), the log-concave MLE \hat{f}_n (top right), the estimated c.d.f. (bottom left), and $\hat{\phi}_n = \log \hat{f}_n$ (bottom right).

over compact subsets of the interior of the support. Moreover, $\mathbb{F}_n - n^{-1} \leq \hat{F}_n \leq \mathbb{F}_n$ on the set of knots of $\hat{\phi}_n$. The resulting uniform \sqrt{n} -consistency of \hat{F}_n outperforms, for example, c.d.f.s of kernel estimators using a nonnegative kernel with optimally chosen bandwidth. While empirical evidence suggests that \hat{f}_n performs well over the whole line, establishing the corresponding theoretical results is still an open problem.

Balabdaoui, Rufibach and Wellner (2009) derive the pointwise limiting distributions of $n^{k/(2k+1)}(\hat{f}_n(x_0) - f(x_0))$, $n^{(k-1)/(2k+1)}(\hat{f}'_n(x_0) - f'(x_0))$, and likewise for $\hat{\phi}_n$ and $\hat{\phi}'_n$, where k is the smallest integer such that $\phi^{(k)}(x_0) \neq 0$. They show that these limiting distributions depend on the “lower envelope” of an integrated Brownian motion process minus a drift term that depends on k .

4. COMPUTATIONAL ASPECTS

Maximizing the log-likelihood function under the constraint $\int \exp \phi(x) dx = 1$ is equivalent to maximizing $\sum_{i=1}^n \phi(X_i) - n \int \exp \phi(x) dx$ over the set of all concave functions ϕ ; see Silverman (1982). Due to the piecewise linear form of the solution $\hat{\phi}$, one can write this as a finite-dimensional optimization problem as follows: For the ordered data $x_1 < \dots < x_n$ write $\phi_1 := \phi(x_1)$ and denote the slope between x_{i-1} and x_i by $s_i := (\phi(x_i) - \phi(x_{i-1})) / (x_i - x_{i-1})$, $i = 2, \dots, n$. Then the optimization problem is to maximize

$$\begin{aligned} \Psi_n(\phi_1, s_2, \dots, s_n) &= n\phi_1 + \sum_{i=2}^n (n-i+1)(x_i - x_{i-1})s_i \\ &\quad - n \exp(\phi_1) \sum_{i=2}^n \left(\exp\left(\sum_{k=2}^i (x_k - x_{k-1})s_k\right) \right. \\ &\quad \left. - \exp\left(\sum_{k=2}^{i-1} (x_k - x_{k-1})s_k\right) \right) / s_i \end{aligned}$$

under the constraint that the vector $(\phi_1, s_2, \dots, s_n)$ belongs to the cone $\mathcal{C}_n := \{\mathbf{y} \in \mathbf{R}^n : y_2 \geq \dots \geq y_n\}$. Ψ_n is a concave function on \mathbf{R}^n which needs to be maximized over the convex cone \mathcal{C}_n . This is precisely the type of problem for which the Iterative Convex Minorant Algorithm (ICMA) was developed; see Groeneboom and Wellner (1992) and Jongbloed (1998). The key idea of that algorithm is to approximate the concave function locally around the current candidate solution by a quadratic form, which is then maximized by a Newton procedure over the cone by using the pool-adjacent-violators algorithm. This procedure is then iterated to

the final solution. Walther (2002), Pal, Woodroffe and Meyer (2007) and Rufibach (2007) successfully employ the ICMA for this problem. The last reference gives a very detailed description of the algorithm and also compares the ICMA to several other algorithms that can be used for this problem, such as an interior point method; see, for example, Terlaky and Vial (1998). The ICMA shows a clearly superior performance in these simulation studies. Recently, Dümbgen, Hüsler and Rufibach (2007) have computed the log-concave MLE with an active set algorithm; see, for example, Fletcher (1987). Active set algorithms have the attractive property that they find the solution in finitely many steps, while the iterations of the ICMA have to be terminated by a stopping criterion. It appears that the active set algorithm provides the most efficient method for computing the MLE to date. Both the ICMA and the active set algorithm for computing the log-concave MLE are available with the **R** package “logcondens,” which is accessible from “CRAN.” An alternative way to compute the MLE with convex programming algorithms is described in Koenker and Mizera (2008).

Another advantage of the log-concave MLE \hat{f}_n is that sampling from \hat{f}_n is quite straightforward: First, compute the c.d.f. \hat{F}_n at the ordered sample x_1, \dots, x_n by integrating the piecewise exponential function \hat{f}_n . Next, generate a random index $J \in \{2, \dots, n\}$ with $P(J = j) = \hat{F}_n(x_j) - \hat{F}_n(x_{j-1})$. Then generate $U \sim U[0, 1]$ and set $\Theta := \hat{\phi}_n(x_J) - \hat{\phi}_n(x_{J-1})$. If $\Theta \neq 0$, set $V := \log(1 + (\exp(\Theta) - 1)U) / \Theta$, otherwise set $V := U$. Then $X := x_{J-1} + (x_J - x_{J-1})V$ has density \hat{f}_n .

5. THE MULTIVARIATE CASE

The definition of a log-concave density does not depend on the underlying dimension; see (1). The fact that the MLE does not require the choice of a tuning parameter makes its use even more attractive in a multivariate setting, where, for example, a kernel estimator requires the difficult choice of a bandwidth matrix. The structure of the multivariate MLE is analogous to the univariate case; see, for example, Cule, Samworth and Stewart (2008): The support of the MLE is the convex hull of the data, and there is a triangulation of this convex hull such that $\log \hat{f}_n$ is linear on each simplex of the triangulation. Figure 3 depicts an example for two-dimensional data. The multivariate MLE has already shown promise in a number of applications; see Section 6.

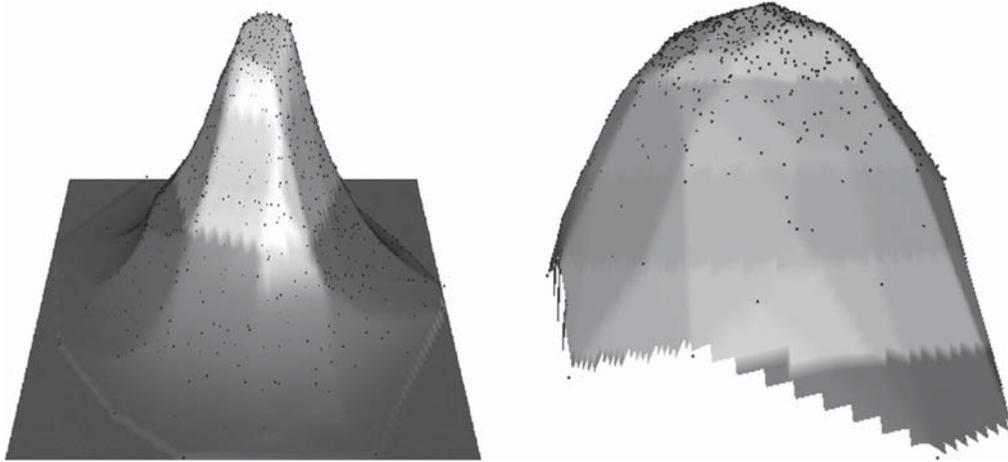


FIG. 3. The MLE \hat{f}_n (left) and $\hat{\phi}_n = \log \hat{f}_n$ (right) for $n = 1000$ observations (plotted as dots) from a standard bivariate normal distribution. The plots are from Cule, Samworth and Stewart (2008).

The computation of the MLE requires an approach that is different from the univariate setting, as the multivariate piecewise linear structure of $\log \hat{f}_n$ does not allow to write this optimization problem in terms of a simple ordering of the slopes. Cule, Samworth and Stewart (2008) show how the MLE can be computed by solving a nondifferentiable convex optimization problem using Shor's r -algorithm; see Kappel and Kuntsevich (2000). Cule, Samworth and Stewart (2008) report a robust and accurate performance of this algorithm, which they implemented in the R package `Log-ConcDEAD`; see Cule, Gramacy and Samworth (2009). However, the computation time increases quickly with sample size and dimension. Cule, Samworth and Stewart (2008) report computation times of about 1 sec for $n = 100$ observations in two dimensions, to 37 min for a sample of size $n = 1000$ in four dimensions. It is therefore desirable to develop faster algorithms for this problem.

Cule, Samworth and Stewart (2008) investigate the finite sample performance of the multivariate MLE via a simulation study. They compare the mean integrated squared error of the MLE with that of a kernel estimator with Gaussian kernel and a bandwidth that is either chosen to minimize the mean integrated squared error (using knowledge about the density that would not be available in practice) or determined by an empirical bandwidth selector based on least squares cross validation. The MLE outperforms both of these estimators except for small sample sizes, and the improvement can be quite dramatic. On the other hand, in view of the work of Birgé and Massart (1993), it seems unlikely that the MLE will achieve optimal rates of convergence in dimensions $d > 4$, due to the richness of

the class of concave functions. It would thus be helpful to have theoretical results about the performance of the multivariate MLE. Deriving such results is an open problem.

6. APPLICATIONS IN MODELING AND INFERENCE

One of the most fruitful applications of log-concave distributions has been in the area of clustering. A principled and successful approach to assign the observations to clusters is via the mixture model $f(x) = \sum_{m=1}^k \pi_m f_m(x)$, where the mixture proportions π_m are nonnegative and sum to unity, and the component distributions f_m model the conditional density of the data in the m th cluster; see, for example, McLachlan and Peel (2000). Typically one assumes a parametric formulation $f_m(x) = f(\theta_m, x)$ for the component distributions, such as the normal model; see, for example, Fraley and Raftery (2002). Then the EM algorithm provides an elegant solution to fit the above mixture model and to assign the data to one of the k components: The EM algorithm iteratively assigns the data based on the current maximum likelihood estimates of the component distributions, and then updates those estimates $\hat{\pi}_m, \hat{\theta}_m$ based on these assignments. An important advantage of using a mixture model for clustering is that it provides not only an assignment of the data to the k components, but also a measure of uncertainty for this assignment via the posterior probabilities that the i th observation belongs to the m th component: $\hat{\pi}_m \hat{f}_m(X_i) / \sum_{j=1}^k \hat{\pi}_j \hat{f}_j(X_i)$.

A disadvantage of this approach is that it depends on the parametric formulation in several important ways:

If the parametric model is misspecified, then the accuracy of the clustering may deteriorate and the measure of uncertainty may be considerably off. For some data, such as those in Figure 2, no appropriate parametric model may be available. Another disadvantage is that each parametric model requires a different implementation of the EM algorithm based on certain theoretical derivations; see, for example, McLachlan and Krishnan (1997).

Therefore, it is desirable to have an EM-type clustering algorithm with nonparametric component distributions. This would allow for a universal software implementation with flexible component distributions. As was expounded in Sections 1 and 2, the class of log-concave distributions provides a flexible model, and, moreover, the MLE exists. Thus, one may attempt to mimic the EM-type clustering algorithm that works so well in the parametric context. This idea was successfully carried out in Chang and Walther (2007) and in Cule, Samworth and Stewart (2008). In related work, Eilers and Borgdorff (2007) use a nonparametric smoother in place of the log-concave MLE in the M-step, with a penalty term that moves the estimate toward a log-concave function. Chang and Walther (2007) report a clear improvement compared to the parametric EM algorithm when the parametric model is not correct, and a performance that is almost similar to the Gaussian EM algorithm in the case where the Gaussian model is correct. Thus, the use of log-concave component distributions provides a flexible methodology for clustering, and this flexibility does

not entail any noticeable penalty in the special case where a parametric model is appropriate.

Chang and Walther (2007) also consider a multivariate extension by modeling each component distribution with log-concave marginals and a normal copula for the dependence structure. This simple multivariate extension avoids the more challenging task of estimating a multivariate log-concave density, but it is flexible enough for many situations. Figure 4 compares the fitted components with those for the Gaussian model for simulated bivariate data. The log-concave model automatically picks up the skewness in the y -direction and results in a noticeably improved error rate for the clustering; see Chang and Walther (2007) for details.

Cule, Samworth and Stewart (2008) extend this approach by using the multivariate log-concave MLE for each component. They apply the log-concave EM algorithm to the Wisconsin breast cancer data of Street et al. (1993) and obtain only 121 misclassified instances compared to 144 with the Gaussian EM algorithm. Figure 5 shows a scatterplot of the data and the fitted log-concave mixture. The contour plots of the fitted components from the Gaussian EM algorithm and the log-concave EM algorithm are given in Figure 1.

Developing principled methodology for selecting an appropriate number of components is an open problem. Methodology for testing for the presence of mixing in the log-concave model is given by Walther (2001) and Walther (2002), where the latter approach uses the fact that a log-concave mixture allows the representation

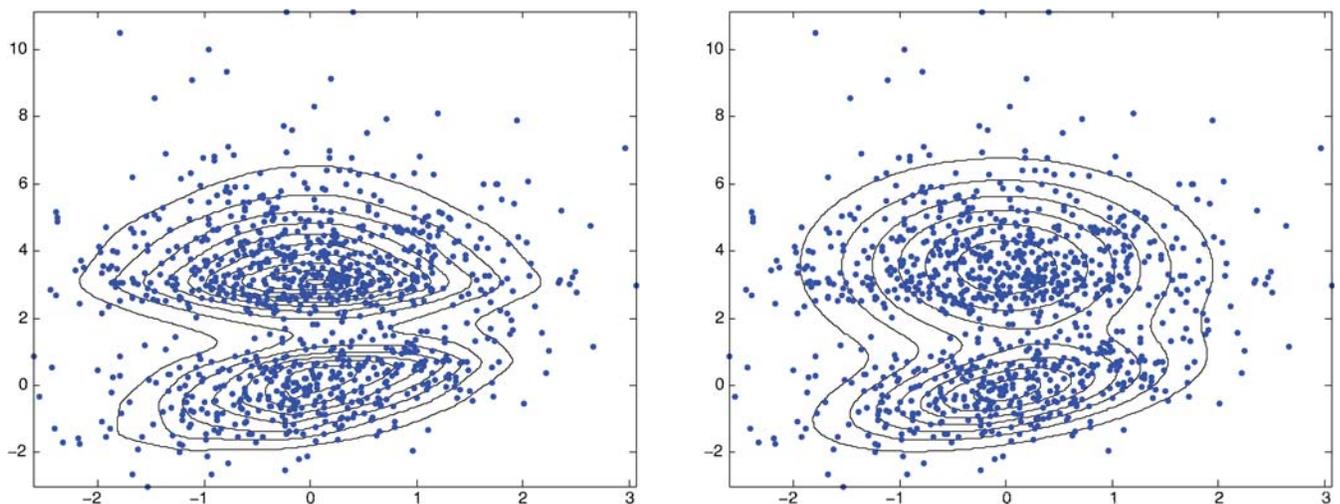


FIG. 4. Contours of the estimated model obtained from the log-concave EM algorithm of Chang and Walther (2007) (left) and from the Gaussian EM algorithm (right) based on the plotted observations. The underlying distribution has a skewed (shifted gamma) distribution in the y -direction of the top component. The plots are from Chang and Walther (2007).

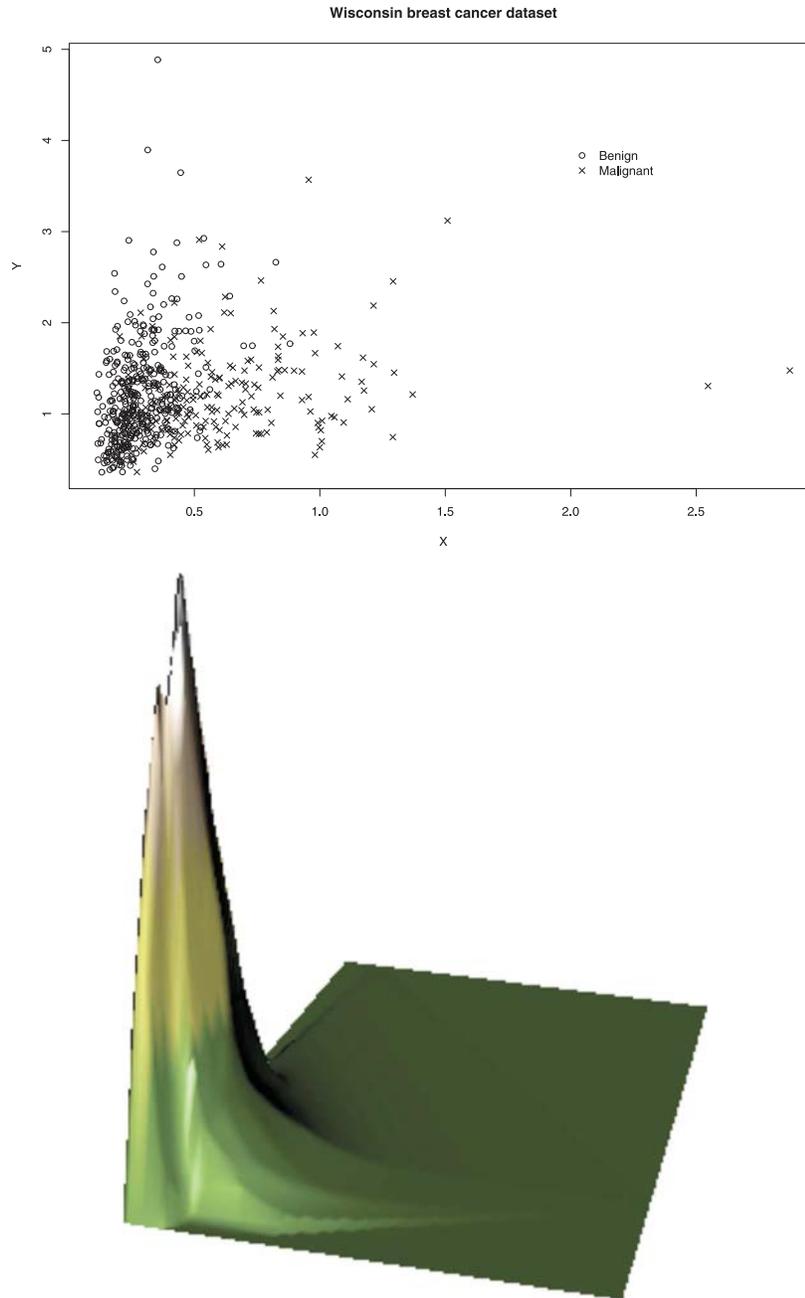


FIG. 5. The Wisconsin breast cancer data (top), with benign cases as open circles and malignant cases as crosses. The bottom plot shows the fitted mixture distribution from the log-concave EM algorithm. The plots are from *Cule, Samworth and Stewart (2008)*.

$\exp(\phi(x) + c\|x\|^2)$ for some $c \geq 0$ and a concave function ϕ .

While log-concave distributions allow for flexible modeling, the structure provided by a log-concave estimator has turned out to result in advantageous properties in a number of other inference problems:

Dümbgen and Rufibach (2009) use the fact that the hazard rate of a log-concave density is automatically monotone and construct a simple plug-in estimator of

the hazard rate which is nondecreasing. Rates of convergence for \hat{f}_n automatically translate to rates for the hazard rate estimator.

Müller and Rufibach (2009) report an improved performance for certain problems in extreme value theory when employing a log-concave estimator.

Dümbgen, Hüsler and Rufibach (2007) show how the assumption of log-concavity allows the estimation of a distribution based on arbitrarily censored data us-

ing the EM algorithm. They replace the log-likelihood function by a function that is linear in ϕ . This function can be interpreted as the conditional expectation of the log-likelihood function given the available data and represents the E-step in the EM algorithm. The M-step consists of maximizing this function using the active set algorithm described in Section 4.

Balabdaoui, Rufibach and Wellner (2009) investigate the mode of \hat{f}_n as an estimator of the mode of f . Estimation of the mode of a unimodal density has received considerable attention in the literature. Typically, some choice of bandwidth or tuning parameter is required due to the problems with the MLE of a univariate density described in Section 1. The MLE of a log-concave density does not suffer from this problem and provides an estimate of the mode as a by-product. Balabdaoui, Rufibach and Wellner (2009) establish the limiting distribution of this estimator and show that the estimator is optimal in the asymptotic minimax sense.

7. SUMMARY AND FUTURE WORK

Log-concave distributions constitute a flexible non-parametric class which allows modeling and inference without a tuning parameter. The MLE has favorable theoretical performance properties and can be computed with available algorithms. These advantageous properties have resulted in tangible improvements in a number of relevant problems, such as in clustering and when handling censored data.

As for future work, there is clearly the potential for similar improvements in a host of other problems, such as regression (see, e.g., Eilers, 2005) or Cox regression under shape constraints on the hazard rate. Further, it would be useful to study the consequences of model misspecification. For example, the mode of the log-concave MLE is a useful tool for data analysis. It would thus be interesting to investigate how far off this mode can be from the population mode in the case where the population distribution is unimodal but not log-concave. The outstanding performance of the multivariate MLE reported in the simulation studies in Cule, Samworth and Stewart (2008) lends importance to a theoretical investigation of its convergence properties. Finally, it would be desirable to develop faster algorithms for computing the multivariate MLE.

For modeling with heavier, algebraic tails, it may be of interest to consider the more general class of ρ -concave densities; see Avriel (1972), Borell (1975) and Dharmadhikari and Joag-Dev (1988). First results about nonparametric estimation and computational issues in this class were obtained in Koenker and Mizera (2008) and Seregin (2008).

ACKNOWLEDGMENTS

Thanks to Kaspar Rufibach and a referee for comments and several references, to Richard Samworth for providing figures, and to Jon Wellner for bringing the work of Arseni Seregin to my attention. Work supported by NSF Grant DMS-05-05682 and NIH Grant 1R21AI069980.

REFERENCES

- AN, M. Y. (1995). Log-concave probability distributions: Theory and statistical testing. Technical report, Economics Dept., Duke Univ.
- AN, M. Y. (1998). Logconcavity versus logconvexity: A complete characterization. *J. Econom. Theory* **80** 350–369. MR1637480
- AVRIEL, M. (1972). r -convex functions. *Math. Program.* **2** 309–323. MR0301151
- BAGNOLI, M. and BERGSTROM, T. (2005). Log-concave probability and its applications. *Econometric Theory* **26** 445–469. MR2213177
- BALABDAOUI, R., RUFIBACH, K. and WELLNER, J. A. (2009). Limit distribution theory for maximum likelihood estimation of a log-concave density. *Ann. Statist.* **37** 1299–1331. MR2509075
- BARLOW, R. E. and PROSCHAN, F. (1975). *Statistical Theory of Reliability and Life Testing*. Holt, Reinhart and Winston, New York. MR0438625
- BIRGÉ, L. (1997). Estimation of unimodal densities without smoothness assumptions. *Ann. Statist.* **25** 970–981. MR1447736
- BIRGÉ, L. and MASSART, P. (1993). Rates of convergence for minimum contrast estimators. *Probab. Theory Related Fields* **97** 113–150. MR1240719
- BORELL, C. (1975). Convex set functions in d -space. *Period. Math. Hungar.* **6** 111–136. MR0404559
- BROOKS, S. P. (1998). MCMC convergence diagnosis via multivariate bounds on log-concave densities. *Ann. Statist.* **26** 398–433. MR1608152
- CAPLIN, A. and NALEBUFF, B. (1991). Aggregation and social choice: A mean voter theorem. *Econometrica* **59** 1–23. MR1085582
- CHANG, G. and WALTHER, G. (2007). Clustering with mixtures of log-concave distributions. *Comput. Statist. Data Anal.* **51** 6242–6251. MR2408591
- CULE, M., SAMWORTH, R. and STEWART, M. (2008). Maximum likelihood estimation of a multidimensional log-concave density. Preprint.
- CULE, M., GRAMACY, R. and SAMWORTH, R. (2009). LogConcDEAD: An R package for maximum likelihood estimation of a multivariate log-concave density. *J. Statist. Software* **29**.
- DELLAPORTAS, P. and SMITH, A. F. M. (1993). Bayesian inference for generalized linear and proportional hazards models via Gibbs sampling. *J. Roy. Statist. Soc. Ser. C* **42** 443–460. MR1222610
- DHARMADHIKARI, S. and JOAG-DEV, K. (1988). *Unimodality, Convexity, and Applications*. Academic Press, Boston. MR0954608
- DÜMBGEN, L., HÜSLER, A. and RUFIBACH, K. (2007). Active set and EM algorithms for log-concave densities based on complete and censored data. Preprint.

- 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.
- EILERS, P. H. C. (2005). Unimodal smoothing. *J. Chemometrics* **19** 317–328.
- EILERS, P. H. C. and BORGDORFF, M. W. (2007). Non-parametric log-concave mixtures. *Comput. Statist. Data Anal.* **51** 5444–5451. MR2370883
- FLETCHER, R. (1987). *Practical Methods for Optimization*, 2nd ed. Wiley, New York. MR0955799
- FRALEY, C. F. and RAFTERY, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *J. Amer. Statist. Assoc.* **97** 611–631. MR1951635
- GILKS, W. R. and WILD, P. (1992). Adaptive rejection sampling for Gibbs sampling. *J. Roy. Statist. Soc. Ser. C* **41** 337–348.
- GRENANDER, U. (1956). On the theory of mortality measurement, part II. *Skandinavisk Aktuarietidskrift* **39** 125–153. MR0093415
- GROENEBOOM, P. and WELLNER, J. A. (1992). *Information Bounds and Nonparametric Maximum Likelihood Estimation*. Birkhäuser, Basel. MR1180321
- HAMPEL, F. R. (1987). Design, modelling and analysis of some biological datasets. In *Design, Data and Analysis, By Some Friends of Cuthbert Daniel* (C. L. Mallows, ed.) 93–128. Wiley, New York.
- IBRAGIMOV, I. A. (1956). On the composition of unimodal distributions. *Theory Probab. Appl.* **1** 255–260. MR0087249
- JONGBLOED, G. (1998). The iterative convex minorant algorithm for nonparametric estimation. *J. Comput. Graph. Statist.* **7** 310–321. MR1646718
- KAPPEL, F. and KUNTSEVICH, A. (2000). An implementation of Shor's r -algorithm. *Comput. Optim. Appl.* **15** 193–205. MR1747059
- KARLIN, S. (1968). *Total Positivity, Vol. I*. Stanford Univ. Press, Stanford. MR0230102
- KOENKER, R. and MIZERA, I. (2008). Quasi-concave density estimation. Technical report.
- MARSHALL, A. W. and OLKIN, I. (1979). *Inequalities: Theory of Majorization and Its Applications*. Academic Press, Orlando. MR0552278
- MCLACHLAN, G. J. and KRISHNAN, T. (1997). *The EM Algorithm and Extensions*. Wiley, New York. MR1417721
- MCLACHLAN, G. J. and PEEL, D. (2000). *Finite Mixture Models*. Wiley, New York. MR1789474
- MÜLLER, S. and RUFIBACH, K. (2009). Smooth tail index estimation. *J. Statist. Comput. Simul.* **79** 1155–1167.
- PAL, J., WOODROOFE, M. and MEYER, M. (2007). Estimating a Polya frequency function. In *Complex Datasets and Inverse Problems: Tomography, Networks and Beyond* (R. Liu, W. Straderman, C.-H. Zhang, eds.) 239–249. IMS, Beachwood, OH. MR2459196
- RUFIBACH, K. (2006). Log-concave density estimation and bump hunting for I.I.D. observations. Ph.D. thesis, Univ. Bern and Göttingen.
- RUFIBACH, K. (2007). Computing maximum likelihood estimators of a log-concave density function. *J. Statist. Comput. Simul.* **77** 561–574. MR2407642
- SEREGIN, A. (2008). Estimation of convex-transformed densities. General Exam Proposal. Statistics Dept., Univ. Washington.
- SILVERMAN, B. W. (1982). On the estimation of a probability density function by the maximum penalized likelihood method. *Ann. Statist.* **10** 795–810. MR0663433
- STREET, W. M., WOLBERG, W. H. and MANGASARIAN, O. L. (1993). Nuclear feature extraction for breast tumor diagnosis. IS& T/SPIE 1993 International Symposium on Electronic Imaging: Science and Technology, San Jose, CA, 1905, 861–870.
- TERLAKY, T. and VIAL, J. P. (1998). Computing maximum likelihood estimators of convex density functions. *SIAM J. Sci. Comput.* **19** 675–694. MR1618864
- WALTHER, G. (2001). Multiscale maximum likelihood analysis of a semiparametric model, with applications. *Ann. Statist.* **29** 1298–1319. MR1873332
- WALTHER, G. (2002). Detecting the presence of mixing with multiscale maximum likelihood. *J. Amer. Statist. Assoc.* **97** 508–513. MR1941467
- WANG, X., WOODROOFE, M., WALKER, M., MATEO, M. and OLSZEWSKI, E. (2005). Estimating dark matter distributions. *Astrophys. J.* **626** 145–158.
- WOODROOFE, M. and SUN, J. (1993). A penalized maximum likelihood estimate of $f(0+)$ when f is non-increasing. *Statist. Sinica* **3** 501–515. MR1243398