

PRACTICAL CONSIDERATIONS IN BOUNDARY ESTIMATION: MODEL-ROBUSTNESS, EFFICIENT COMPUTATION, AND BOOTSTRAPPING

BY CHARU KRISHNAMOORTHY AND ED CARLSTEIN*
Glaxo Research Institute and University of North Carolina

A data-set consists of independent observations taken at the nodes of a grid. An unknown *boundary* partitions the grid into two regions: All the observations coming from a particular region share a common distribution, but the distributions are different for the two different regions. These two distributions are entirely unknown, and the grid is of arbitrary dimension with rectangular mesh. In this scenario, the boundary can be consistently estimated.

The boundary estimate is selected from an appropriate collection \mathcal{T} of *candidate boundaries* which must be specified by the user. The candidate boundaries must satisfy certain regularity assumptions, including a “richness” condition. In practice, one may be faced with a \mathcal{T} that is not sufficiently “rich”. Is the estimator “robust” with respect to such misspecification of \mathcal{T} ? This question is addressed via a.s. results comparing the asymptotic error of the estimator with the smallest possible error in \mathcal{T} .

Because the boundary estimator requires a search (through \mathcal{T}) and calculation over the whole data grid, efficient computation is a practical necessity. An algorithm is presented which reduces the computational burden by an order of magnitude relative to the naive approach.

A graphical bootstrap procedure is proposed for studying the variability of the boundary estimator. Simulations of this procedure indicate that it does accurately reflect the true sampling behavior of the boundary estimator.

1. Introduction.

1.1. The Boundary Estimation Problem. We observe a collection of independent rv's $\{X_i^I\}$, indexed by nodes i of a finite d -dimensional grid I within the d -dimensional unit cube $\mathcal{U}_d := [0, 1]^d$. The unknown *boundary* Θ is a $(d - 1)$ -dimensional surface that partitions \mathcal{U}_d into two regions, $\bar{\Theta}$ and $\underline{\Theta}$.

* Research supported by National Science Foundation Grants DMS-8902973 and DMS-9205881.

AMS 1991 Subject Classification: Primary 62G05; Secondary 62G20.

Key words and phrases: Change-point, grid, nonparametric, template.

All observations X_i^I made at nodes $i \in \bar{\Theta}$ are from distribution F , while all observations X_i^I made at nodes $i \in \underline{\Theta}$ are from distribution G . The distributions F and G are entirely unknown; the only distributional assumption is that $F \neq G$. The objective is to estimate the unknown boundary Θ , using the observed data $\{X_i^I : i \in I\}$.

The grid is generated by divisions along each coordinate axis in \mathcal{U}_d . Along the j th axis ($1 \leq j \leq d$), there are n_j divisions which are equally spaced at $1/n_j, 2/n_j, \dots, n_j/n_j$. Observations are made at the resulting grid nodes $i := (i_1/n_1, i_2/n_2, \dots, i_d/n_d) \in \mathcal{U}_d$, where $i_j \in \{1, 2, \dots, n_j\}$. The collection of all nodes i is denoted by I , and the total number of observations is $|I| := \prod_{j=1}^d n_j$. In any set $A \subseteq \mathcal{U}_d$, the number of observations (i.e., grid nodes) is $|A|_I := \#\{i \in A\}$.

The notion of a *boundary* in \mathcal{U}_d is formulated in a set-theoretic way: the unknown boundary Θ is identified with the corresponding partition $(\bar{\Theta}, \underline{\Theta})$ of \mathcal{U}_d . The sample-based estimate of Θ will be selected from a finite collection \mathcal{T}_I of *candidate boundaries*, with generic element T . Again, each candidate T is identified with its corresponding partition (\bar{T}, \underline{T}) of \mathcal{U}_d . The total number of candidates considered is $|\mathcal{T}_I| := \#\{T \in \mathcal{T}_I\}$.

1.2. The Boundary Estimator. Our tool for selecting an estimate $\hat{\Theta}_I$ from \mathcal{T}_I is the *empirical cumulative distribution function (ecdf)*. Let $\mathbb{I}\{\cdot\}$ denote the indicator function. For a candidate boundary $T \in \mathcal{T}_I$, compute the ecdf's

$$\bar{h}_T^I(x) := \sum_{i \in \bar{T}} \mathbb{I}\{X_i^I \leq x\} / |\bar{T}|_I \quad \text{and} \quad \underline{h}_T^I(x) := \sum_{i \in \underline{T}} \mathbb{I}\{X_i^I \leq x\} / |\underline{T}|_I,$$

and consider the differences $d_{I_i}^T := |\bar{h}_T^I(X_i^I) - \underline{h}_T^I(X_i^I)|$ for each $i \in I$. Now combine these differences $d_{I_i}^T$ using a “mean-dominant norm” $S_{|I|}(d_{I_1}^T, d_{I_2}^T, \dots, d_{I_{|I|}}^T)$ [see Carlstein (1988)], i.e., a function $S_{|I|}(\cdot) : \mathbb{R}_+^{|I|} \mapsto \mathbb{R}_+^1$ satisfying the following definition:

- (D.a) [Symmetry] $S_{|I|}(\cdot)$ is symmetric in its $|I|$ arguments;
- (D.b) [Homogeneity] $S_{|I|}(\alpha d_1, \alpha d_2, \dots, \alpha d_{|I|}) = \alpha S_{|I|}(d_1, d_2, \dots, d_{|I|})$ whenever $\alpha \geq 0$;
- (D.c) [Triangle Inequality] $S_{|I|}(d_1 + d'_1, d_2 + d'_2, \dots, d_{|I|} + d'_{|I|}) \leq S_{|I|}(d_1, d_2, \dots, d_{|I|}) + S_{|I|}(d'_1, d'_2, \dots, d'_{|I|})$;
- (D.d) [Identity] $S_{|I|}(1, 1, \dots, 1) = 1$;

(D.e) [Monotonicity] $S_{|I|}(d_1, d_2, \dots, d_{|I|}) \leq S_{|I|}(d'_1, d'_2, \dots, d'_{|I|})$ whenever $d_i \leq d'_i \forall i$;

(D.f) [Mean Dominance] $S_{|I|}(d_1, d_2, \dots, d_{|I|}) \geq \sum_{1 \leq i \leq |I|} d_i / |I|$.

Finally, we standardize $S_{|I|}(\cdot)$ by a multiplicative factor to account for the inherent instability in the ecdf. The boundary estimator $\widehat{\Theta}_I$ is defined as the candidate boundary in \mathcal{T}_I which maximizes the criterion function $D_I(T) := (|\overline{T}|_I / |I|)(|\underline{T}|_I / |I|) \cdot S_{|I|}(d_{I1}^T, d_{I2}^T, \dots, d_{I|I|}^T)$ over all $T \in \mathcal{T}_I$. Formally, $\widehat{\Theta}_I := \operatorname{argmax}_{T \in \mathcal{T}_I} D_I(T)$.

The intuitive motivation for this estimator is given by Carlstein and Krishnamoorthy (1992) [referred to as C&K from here on]. A literature review comparing $\widehat{\Theta}_I$ to other related methods is provided in C&K, as is an extensive discussion of examples in the $d = 1$ and $d = 2$ cases (including the change-point problem, the epidemic-change model, linear bisection of the plane, templates, and Lipschitz boundaries).

1.3. Performance of the Boundary Estimator. In order to assess the performance of $\widehat{\Theta}_I$, we must quantify the notion of “distance” between two boundaries (say, T and Θ). Our “distance” measure is the pseudometric $\partial(T, \Theta) := \min\{\lambda(\overline{T} \circ \overline{\Theta}), \lambda(\underline{T} \circ \underline{\Theta})\}$, where \circ denotes set-theoretic symmetric difference, i.e., $(A \circ B) := (A \cap B^c) \cup (A^c \cap B)$, and $\lambda(\cdot)$ is Lebesgue measure over \mathcal{U}_d .

Consider the following set-theoretic regularity conditions on the boundaries:

REGULARITY CONDITION (R.1): [Non-trivial Partitions] For each $T \in \mathcal{T}_I$, $0 < \lambda(\overline{T}) < 1$ and $0 < |\overline{T}|_I / |I| < 1$. Also, $0 < \lambda(\underline{\Theta}) < 1$.

REGULARITY CONDITION (R.2): [Richness of \mathcal{T}_I] For each I , $\exists T_I \in \mathcal{T}_I$ such that the sequence $\{T_I\}$ satisfies: $\partial(\Theta, T_I) \rightarrow 0$ as $|I| \rightarrow \infty$.

REGULARITY CONDITION (R.3): [Cardinality of \mathcal{T}_I] For each $\gamma > 0$, $|\mathcal{T}_I| \cdot \exp\{-\gamma \cdot |I|\} \rightarrow 0$ as $|I| \rightarrow \infty$.

REGULARITY CONDITION (R.4): [Smoothness of Perimeter] Denote $\mathcal{F}_I := \{\overline{\Theta}, \overline{T} : T \in \mathcal{T}_I\}$ and $\mathcal{P}_I(A) := \{C \in \mathcal{C}_I : C \cap A \neq \phi \text{ and } C \cap A^c \neq \phi\}$, where \mathcal{C}_I is the collection of $|I|$ “rectangular” cells C induced by the grid partition of \mathcal{U}_d , and where $A \subseteq \mathcal{U}_d$. We require $\sup_{A \in \mathcal{F}_I} \lambda(\mathcal{P}_I(A)) \rightarrow 0$ as $|I| \rightarrow \infty$.

In C&K, these regularity conditions are seen to be intuitively natural and are explicitly checked for several examples. Moreover, R.1–R.4 are shown to imply strong consistency of the boundary estimator $\hat{\Theta}_I$, i.e., $\partial(\Theta, \hat{\Theta}_I) \xrightarrow{\text{a.s.}} 0$ as $|I| \rightarrow \infty$.

1.4. Some Practical Considerations in Boundary Estimation.

Model-Robustness: Condition R.2 requires \mathcal{T}_I to contain some candidate boundary T_I that “gets close” to the true Θ . If no such “ideal” candidate were available, we could not possibly hope to statistically select an estimator $\hat{\Theta}_I$ from \mathcal{T}_I in such a way that consistency holds. Yet, in practice, the collection \mathcal{T}_I may be misspecified by the user and may be incompatible with the true unknown boundary Θ . Even though R.2 is violated, one still hopes that $\hat{\Theta}_I$ will contain some useful information about Θ . The performance of the estimator $\hat{\Theta}_I$ when faced with misspecification of \mathcal{T}_I is called “model-robustness” (as opposed to “distributional-robustness”, which focuses on the effects of, e.g., heavy-tailed distributions). In Section 2, this model-robustness issue is quantified by comparing the error of $\hat{\Theta}_I$ with the minimal possible error in \mathcal{T}_I ; theoretical results are given in terms of a.s. asymptotic behavior and in terms of probability of error.

Efficient Computation: Condition R.3 restricts the number of candidate boundaries relative to the sample size, but still allows for extremely large $|\mathcal{T}_I|$; if the user is concerned about violating R.2, then caution should be exercised in attempting to reduce $|\mathcal{T}_I|$. And, for higher dimensional data-sets ($d \geq 2$), $|I|$ tends to be large. Therefore, since the estimator requires calculations over all $i \in I$ and over all $T \in \mathcal{T}_I$, it behooves us to develop efficient computational schemes. In Section 3, an algorithm is presented which improves dramatically on the naive approach.

Bootstrapping: Although $\hat{\Theta}_I$ is a consistent estimator of Θ (in the sense described above), it is difficult to characterize the sampling variability of the boundary estimator – especially in cases where the boundary is not readily expressible as a parametric function. Therefore, in Section 4, we suggest a graphical bootstrap procedure for visually assessing the sampling variability of $\hat{\Theta}_I$; this procedure was simulated, yielding encouraging results.

2. Model-Robustness.

2.1. *Motivation.* Knowledge about the form of the boundary Θ (e.g.,

rectangular template, Lipschitz function, etc.) has been assumed in the formulation of $\widehat{\Theta}_I$. This information is used to construct “rich enough” collections \mathcal{T}_I (see R.2), so that boundaries which are “close” to the true boundary Θ are available as possible estimates. Thus, if the user has knowledge that Θ defines a circle, then \mathcal{T}_I is constructed to include only circles. If further information is available on the *location* of Θ (say the center of the circle is known to be in the left half of \mathcal{U}_2), then a *smaller* set \mathcal{T}_I of circles can be constructed which will still satisfy R.2. Smaller $|\mathcal{T}_I|$ means easier computability (see Section 3) and sharper bounds on the error probability (see Theorem 2 of C&K). Therefore, there are practical incentives for small \mathcal{T}_I . It is, however, possible for \mathcal{T}_I to be *too small* if it violates R.2. This may happen due to overzealousness in reducing the computational burden, or due to incorrect prior information. Thus, a question of practical interest is: What is the price paid if the collections \mathcal{T}_I do not satisfy R.2? A useful illustration to keep in mind is the case where the true Θ is a circle, but \mathcal{T}_I contains only squares.

Formally, let $\Theta'_I := \arg \min_{T \in \mathcal{T}_I} \partial(\Theta, T)$ be the element in \mathcal{T}_I closest to Θ . Then $\partial(\Theta'_I, \Theta)$ is the minimal error that we can expect to make in estimating Θ ; if the error of our estimator, $\partial(\widehat{\Theta}_I, \Theta)$, is of the same order as $\partial(\Theta'_I, \Theta)$ we shall be reasonably satisfied, i.e., $\widehat{\Theta}_I$ is “model-robust”. We will study the asymptotic behavior of $\partial(\widehat{\Theta}_I, \Theta)$ and compare it to the asymptotic behavior of $\partial(\Theta'_I, \Theta)$. Notice that Θ'_I is not necessarily *unique*: There might well be a class of candidates Θ'_I that are closest in terms of the pseudometric ∂ , but each of them might yield a different value of $\partial(\widehat{\Theta}_I, \Theta'_I)$. Because of this possible ambiguity the study of $\partial(\widehat{\Theta}_I, \Theta'_I)$ is not appropriate here. Also note that alternative formulations of the original boundary estimation problem (Section 1.1) could allow the “true” boundary to depend on the grid I , yielding better compatibility with \mathcal{T}_I .

2.2. Asymptotic Target Error. Assume that R.1 holds. Given Θ and the candidate family $\{\mathcal{T}_I : \forall I\}$, define the *asymptotic target error*, η , as follows:

$$\eta := \overline{\lim}_{|I| \rightarrow \infty} \min_{T \in \mathcal{T}_I} \partial(\Theta, T) = \overline{\lim}_{|I| \rightarrow \infty} \partial(\Theta, \Theta'_I).$$

We want to study the model-robustness of $\widehat{\Theta}_I$ with respect to violation of R.2. The quantity η is the natural parameter for measuring the severity of the failure of R.2, because:

PROPOSITION 1. R.2 $\Leftrightarrow \eta = 0$.

This also justifies the use of $\overline{\lim}$ rather than $\underline{\lim}$ in η 's definition, since the latter could yield zero even if R.2 fails and hence does not adequately quantify the departure from our conditions. We will assess the model-robustness of $\widehat{\Theta}_I$ by comparing $\partial(\widehat{\Theta}_I, \Theta)$ to the target η .

Another quantity which arises naturally in the model-robustness analysis is:

$$\rho(T) := |\lambda(\overline{T} \cap \overline{\Theta})\lambda(\underline{T}) - \lambda(\underline{T} \cap \overline{\Theta})\lambda(\overline{T})|.$$

The difference $\rho(\Theta) - \rho(T)$ behaves similarly to $\partial(\Theta, T)$, as formalized by:

PROPOSITION 2. For $\gamma > 0$,

$$\partial(\Theta, T) < \gamma \Rightarrow \rho(\Theta) - \rho(T) < \gamma,$$

$$\partial(\Theta, T) > \gamma \Rightarrow \rho(\Theta) - \rho(T) > \sigma \cdot \gamma > 0,$$

where $\sigma := \min\{\lambda(\overline{\Theta}), \lambda(\underline{\Theta})\}$. It follows that:

$$\rho(\Theta) - \rho(T) \leq \partial(\Theta, T) \leq [\rho(\Theta) - \rho(T)]/\sigma.$$

When the partition defined by Θ is “balanced”, i.e., $\sigma = \frac{1}{2}$, we find $\partial(\Theta, T) = 2[\rho(\Theta) - \rho(T)]$.

[Proofs of results are given in Section 2.4 below.] When R.2 does hold, then

$$\overline{\lim}_{|I| \rightarrow \infty} \min_{T \in \mathcal{T}_I} \{\rho(\Theta) - \rho(T)\} = 0.$$

However, unlike $\partial(\Theta, T)$, the quantity $\rho(\Theta) - \rho(T)$ is *not* a pseudometric; in fact, it is not even symmetric in its arguments. Nevertheless, to analyze the model-robustness of $\widehat{\Theta}_I$, it is useful to define the analog of η for $\rho(\Theta) - \rho(T)$:

$$\eta^* := \overline{\lim}_{|I| \rightarrow \infty} \min_{T \in \mathcal{T}_I} \{\rho(\Theta) - \rho(T)\} = \overline{\lim}_{|I| \rightarrow \infty} [\rho(\Theta) - \rho(\Theta''_I)],$$

where $\Theta''_I := \arg \min_{T \in \mathcal{T}_I} \{\rho(\Theta) - \rho(T)\}$. The relationship between η and η^* is given by:

PROPOSITION 3. $0 \leq \eta^* \leq \eta \leq \eta^*/\sigma$.

Notice that R.2 $\Leftrightarrow \eta = 0 \Leftrightarrow \eta^* = 0$; also, when the partition defined by Θ is balanced, the third inequality becomes an equality.

2.3. *Model-Robustness Results.* Assume now that: $F \neq G$; $\widehat{\Theta}_I$ is based on a mean-dominant norm; regularity conditions R.1, R.3, R.4 hold. The main result provides an a.s. bound on $\widehat{\Theta}_I$'s error that is linear in the target error.

THEOREM 1. [*Model-Robustness*]

$$\eta \leq \overline{\lim}_{|I| \rightarrow \infty} \partial(\Theta, \widehat{\Theta}_I) \stackrel{\text{a.s.}}{\leq} \eta^* / \sigma\mu,$$

where $\mu := \mu_F \lambda(\overline{\Theta}) + \mu_G \lambda(\underline{\Theta})$ and $\mu_F := \int_{-\infty}^{\infty} |F(x) - G(x)| dF(x)$, $\mu_G := \int_{-\infty}^{\infty} |F(x) - G(x)| dG(x)$. When the partition defined by Θ is balanced,

$$\eta \leq \overline{\lim}_{|I| \rightarrow \infty} \partial(\Theta, \widehat{\Theta}_I) \stackrel{\text{a.s.}}{\leq} \eta / \mu.$$

Notice that $\sigma\mu > 0$ (by R.1 and Lemma 7 of C&K). The left inequalities in Theorem 1 hold *always*, while the right inequalities hold *a.s.* The upper bounds in Theorem 1 show that the model-robustness of $\widehat{\Theta}_I$ is favorably influenced by two intuitively natural factors: (i) *balance* between the amount of data in $\overline{\Theta}$ versus $\underline{\Theta}$, as measured by σ ; (ii) *disparity* between the two distributions F and G , as measured by μ .

Example. Let F be the distribution with point mass at the origin, and let G be the Uniform[0, 1] distribution. Then direct calculations yield $\mu_F + \mu_G = \frac{3}{2}$. Therefore, in the case of a balanced partition $(\overline{\Theta}, \underline{\Theta})$, we obtain

$$\eta \leq \overline{\lim}_{|I| \rightarrow \infty} \partial(\Theta, \widehat{\Theta}_I) \stackrel{\text{a.s.}}{\leq} \frac{4}{3} \eta.$$

Finally, we bound the probability of $\widehat{\Theta}_I$ behaving in a nonrobust way (as compared to the rhs of Theorem 1).

THEOREM 2. [*Bound on Error Probability*]. For any $0 < \alpha < 1$ and $0 < \epsilon \leq \epsilon(\alpha)$,

$$\mathbb{P}\{\partial(\Theta, \widehat{\Theta}_I) > (\eta^* + \epsilon) / \sigma\mu(1 - \alpha)\} \leq \widetilde{K} \cdot |T_I| \cdot \exp\{-K \cdot \epsilon^2 \cdot |I|\}$$

for $|I|$ sufficiently large, where K and \widetilde{K} are positive constants.

The probability of nonrobustness decreases exponentially as a function of sample size. However, this effect is counterbalanced by the number of candidate boundaries considered (i.e., for fixed target error, it is easier to

grossly mislead $\widehat{\Theta}_I$ when the collection \mathcal{T}_I is larger). The bound in Theorem 2 explains why the joint growth of $|I|$ and $|\mathcal{T}_I|$ must be controlled via R.3. Lastly, notice that the bound does not depend on the actual target error, i.e., it is unaffected by the severity of the R.2 violation.

2.4. Proofs.

PROOF OF PROPOSITION 1. Immediate from the definitions. ■

PROOF OF PROPOSITION 2. Note that

$$\rho(\Theta) - \rho(T) = \min\{\lambda(\underline{\Theta} \cap \overline{T})\lambda(\overline{\Theta}) + \lambda(\overline{\Theta} \cap \underline{T})\lambda(\underline{\Theta}), \lambda(\overline{\Theta} \cap \overline{T})\lambda(\underline{\Theta}) + \lambda(\underline{\Theta} \cap \underline{T})\lambda(\overline{\Theta})\}.$$

Comparing this expression to the definition of $\partial(\Theta, T)$, the first implication is clear. For the second implication, write $\partial(\Theta, T) = \min\{x + x', y + y'\}$ and $\rho(\Theta) - \rho(T) = \min\{\Gamma_{\Theta}(x, x'), \Gamma_{\Theta}(y, y')\}$, where $x := \lambda(\underline{\Theta} \cap \overline{T})$, $x' := \lambda(\overline{\Theta} \cap \underline{T})$, $y := \lambda(\underline{\Theta} \cap \underline{T})$, $y' := \lambda(\overline{\Theta} \cap \overline{T})$, and $\Gamma_{\Theta}(z, z') := \lambda(\overline{\Theta})z + \lambda(\underline{\Theta})z'$. Then

$$\begin{aligned} \partial(\Theta, T) > \gamma &\Rightarrow x + x' > \gamma \text{ and } y + y' > \gamma \\ &\Rightarrow \sigma(x + x') > \sigma\gamma \text{ and } \sigma(y + y') > \sigma\gamma \\ &\Rightarrow \Gamma_{\Theta}(x, x') > \sigma\gamma \text{ and } \Gamma_{\Theta}(y, y') > \sigma\gamma \Rightarrow \rho(\Theta) - \rho(T) > \sigma\gamma. \end{aligned}$$

The final two assertions of Proposition 2 are now immediate. ■

PROOF OF PROPOSITION 3. Since $\Theta'_I, \Theta''_I \in \mathcal{T}_I$, we have by definition and by Proposition 2 that

$$0 \leq \rho(\Theta) - \rho(\Theta''_I) \leq \rho(\Theta) - \rho(\Theta'_I) \leq \partial(\Theta, \Theta'_I) \leq \partial(\Theta, \Theta''_I) \leq [\rho(\Theta) - \rho(\Theta''_I)]/\sigma.$$

Now take $\overline{\lim}_{|I| \rightarrow \infty}$ throughout. ■

PROOF OF THEOREM 1. For $0 < \alpha < 1$ and $0 < \epsilon \leq \epsilon(\alpha)$, we have by Theorem 2 and R.3 that

$$\sum_{|I|} \mathbb{P}\{\partial(\Theta, \widehat{\Theta}_I) > (\eta^* + \epsilon)/\sigma\mu(1 - \alpha)\} < \infty;$$

hence by the Borel-Cantelli Lemma $\mathbb{P}\{\partial(\Theta, \widehat{\Theta}_I) > (\eta^* + \epsilon)/\sigma\mu(1 - \alpha)$ infinitely often $[[|I|]] = 0$. For $\alpha_k = \frac{1}{k}$ and $\epsilon_k = \min\{\frac{1}{k}, \epsilon(\alpha_k)\}$, $k > 1$, \exists a null set A_k s.t. $\forall \omega \notin A_k$ we have $\partial(\Theta, \widehat{\Theta}_I) \leq (\eta^* + \epsilon_k)/\sigma\mu(1 - \alpha_k)$ when $|I| \geq i(k, \omega)$. Thus, for $\omega \notin \cup_{k>1} A_k$, $\overline{\lim}_{|I| \rightarrow \infty} \partial(\Theta, \widehat{\Theta}_I) \leq \eta^*/\sigma\mu$. ■

PROOF OF THEOREM 2. A series of Lemmas will be presented, leading to Theorem 2.

LEMMA 1. Define $\mathcal{F}_I^* := \{\bar{\Theta}, \underline{\Theta}; \bar{T}, \underline{T}, \bar{T} \cap \bar{\Theta}, \underline{T} \cap \bar{\Theta}, \bar{T} \cap \underline{\Theta}, \underline{T} \cap \underline{\Theta} : T \in \mathcal{T}_I\}$.

Then:

$$\sup_{A \in \mathcal{F}_I^*} |\lambda(A) - |A|_I / |I|| \rightarrow 0 \text{ as } |I| \rightarrow \infty.$$

PROOF. Follows from R.4, using the same argument as in the proof of C&K's Lemma 1. ■

Define the following notation:

$$\bar{\eta}_T(x) := [\lambda(\bar{T} \cap \bar{\Theta})F(x) + \lambda(\bar{T} \cap \underline{\Theta})G(x)] / \lambda(\bar{T}),$$

$$\underline{\eta}_T(x) := [\lambda(\underline{T} \cap \bar{\Theta})F(x) + \lambda(\underline{T} \cap \underline{\Theta})G(x)] / \lambda(\underline{T}),$$

$$\delta_{Ii}^T := |\bar{\eta}_T(X_i^I) - \underline{\eta}_T(X_i^I)|, \Delta_I(T) := \lambda(\bar{T})\lambda(\underline{T})S_{|I|}(\delta_{Ii}^T : i \in I).$$

LEMMA 2. For $|I| \geq N_1(\epsilon)$,

$$\mathbb{P}\{\sup_{T \in \mathcal{T}_I} |D_I(T) - \Delta_I(T)| > \epsilon\} \leq K_1 |\mathcal{T}_I| \exp\{-K_2 \cdot \epsilon^2 \cdot |I|\}.$$

PROOF. Follows from Lemma 1 and Dvoretzky, Kiefer and Wolfowitz (1956), using the same argument as in the proof of C&K's Lemma 2 (with $\delta = 0$). ■

LEMMA 3. We can write $\Delta_I(T) = \rho(T) \cdot S_{|I|}(\delta_{Ii}^\Theta : i \in I)$.

PROOF. Exactly as in the proof of C&K's Lemma 3. ■

LEMMA 4. For every $T \in \mathcal{T}_I$, we have $\Delta_I(T) \leq \Delta_I(\Theta)$.

PROOF. Follows from Lemma 3, exactly as in the proof of C&K's Lemma 4. ■

LEMMA 5. For $|I| \geq N_2(\epsilon)$,

$$\mathbb{P}\{|\Delta_I(\hat{\Theta}_I) - \Delta_I(\Theta)| > \eta^* + \epsilon\} \leq K_3 |\mathcal{T}_I| \exp\{-K_4 \cdot \epsilon^2 \cdot |I|\}.$$

PROOF. By definition $\Theta_I'' \in \mathcal{T}_I$ is the maximizer of $\rho(\cdot)$, and hence, by Lemma 3, it is the maximizer of $\Delta_I(\cdot)$, over \mathcal{T}_I . Then by Lemma 4 we have $\Delta_I(\Theta) \geq \Delta_I(\Theta_I'') \geq \Delta_I(T) \forall T \in \mathcal{T}_I$, and by definition we have $D_I(\hat{\Theta}_I) \geq D_I(T) \forall T \in \mathcal{T}_I$. Now,

$$\begin{aligned} |\Delta_I(\hat{\Theta}_I) - \Delta_I(\Theta)| &\leq |\Delta_I(\hat{\Theta}_I) - D_I(\hat{\Theta}_I)| + |D_I(\hat{\Theta}_I) - \Delta_I(\Theta_I'')| \\ &\quad + |\Delta_I(\Theta_I'') - \Delta_I(\Theta)|. \end{aligned} \tag{†}$$

The second modulus on the rhs is bounded by $\sup_{T \in \mathcal{T}_I} |D_I(T) - \Delta_I(T)|$, because either $D_I(\widehat{\Theta}_I) \geq \Delta_I(\Theta'_I) \geq \Delta_I(\widehat{\Theta}_I)$ or $\Delta_I(\Theta'_I) \geq D_I(\widehat{\Theta}_I) \geq D_I(\Theta''_I)$. The same bound applies to the first modulus on the rhs of (†). Using Lemmas 4 and 3, and properties (D.e) and (D.d) of $S_{|I|}$, the third modulus on the rhs is $\Delta_I(\Theta) - \Delta_I(\Theta''_I) = [\rho(\Theta) - \rho(\Theta''_I)]S_{|I|}(\delta_{I_i}^\ominus : i \in I) \leq \rho(\Theta) - \rho(\Theta''_I)$. Therefore,

$$|\Delta_I(\widehat{\Theta}_I) - \Delta_I(\Theta)| \leq 2 \cdot \sup_{T \in \mathcal{T}_I} |D_I(T) - \Delta_I(T)| + \rho(\Theta) - \rho(\Theta''_I).$$

For $|I|$ sufficiently large, by the definition of η^* , we have the *deterministic* bound $\rho(\Theta) - \rho(\Theta''_I) < \eta^* + \epsilon/3$. Now combine this with the probabilistic bound from Lemma 2. \blacksquare

To prove Theorem 2, begin by applying Proposition 2, Lemma 4, Lemma 3 and property (D.f) of $S_{|I|}$, obtaining

$$\partial(\Theta, T) > \gamma \Rightarrow |\Delta_I(\Theta) - \Delta_I(T)| = [\rho(\Theta) - \rho(T)]S_{|I|}(\delta_{I_i}^\ominus : i \in I) > \sigma\gamma\delta_I^\ominus,$$

where $\delta_I^\ominus := \sum_{i \in I} \delta_{I_i}^\ominus / |I|$. Thus

$$\begin{aligned} \mathbb{P}\{\partial(\Theta, \widehat{\Theta}_I) > (\eta^* + \epsilon)/\sigma\mu(1 - \alpha)\} &\leq \mathbb{P}\{|\Delta_I(\Theta) - \Delta_I(\widehat{\Theta}_I)| \\ &> (\eta^* + \epsilon)\delta_I^\ominus/\mu(1 - \alpha)\} \\ &\leq \mathbb{P}\{|\Delta_I(\Theta) - \Delta_I(\widehat{\Theta}_I)| > \eta^* + \epsilon\} + \mathbb{P}\{\delta_I^\ominus < \mu(1 - \alpha)\}. \end{aligned}$$

The first probability on the rhs is immediately handled by Lemma 5. The second probability on the rhs is bounded by $\mathbb{P}\{|\delta_I^\ominus - \mu| > \mu\alpha\}$. Denote $\bar{\delta}_I^\ominus := \sum_{i \in \bar{\Theta}} \delta_{I_i}^\ominus / |\bar{\Theta}|_I$ and $\underline{\delta}_I^\ominus := \sum_{i \in \underline{\Theta}} \delta_{I_i}^\ominus / |\underline{\Theta}|_I$, so that $\delta_I^\ominus = \bar{\delta}_I^\ominus (|\bar{\Theta}|_I / |I|) + \underline{\delta}_I^\ominus (|\underline{\Theta}|_I / |I|)$. Notice that

$$|\delta_I^\ominus - \mu| \leq \bar{\delta}_I^\ominus \left| \frac{|\bar{\Theta}|_I}{|I|} - \lambda(\bar{\Theta}) \right| + \lambda(\bar{\Theta}) |\bar{\delta}_I^\ominus - \mu_F| + \underline{\delta}_I^\ominus \left| \frac{|\underline{\Theta}|_I}{|I|} - \lambda(\underline{\Theta}) \right| + \lambda(\underline{\Theta}) |\underline{\delta}_I^\ominus - \mu_G|.$$

The first and third summands on the rhs are handled using Lemma 1. Consider the second summand on the rhs (a similar argument holds for the fourth). By equation (2.3) of Hoeffding (1963), we have $\mathbb{P}\{|\bar{\delta}_I^\ominus - \mu_F| > \mu\alpha/4\} \leq 2 \cdot \exp\{-c_\alpha |\bar{\Theta}|_I\}$. Lemma 1 ensures that $|\bar{\Theta}|_I$ eventually exceeds $|I|\lambda(\bar{\Theta})/2$ (say); therefore the bound from Hoeffding (1963) can be combined with the earlier bound from Lemma 5, for $\epsilon > 0$ sufficiently small (given α). \blacksquare

3. Efficient Computation. To calculate the boundary estimator $\widehat{\Theta}_I$ based on data $\{X_i^I : i \in I\}$, we need to calculate the criterion function $D_I(T)$

for each candidate $T \in \mathcal{T}_I$. Let C_I denote the number of simple computations (e.g., scalar addition, scalar multiplication, pairwise comparison of scalars, etc.) required to calculate $\widehat{\Theta}_I$. Essentially, we have

$$C_I \approx |\mathcal{T}_I| \times \#\{\text{computations for calculating } D_I(T) \text{ on a single } T\}.$$

The number of computations for calculating a single $D_I(T)$ seems on first thought to be of the order $|I|^2$, since we have to calculate $|I|$ distinct $d_{I_i}^T$'s, and each $d_{I_i}^T$ requires us to find the rank of X_i^I ($i \in I$) relative to the subsets $\{X_j^I : j \in \overline{T}\}$ and $\{X_j^I : j \in \underline{T}\}$. This would yield

$$C_I^0 \approx |\mathcal{T}_I| \times |I|^2,$$

where C_I^0 denotes the number of computations required to calculate $\widehat{\Theta}_I$ using the naive approach.

Algorithm. Notice that the calculation of the $d_{I_i}^T$'s involves ranking the X_i^I 's within each individual candidate region, \overline{T} and \underline{T} . For fixed T , the $\overline{h}_T^I(X_i^I)$'s ($1 \leq i \leq |I|$) can be calculated from an ordered subsequence of the ordered sequence $\{X_{(r)}^I : 1 \leq r \leq |I|\}$; similarly for the $\underline{h}_T^I(X_i^I)$'s. Thus $\widehat{\Theta}_I$ is a function only of $\{R(i), \mathbb{I}\{i \in \overline{T}\} : i \in I, T \in \mathcal{T}_I\}$, where $R(i)$ is the rank of X_i^I in the ordered sequence. The suggested computational algorithm exploits this structure.

First sort the X_i^I 's so that $X_{(1)}^I \leq X_{(2)}^I \leq \dots \leq X_{(r)}^I \leq \dots \leq X_{(|I|-1)}^I \leq X_{(|I|)}^I$. Since $S_{|I|}(\cdot)$ is symmetric [property (D.a)], its arguments $d_{I_i}^T$ may be entered according to rank order (r) rather than grid location i . The rank of $X_{(r)}^I$ is r , and its original grid index is stored in $\text{INDEX}(r)$. Now $\widehat{\Theta}_I$ is a function only of $\{\mathbb{I}\{\text{INDEX}(r) \in \overline{T}\} : 1 \leq r \leq |I|, T \in \mathcal{T}_I\}$. Sorting of the X_i^I 's and calculation of the vector $\text{INDEX}(r)$ together take at most order $|I|^2$ computations.

The collection of $\overline{h}_T^I(X_{(r)}^I)$'s (and $\underline{h}_T^I(X_{(r)}^I)$'s) is now calculated recursively, as follows. The quantity $|\overline{T}|_I \overline{h}_T^I(X_{(1)}^I)$ is 1 if the smallest observation $X_{(1)}^I$ comes from \overline{T} , and 0 if it does not. Given $|\overline{T}|_I \overline{h}_T^I(X_{(r)}^I)$, the quantity $|\overline{T}|_I \overline{h}_T^I(X_{(r+1)}^I)$ is one more than $|\overline{T}|_I \overline{h}_T^I(X_{(r)}^I)$ if $X_{(r+1)}^I$ comes from \overline{T} , and the same as $|\overline{T}|_I \overline{h}_T^I(X_{(r)}^I)$ if it does not. Formally,

$$|\overline{T}|_I \overline{h}_T^I(X_{(1)}^I) = \mathbb{I}\{\text{INDEX}(1) \in \overline{T}\},$$

$$|\bar{T}|_I \bar{h}_T^I(X_{(r+1)}^I) = |\bar{T}|_I \bar{h}_T^I(X_{(r)}^I) + \mathbb{I}\{\text{INDEX}(r+1) \in \bar{T}\}, \text{ for } r = 1, 2, \dots, |I|-1.$$

Also, $|\underline{T}|_I \underline{h}_T^I(X_{(r)}^I) = r - |\bar{T}|_I \bar{h}_T^I(X_{(r)}^I)$. Although the symbols $|\underline{T}|_I$, $|\bar{T}|_I$, and $X_{(r)}^I$ are used in the explanation above, the recursive calculations (rhs) depend only on $\mathbb{I}\{\text{INDEX}(r) \in \bar{T}\}$. The numerical values of $|\bar{T}|_I$ and $|\underline{T}|_I$ are obtained “for free”, since $|\bar{T}|_I = |\bar{T}|_I \bar{h}_T^I(X_{(|I|)}^I)$ and $|\underline{T}|_I = |I| - |\bar{T}|_I$. Thus, for fixed $T \in \mathcal{T}_I$, calculation of $D_I(T)$ requires us to go through $\{\mathbb{I}\{\text{INDEX}(r) \in \bar{T}\} : 1 \leq r \leq |I|\}$ exactly once.

Combining both phases of the algorithm, we find that the number of computations required to calculate $\hat{\Theta}_I$ is now

$$C_I^* \approx |I|^2 + |\mathcal{T}_I| \times |I|,$$

which is an order of magnitude less than the initial naive approach, i.e.,

$$C_I^*/C_I^0 \approx |\mathcal{T}_I|^{-1} + |I|^{-1}.$$

Examples. In most applications, $|\mathcal{T}_I|$ is a function of $|I|$. Table 1 gives the approximate number of computations C_I^* for several specific examples. [See C&K for precise descriptions of \mathcal{T}_I in each case.]

Table 1: Number of Simple Computations

Boundary Type Θ, T	Dimension d	# of Candidates $ \mathcal{T}_I $	# of Computations C_I^*
Change-Point	1	$ I $	$ I ^2$
Epidemic-Change	1	$ I ^2$	$ I ^3$
Template (Oriented Rectangle)	2	$ I ^2$	$ I ^3$

4. Bootstrapping. The boundary estimator $\hat{\Theta}_I$ is applicable in a broad range of situations: there are no restrictions on the distributions F and G ; the class of boundaries (Θ, \mathcal{T}_I) is quite general. In keeping with this omnibus spirit, it is appropriate to use Efron’s (1979) bootstrap for assessing the sampling variability of $\hat{\Theta}_I$. Dümbgen (1991) has studied bootstrap procedures for a related class of estimators in the one-dimensional single change-point case.

We now propose a graphical bootstrap procedure that is particularly useful in the planar ($d = 2$) case.

The true underlying probability structure is determined by the unknown triplet $(F, G; \Theta)$. Having estimated Θ with $\widehat{\Theta}_I$, the ecdf's constructed from the two regions induced by $\widehat{\Theta}_I$ provide estimates of F and G :

$$\widehat{F}_I(x) := \sum_{i \in \widehat{\Theta}_I} \mathbb{I}\{X_i^I \leq x\} / |\widehat{\Theta}_I|_I,$$

$$\widehat{G}_I(x) := \sum_{i \in \widehat{\Theta}_I} \mathbb{I}\{X_i^I \leq x\} / |\widehat{\Theta}_I|_I.$$

[Purer estimates of F and G can usually be obtained by excluding a buffer region of grid points around $\widehat{\Theta}_I$ when constructing these ecdf's.] We now have the estimated probability structure $(\widehat{F}_I, \widehat{G}_I; \widehat{\Theta}_I)$, from which bootstrap samples can be drawn. Generate $\{X_i^* : i \in \widehat{\Theta}_I\}$ as iid observations from \widehat{F}_I , and $\{X_i^* : i \in \widehat{\Theta}_I\}$ as iid observations from \widehat{G}_I . From this bootstrap sample $\{X_i^* : i \in I\}$, calculate $\widehat{\Theta}_I^*$ via the same formula as $\widehat{\Theta}_I$. By repeating this process, say B times, we obtain bootstrap realizations $\{\widehat{\Theta}_{I1}^*, \widehat{\Theta}_{I2}^*, \dots, \widehat{\Theta}_{IB}^*\}$ from $(\widehat{F}_I, \widehat{G}_I; \widehat{\Theta}_I)$. These bootstrap realizations $\widehat{\Theta}_{Ik}^*$ are used to model the sampling variability of $\widehat{\Theta}_I$.

Each $\widehat{\Theta}_{Ik}^*$ determines a partition of the grid I . If each of these $\widehat{\Theta}_{Ik}^*$'s were a number, we could look at the variance, percentiles, and histogram of these numbers in order to assess the variability of $\widehat{\Theta}_I$. But in some situations (e.g., Lipschitz curves [see C&K]), the boundaries are not readily expressible in terms of a fixed finite set of parameters. Moreover, in the two-dimensional case it is natural to want a direct visual representation of $\widehat{\Theta}_I$'s variability in \mathcal{U}_2 . For this purpose we introduce the $(1 - \alpha)100\%$ bootstrap indifference zone, $Z_{(1-\alpha)}^*$, which is defined as follows:

$$p_i^* := \sum_{k=1}^B \mathbb{I}\{i \in \widehat{\Theta}_{Ik}^*\} / B, \quad i \in I;$$

$$Z_{(1-\alpha)}^* := \{i \in I : \alpha/2 \leq p_i^* \leq 1 - \alpha/2\};$$

i.e., $Z_{(1-\alpha)}^*$ excludes those points $i \in I$ which either belong to more than $(1 - \alpha/2)100\%$ of the $\widehat{\Theta}_{Ik}^*$'s or belong to more than $(1 - \alpha/2)100\%$ of the $\widehat{\Theta}_{Ik}^*$'s.

The zone $Z_{(1-\alpha)}^*$ is conditional on the estimated probability structure $(\hat{F}_I, \hat{G}_I; \hat{\Theta}_I)$. A wide $Z_{(1-\alpha)}^*$ typically indicates high spatial variability among the $\hat{\Theta}_{Ik}^*$'s, and similarly a nonexistent or narrow $Z_{(1-\alpha)}^*$ means that the $\hat{\Theta}_{Ik}^*$'s have little spatial variability. Furthermore, if the boundary $\hat{\Theta}_I$ lies completely inside $Z_{(1-\alpha)}^*$, this suggests that the spatial "bias" of $\hat{\Theta}_I^*$ (as an "estimator" of $\hat{\Theta}_I$) is small. Finally, by analogy, we pass these conclusions about the bootstrap distribution of $\hat{\Theta}_I^*$ (variability and bias) back to the sampling distribution of $\hat{\Theta}_I$.

In practice, the entire bootstrap procedure described above is implemented on a single set of data $\{X_i^I : i \in I\}$.

Simulations. We simulated the entire bootstrap procedure 100 times for the case of a linear bisecting boundary [see C&K] in a two-dimensional 15×15 grid; the true boundary Θ connects the points (0.67,0.00) and (0.40,1.00) in \mathcal{U}_2 . For each of these 100 simulations, we obtained first $\hat{\Theta}_I$, then a picture of $Z_{(.90)}^*$, and the proportion (q^*) of grid points falling inside $Z_{(.90)}^*$; the bootstrap procedure used $B = 1000$. Figures 1.a, 1.b, 1.c show the results for three of the 100 simulations; these were selected to represent the performance of the bootstrap under "good", "average", and "poor" estimates $\hat{\Theta}_I$. Each figure illustrates $\hat{\Theta}_I$, $Z_{(.90)}^*$, and q^* .

As a basis of comparison for judging the bootstrap, we generated 1000 realizations of $\hat{\Theta}_I$ from the true $(F, G; \Theta)$. The true sampling variability of $\hat{\Theta}_I$ is thus measured by the analogous quantities $Z_{(.90)}$ and q , which are shown in Figure 2 (along with Θ).

Notice that the bootstrapped zones $Z_{(.90)}^*$ in Figure 1 are reasonably comparable in shape and size (q^*) to the true $Z_{(.90)}$ and q in Figure 2. Also note that each $Z_{(.90)}^*$ in Figure 1 contains its $\hat{\Theta}_I$, as does $Z_{(.90)}$ contain Θ in Figure 2. So it seems that the proposed graphical bootstrap procedure is useful for visually assessing the sampling behavior of $\hat{\Theta}_I$.

Figure 1: Simulated Bootstrap Indifference Zones.

$Z^*_{(.90)}$ is indicated by: $\hat{\Theta}_I$ is indicated by: _____

Figure 1.a
($q^* = .16$)





Figure 1.b
($q^* = .28$)



Acknowledgements. The authors are grateful to two referees and the editor for several helpful comments which improved this paper.

REFERENCES

- CARLSTEIN, E. (1988). Nonparametric change-point estimation. *Ann. Statist.* **16**, 188–197.
- CARLSTEIN, E., and KRISHNAMOORTHY, C. (1992). Boundary estimation. *J. Amer. Statist. Assn.* **87**, 430–438.
- DÜMBGEN, L. (1991). The asymptotic behavior of some nonparametric change-point estimators. *Ann. Statist.* **19**, 1471–1495.
- DVORETZKY, A., KIEFER, J., and WOLFOWITZ, J. (1956). Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *Ann. Math. Statist.* **27**, 642–669.
- EFRON, B. (1979). Bootstrap methods: another look at the jackknife. *Ann. Statist.* **7**, 1–26.
- HOEFFDING, W. (1963). Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assn.* **58**, 13–30.

RESEARCH COMPUTING, 2-3224
GLAXO RESEARCH INSTITUTE
RESEARCH TRIANGLE PARK, NC 27709

DEPARTMENT OF STATISTICS
UNIVERSITY OF NORTH CAROLINA
CB #3260, PHILLIPS HALL
CHAPEL HILL, NC 27599-3260