STATISTICS, OPTIMIZATION AND INFORMATION COMPUTING
Stat., Optim. Inf. Comput., Vol. **16**, July 2026, pp 581–598.
Published online in International Academic Press (www.IAPress.org)

Adaptive Generative Bootstrap: A Robust Inference Method for Small and Multimodal Samples

Saifuldeen Dheyauldeen Alrefaee*

Department of Operations Research and Intelligent Technologies, College of Computer Science and Mathematics, University of Mosul, Mosul, Iraq

Abstract The standard nonparametric bootstrap performs reliably across many settings, but small samples $n < 30$ and multimodal data tend expose its limitations under-coverage and discreetness artifacts are the most commonly reported problems. We developed an Adaptive Generative Bootstrap (AGB) method to directly address these gaps. Rather than resampling from the observed data, generative sampling is drawn from a variable-bandwidth kernel density estimate (KDE), with bandwidths adapted locally through Abramson's square root scaling. That adaptation carries a concrete theoretical payoff: under the regularity conditions set out in Section 2 and Appendix A, the leading pointwise bias of the adaptive KDE is reduced from $O\left(h^2\right)$ to $O\left(h^4\right)$. It should be noted, however, that this result pertains to the density estimator itself, not to bootstrap coverage properties directly. For that reason, the finite-sample inferential behavior of AGB is assessed empirically through Monte Carlo simulations and benchmark examples, with coverage probability, interval length, bias, and RMSE serving as the primary performance criteria.

Keywords Bootstrap, Kernel Density Estimation, Adaptive Bandwidth, Small Sample Inference, Bias Reduction.

DOI: 10.19139/soic-2310-5070-3661

1. Introduction

*S*ince its inception, the bootstrap method [1] has been widely adopted for nonparametric inference, and sampling distributions have been routinely approximated without strict parametric assumptions. Comprehensive treatments of resampling and density-based methods have been provided in [2] and [3]. Higher-order properties and comparative analyses of bootstrap confidence intervals have also been documented, including bias-corrected and accelerated (BCa) and studentized constructions [4, 5, 6, 7]. However, standard resampling from the empirical distribution has often been found to degrade in small-sample settings or when complex structure, including multimodality, has been present in the underlying density [8].

To reduce discreteness in small samples, the smoothed bootstrap has been introduced by [9]. Theoretical implications of smoothing for bootstrap accuracy and coverage have been examined in depth [10, 11, 12]. Nevertheless, performance limitations have often been induced by the use of a single global bandwidth, because peaks have been over-smoothed while tails have been under-smoothed in many practical settings [13, 14]. Alternative resampling strategies, including the Bayesian bootstrap [15] and the wild bootstrap [16], have been proposed, but geometric reconstruction in sparse regions has not been addressed directly by these approaches.

More recently, attention has been directed toward adaptive techniques. Weighted kernel estimators have been studied to improve convergence behavior [17], and new bandwidth-selection criteria have been proposed for regression contexts [18]. Additional evidence for the practical advantages of adaptive kernels has been reported in functional data settings [19]. Building on Abramson's square-root law [20] and being motivated by modern

*Correspondence to: Saifuldeen Dheyauldeen Alrefaee (Email: saifldeen.alrefaee@uomosul.edu.iq). Department of Operations Research and Intelligent Technologies, College of Computer Science and Mathematics, University of Mosul, Mosul, Iraq.

ISSN 2310-5070 (online) ISSN 2311-004X (print)

applications in dynamic density estimation [21], the Adaptive Generative Bootstrap (AGB) has been proposed in this paper. The proposed method is intended to address the coverage–length trade-off in small and irregular samples by combining classical adaptive smoothing ideas with resampling-based inference.

Recent contribution in Statistics, Optimization and Information Computing further demonstrate the continuing relevance of kernel-based and nonparametric inference in complex statistical settings. Bouzebda et al. [30] studied nonparametric recursive kernel-type estimators for the moment-generating function under censored data and established asymptotic bias, variance, and central limit results, together with a finite-sample simulation assessment. Slama et al. [31] developed a multivariate recursive estimator for conditional cumulative distribution functions using stochastic approximation, emphasizing the roles of bandwidth and step size selection and showing that the resulting mean squared error can improve over the classical Nadaraya-Watson estimator under suitable conditions. These studies reinforce the importance of combining nonparametric smoothing, careful bandwidth control, theoretical regularity, and empirical validation when inference is conducted under small-sample or distributionally complex conditions.

Motivated by this line of work, we developed the Adaptive Generative Bootstrap (AGB) as a kernel-based generative resampling procedure that uses local bandwidth adaptation to improve empirical interval calibration in small-sample and distributionally irregular settings.

This paper presents the structured framework as follows: Section 2 lays out the theoretical foundation, which also addresses the algorithmic details of AGB. Section 3 presents the design of the simulations performed and also presents the experimental results compared with those of standard methods. Section 4 presents the conclusions and offers some practical recommendations.

2. Methodology

Let $X_1, \ldots, X_n$ be a random sample from an unknown probability density function f. A fixed-bandwidth kernel density estimator has commonly been defined as $\hat{f}(x) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{x-X_i}{h}\right)$. Although broad applicability has been achieved with this estimator, a global bias–variance compromise has typically been imposed. To mitigate this limitation, a variable kernel density estimator (VKDE) framework has been adopted in what follows. Standard treatments of variable adaptive kernel density estimation and kernel smoothing have been provided, and closely related formulations have been studied for variable kernel estimators and variable location–scale approaches [22, 23, 24, 25, 26, 27].

2.1. Bootstrap Methods

For B bootstrap samples X_b^* drawn with replacement from X, the bootstrap estimator is:

$$\hat{\theta}_B = \frac{1}{B}\sum_{b=1}^{B} T\left(X_b^*\right) \tag{1}$$

with bootstrap standard error: [3, 27]

$$SE_B = \sqrt{\frac{1}{B-1}\sum_{b=1}^{B}\left(T\left(X_b^*\right) - \bar{T}_B\right)^2} \tag{2}$$

where $\bar{T}_B = \frac{1}{B}\sum_{b=1}^{B} T\left(X_b^*\right)$.

However, in small sample sizes $n < 30$, bias has often been induced in the standard error estimate in Eq. 2, and under-coverage has frequently been produced by the resulting confidence intervals. This limitation has been used to motivate the AGB procedure that is described next. Coverage properties of commonly used bootstrap confidence intervals *e.g.*, percentile, studentized, and BCa have been compared in the literature, and systematic under-coverage in small samples has been reported in several settings [5, 6, 7, 28].

2.2. Adaptive Bandwidth Construction

Following Abramson's square root law [20], an adaptive bandwidth h_i has been constructed for each observation X_i as:

$$h_i = h_{\text{global}} \left(\frac{\widetilde{f}(X_i)}{g} \right)^{-1/2} \tag{3}$$

where $\widetilde{f}(X_i)$ has denoted a pilot density estimate, and g has denoted the geometric mean of the pilot densities, with $\log g = n^{-1} \sum_{i=1}^{n} \log \widetilde{f}(X_i)$. By this construction, smaller bandwidths have been assigned in dense regions (so that bias has been reduced), while larger bandwidths have been assigned in sparse regions (so that variance has been stabilized).

2.3. The AGB Algorithm

The proposed Adaptive Generative Bootstrap algorithm has been carried out as follows:

1. **Pilot estimation:** A pilot density $\widetilde{f}$ is computed using a rule-of-thumb bandwidth [2], and bootstrap-oriented bandwidth selection has been discussed in [29].
2. **Adaptation:** Local bandwidth factors are computed as $\lambda_i = \left(\widetilde{f}(X_i)/g \right)^{-1/2}$, and local bandwidths are set as $h_i = h_{pilot} \lambda_i$.
3. **Generative resampling:** One bootstrap sample X_b^* of size n is generated as follows:
   - For $k = 1, \ldots, n$, an index J is selected uniformly from $\{1, \ldots, n\}$.
   - A noise term ϵ is sampled from the kernel K (e.g., $N(0,1)$).
   - A bootstrap draw is formed as $X_{bk}^* = X_J + h_J \epsilon$.
4. **Inference:** The steps above are repeated for $b = 1, \ldots, B$ (in this study, $B = 2,000$) so that a bootstrap distribution $\widehat{F}_{AGB}^*$ is obtained and percentile confidence intervals are computed (e.g., the 2.5th and 97.5th percentiles).

2.4. Theoretical Properties

The following theoretical results concern the adaptive density estimation component used within the AGB resampling mechanism. They provide motivation for the proposed generative construction but should not be interpreted as direct theoretical guarantees for the coverage probability or bias of a bootstrap estimator of a general statistic $T(\cdot)$.

Theorem 1 (Bias Reduction of the Adaptive Density Estimator Underlying AGB): Let f be strictly positive in a neighborhood of x and $f \in C^4$. Let K be a symmetric kernel with $\int K(u)\,du = 1$, $\int uK(u)\,du = 0$, and finite moments $\mu_2(K) = \int u^2 K(u)\,du < \infty$, $\mu_4(K) = \int u^4 K(u)\,du < \infty$. Consider the sample-point adaptive KDE

$$\widehat{f}_{\text{adapt}}(x) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{h\,\lambda(X_i)} K\left(\frac{x - X_i}{h\,\lambda(X_i)} \right) \tag{4}$$

where $\lambda(\cdot) \in C^3$, $\lambda(x) > 0$, and $h \to 0$ with $nh \to \infty$. If $\lambda(x) = c\,f(x)^{-1/2}$ (Abramson's square root law) for some constant $c > 0$, then the pointwise bias satisfies

$$\mathbb{E}\left[\widehat{f}_{\text{adapt}}(x) \right] - f(x) = O\left(h^4\right) \tag{5}$$

In contrast, the standard fixed-bandwidth KDE has pointwise bias of order $O\left(h^2\right)$.

The scope of Theorem 1 should be interpreted carefully. The result concerns the pointwise bias of the adaptive kernel density estimator used to generate AGB resamples. It does not directly establish a corresponding bias

reduction for a bootstrap estimator of a statistic $T(\cdot)$, such as the sample mean, nor does it provide a direct theoretical guarantee of improved percentile-interval coverage. Rather, it provides theoretical motivation for using adaptive local smoothing as a generative approximation to the underlying distribution. The extent to which this approximation improves finite sample bootstrap inference is assessed empirically in the simulation study.

Theorem 2 (Asymptotic Variance): Under the regularity conditions in Appendix A,

$$\mathrm{Var}\left[\widehat{f}_{\text{adapt}}(x)\right] = \frac{1}{nh\,\lambda(x)} f(x) R(K) + o\left(\frac{1}{nh}\right), \qquad R(K) = \int K(u)^2\,du \tag{6}$$

Hence,

$$\mathrm{Var}\left[\widehat{f}_{\text{adapt}}(x)\right] = O\left((nh)^{-1}\right) \tag{7}$$

The proofs for both theorems are provided in the Appendix A.

3. Results and Discussion

To evaluate the finite sample behavior of AGB, a comprehensive Monte Carlo simulation study was conducted and complemented with benchmark data examples. The theoretical results in Section 2 motivate the adaptive generative mechanism, whereas coverage probability, interval length, and estimation accuracy are assessed empirically.

3.1. Performance Metrics and Interval Construction

Let θ denote the estimand of interest, which is the population mean in the simulation study. For each Monte Carlo replication $(r = 1, \ldots, M)$, let $\hat{\theta}_r^{*(b)}$, for $b = 1, \ldots, B$, denote the bootstrap replicate estimates obtained from a given resampling procedure. The corresponding $100(1-\alpha)\%$ percentile confidence interval is defined as

$$L_r = Q_{\alpha/2}\left(\left\{\widehat{\theta}_r^{*(b)}\right\}_{b=1}^{B}\right) \tag{8}$$

$$U_r = Q_{1-\alpha/2}\left(\left\{\widehat{\theta}_r^{*(b)}\right\}_{b=1}^{B}\right) \tag{9}$$

where $Q_p(\cdot)$ denotes the empirical (p)-quantile of the bootstrap replicate estimates. The empirical coverage probability (CP) is computed as

$$CP = \frac{1}{M}\sum_{r=1}^{M} \mathbb{I}\left(\theta \in [L_r, U_r]\right) \tag{10}$$

To quantify Monte Carlo uncertainty, the Monte Carlo standard error of the coverage estimate is reported as

$$MCSE(CP) = \sqrt{\frac{CP(1-CP)}{M}} \tag{11}$$

The average interval length is defined as

$$\bar{L} = \frac{1}{M}\sum_{r=1}^{M}(U_r - L_r) \tag{12}$$

and its Monte Carlo standard error is computed as the standard deviation of the interval lengths across replications divided by $\sqrt{M}$. Point estimation accuracy was also evaluated using the bias and root mean squared

error (RMSE) of the bootstrap based point estimate. Let

$$\hat{\theta}_{r,\text{boot}} = \frac{1}{B}\sum_{b=1}^{B}\hat{\theta}_r^{*(b)} \tag{13}$$

Then

$$Bias = \frac{1}{M}\sum_{r=1}^{M}\left(\hat{\theta}_{r,\text{boot}} - \theta\right) \tag{14}$$

and

$$RMSE = \left\{\frac{1}{M}\sum_{r=1}^{M}\left(\hat{\theta}_{r,\text{boot}} - \theta\right)^2\right\}^{1/2} \tag{15}$$

The RMSE of the estimated standard error was also computed relative to the known Monte Carlo target $(SE_{\text{true}} = \sigma/\sqrt{n})$ in each simulation scenario.

3.2. Expanded Comparative Simulation

The revised simulation study compares four resampling procedures: the standard empirical bootstrap, the globally smoothed bootstrap, the Bayesian bootstrap, and the proposed AGB method. The comparison was conducted under five data-generating scenarios representing regular and irregular settings: normal, heavy-tailed t_3, centered lognormal, bimodal mixture, and trimodal mixture distributions. The main small-sample results are reported for $n = 15$, with $C = 500$ Monte Carlo replications and $B = 2000$ bootstrap resamples per replication. The corresponding results for $n = 30$ are reported in the supplementary material.

Table 1. Expanded comparative simulation results for small sample inference ($n = 15,\ M = 500,\ B = 2000$).

Scenario	Method	CP	MCSE(CP)	Avg. Length	Bias	RMSE	RMSE(SE)
Normal	Standard Bootstrap	0.920	0.0121	0.9447	-0.0045	0.2551	0.0500
Normal	Global Smoothed Bootstrap	0.940	0.0106	1.0512	-0.0047	0.2549	0.0541
Normal	Bayesian Bootstrap	0.902	0.0133	0.9207	-0.0043	0.2549	0.0516
Normal	AGB	0.936	0.0109	1.0627	-0.0049	0.2549	0.0553
Heavy-tailed (t_3)	Standard Bootstrap	0.896	0.0137	1.5297	0.0289	0.4901	0.2347
Heavy-tailed (t_3)	Global Smoothed Bootstrap	0.920	0.0121	1.6453	0.0288	0.4925	0.2342
Heavy-tailed (t_3)	Bayesian Bootstrap	0.872	0.0149	1.5096	0.0299	0.4928	0.2313
Heavy-tailed (t_3)	AGB	0.926	0.0117	1.6625	0.0287	0.4925	0.2299
Skewed lognormal	Standard Bootstrap	0.854	0.0158	1.6587	-0.0250	0.5113	0.2986
Skewed lognormal	Global Smoothed Bootstrap	0.870	0.0150	1.7578	-0.0250	0.5107	0.2906
Skewed lognormal	Bayesian Bootstrap	0.844	0.0162	1.6282	-0.0244	0.5111	0.2965
Skewed lognormal	AGB	0.874	0.0148	1.7777	-0.0242	0.5149	0.2919
Bimodal mixture	Standard Bootstrap	0.918	0.0123	2.1544	-0.0525	0.5840	0.0695
Bimodal mixture	Global Smoothed Bootstrap	0.936	0.0109	2.4412	-0.0521	0.5834	0.0897
Bimodal mixture	Bayesian Bootstrap	0.910	0.0128	2.0788	-0.0514	0.5834	0.0759
Bimodal mixture	AGB	0.940	0.0106	2.4516	-0.0530	0.5833	0.0907
Trimodal mixture	Standard Bootstrap	0.922	0.0120	2.4649	0.0324	0.6657	0.0808
Trimodal mixture	Global Smoothed Bootstrap	0.948	0.0099	2.7967	0.0329	0.6679	0.1080
Trimodal mixture	Bayesian Bootstrap	0.920	0.0121	2.3886	0.0322	0.6668	0.0860
Trimodal mixture	AGB	0.954	0.0094	2.8091	0.0328	0.6677	0.1090

Note. CP denotes empirical coverage probability; MCSE(CP) is the Monte Carlo standard error of CP; Avg. Length is the average 95% percentile confidence-interval length; Bias and RMSE refer to the bootstrap-based point estimate of the mean; RMSE(SE) denotes the RMSE of the estimated standard error. Simulations used $M = 500$ Monte Carlo replications and $B = 2000$ bootstrap resamples per replication.

Table 1 reports the expanded small sample comparison at $n = 15$. Across the examined scenarios, AGB generally produced higher empirical coverage than the standard and Bayesian bootstrap methods in several small-sample irregular settings, particularly in the heavy-tailed and multimodal scenarios. The globally smoothed bootstrap also improved coverage relative to the empirical bootstrap, indicating that part of the improvement is attributable to

smoothing. However, AGB was generally comparable to the globally smoothed bootstrap and, in some irregular settings, produced slightly higher empirical coverage. This suggests that local bandwidth adaptation may provide an additional calibration benefit, although the gain is not uniform across all scenarios.

The coverage improvement was accompanied by moderately wider confidence intervals. This confirms that AGB should not be interpreted as uniformly dominating competing methods, but rather as improving interval calibration through a coverage length trade-off. The RMSE values of the bootstrap-based point estimates were broadly comparable across methods, indicating that the improved coverage was not associated with a substantial loss in point-estimation accuracy. In the skewed lognormal setting, all percentile-based methods exhibited under-coverage, suggesting that strong skewness remains challenging for percentile interval construction even under adaptive smoothing.

Thus, the empirical advantage of AGB is best interpreted as improved interval calibration in selected small sample and irregular settings, rather than as uniform superiority over all competing resampling methods.

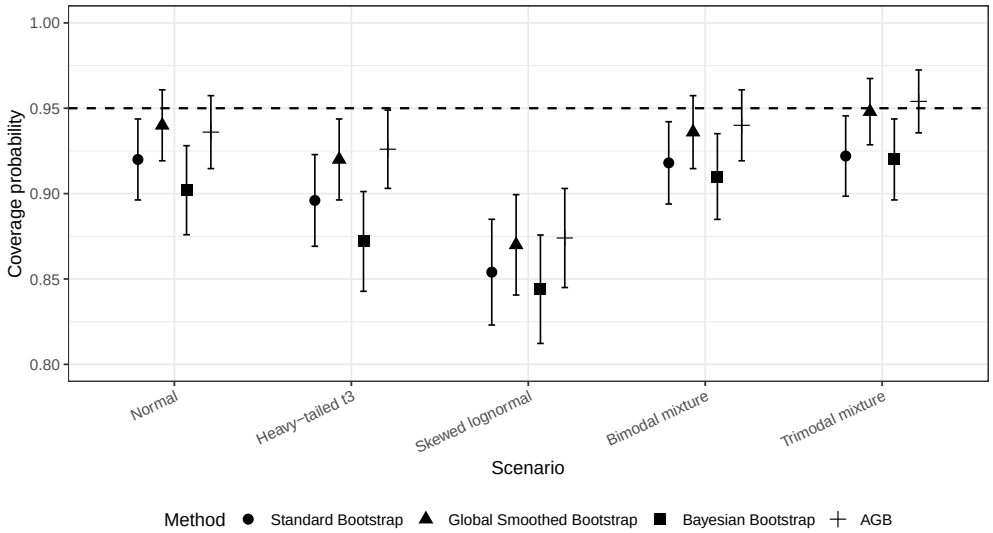

Figure 1. Empirical coverage probability across small-sample scenarios at $n = 15$. The dashed horizontal line denotes the nominal 95% coverage level. Error bars represent approximate 95% Monte Carlo uncertainty intervals based on $\pm 1.96\,\mathrm{MCSE}$. Results are based on $M = 500$ Monte Carlo replications and $B = 2000$ bootstrap resamples per replication.

3.3. Sensitivity Analysis

The robustness of AGB to key tuning and implementation choices was probed through a set of targeted sensitivity checks. The pilot bandwidth was scaled by multiplying the rule-of-thumb bandwidth by factors of 0.5, 1, 2, and 5. The impact of truncating the local bandwidth factor λ was then assessed under both fixed and data-dependent clipping rules. The kernel used in the generative perturbation step was also varied Gaussian and Epanechnikov kernels were directly compared. All of these checks were carried out to determine whether the observed behavior of AGB is tied to any specific tuning convention or implementation decision.

Table 2 reports the sensitivity of AGB to the pilot bandwidth in small samples. The results show a clear coverage length trade-off. Smaller pilot bandwidths generally produced shorter intervals but were more likely to under cover, particularly in the heavy-tailed and skewed settings. Increasing the pilot multiplier improved coverage, but this improvement was obtained through wider intervals. The multiplier-5 configuration frequently produced nearly conservative or fully conservative coverage, but at the cost of substantially inflated interval length, especially in the bimodal and trimodal mixtures. Therefore, the default rule-of-thumb bandwidth provides a practical compromise

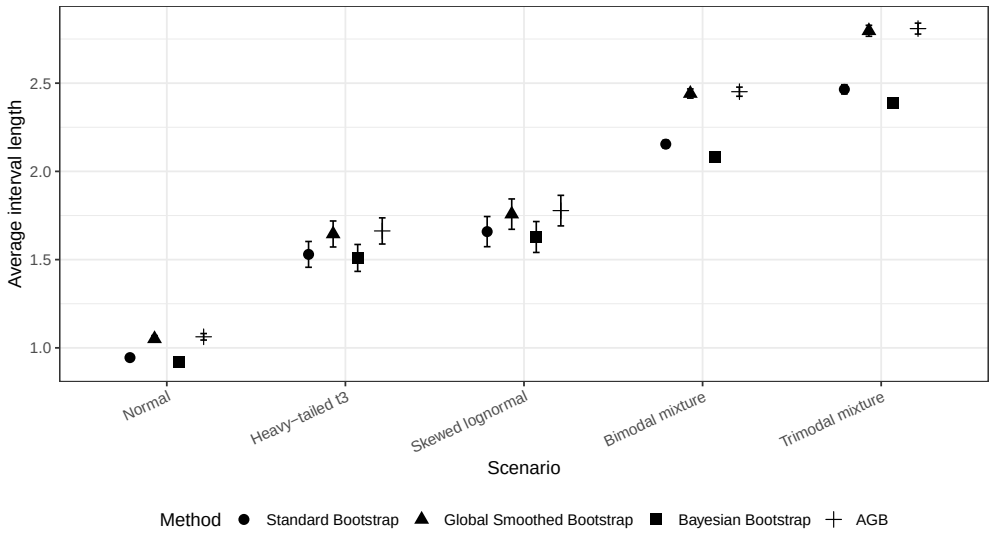

Figure 2. Average 95% percentile confidence-interval length across small-sample scenarios at $n = 15$. Error bars represent approximate 95% Monte Carlo uncertainty intervals based on $\pm 1.96\,\mathrm{MCSE}$. Results are based on $M = 500$ Monte Carlo replications and $B = 2000$ bootstrap resamples per replication.

Table 2. Pilot-bandwidth sensitivity analysis for AGB in small samples ($n = 15,\ M = 500,\ B = 2000$).

Scenario	Pilot multiplier	CP	MCSE(CP)	Avg. Length	MCSE(Length)
Heavy-tailed (t_3)	0.5	0.896	0.0137	1.5156	0.0299
Heavy-tailed (t_3)	1	0.954	0.0094	1.6591	0.0293
Heavy-tailed (t_3)	2	0.970	0.0076	2.0087	0.0339
Heavy-tailed (t_3)	5	1.000	0.0000	3.4899	0.0474
Skewed lognormal	0.5	0.846	0.0161	1.7164	0.0440
Skewed lognormal	1	0.830	0.0168	1.7518	0.0434
Skewed lognormal	2	0.922	0.0120	2.1931	0.0467
Skewed lognormal	5	0.964	0.0083	3.5154	0.0680
Bimodal mixture	0.5	0.934	0.0111	2.2333	0.0119
Bimodal mixture	1	0.932	0.0113	2.4330	0.0139
Bimodal mixture	2	0.982	0.0059	3.2048	0.0187
Bimodal mixture	5	1.000	0.0000	6.3940	0.0370
Trimodal mixture	0.5	0.922	0.0120	2.5533	0.0152
Trimodal mixture	1	0.948	0.0099	2.7838	0.0165
Trimodal mixture	2	0.980	0.0063	3.6392	0.0211
Trimodal mixture	5	1.000	0.0000	7.1377	0.0452

Note. The pilot multiplier scales the rule-of-thumb pilot bandwidth used in the adaptive density-estimation step of AGB. CP denotes empirical coverage probability, and MCSE denotes Monte Carlo standard error. Larger pilot bandwidths generally increase coverage by widening the resulting percentile confidence intervals.

between interval calibration and interval width in most settings, while very large pilot multipliers should be used cautiously.

The corresponding RMSE and RMSE(SE) values are reported in Supplementary Table S2 and show that the main effect of increasing the pilot bandwidth is on interval width and estimated standard error rather than on the bootstrap based point estimate.

Because the practical implementation of AGB requires several finite-sample tuning choices that are not specified by the asymptotic Abramson square root law, additional sensitivity analyses were conducted. First, the influence of truncating the local scaling factor λ was investigated by comparing no clipping, two fixed clipping intervals, and a data-dependent quantile-based clipping rule. Second, the role of kernel choice in the generative perturbation step was evaluated by comparing Gaussian and Epanechnikov kernels. These analyses were intended to assess whether the empirical behavior of AGB is robust to implementation choices or depends critically on a specific computational convention.

Table 3. Implementation sensitivity analysis for AGB in small samples ($n = 15,\ M = 500,\ B = 2000$).

Panel	Scenario	Configuration	CP	MCSE(CP)	Avg. Length	MCSE(Length)
A: λ-clipping	Heavy-tailed (t_3)	No clipping	0.952	0.0096	1.6615	0.0301
A: λ-clipping	Heavy-tailed (t_3)	Fixed ($[0.1, 5.0]$)	0.942	0.0105	1.6294	0.0302
A: λ-clipping	Heavy-tailed (t_3)	Fixed ($[0.2, 3.0]$)	0.952	0.0096	1.5882	0.0282
A: λ-clipping	Heavy-tailed (t_3)	($[Q_{0.05}, Q_{0.95}]$)	0.950	0.0097	1.6404	0.0264
A: λ-clipping	Skewed lognormal	No clipping	0.848	0.0161	1.7873	0.0409
A: λ-clipping	Skewed lognormal	Fixed ($[0.1, 5.0]$)	0.842	0.0163	1.9266	0.0540
A: λ-clipping	Skewed lognormal	Fixed ($[0.2, 3.0]$)	0.852	0.0159	1.8948	0.0626
A: λ-clipping	Skewed lognormal	($[Q_{0.05}, Q_{0.95}]$)	0.862	0.0154	1.8556	0.0483
A: λ-clipping	Bimodal mixture	No clipping	0.954	0.0094	2.4689	0.0144
A: λ-clipping	Bimodal mixture	Fixed ($[0.1, 5.0]$)	0.950	0.0097	2.4709	0.0137
A: λ-clipping	Bimodal mixture	Fixed ($[0.2, 3.0]$)	0.952	0.0096	2.4804	0.0141
A: λ-clipping	Bimodal mixture	($[Q_{0.05}, Q_{0.95}]$)	0.956	0.0092	2.4334	0.0143
B: Kernel	Heavy-tailed (t_3)	Gaussian	0.940	0.0106	1.6532	0.0308
B: Kernel	Heavy-tailed (t_3)	Epanechnikov	0.944	0.0103	1.6973	0.0405
B: Kernel	Skewed lognormal	Gaussian	0.876	0.0147	1.9382	0.1044
B: Kernel	Skewed lognormal	Epanechnikov	0.892	0.0139	1.8606	0.0465
B: Kernel	Bimodal mixture	Gaussian	0.944	0.0103	2.4614	0.0136
B: Kernel	Bimodal mixture	Epanechnikov	0.954	0.0094	2.4476	0.0138
B: Kernel	Trimodal mixture	Gaussian	0.958	0.0090	2.7931	0.0155
B: Kernel	Trimodal mixture	Epanechnikov	0.946	0.0101	2.8056	0.0163

Note. Panel A compares different rules for truncating the adaptive local bandwidth factor λ, where $[Q_{0.05}, Q_{0.95}]$ denotes quantile-based clipping. Panel B compares Gaussian and Epanechnikov kernels in the generative perturbation step. CP denotes empirical coverage probability, and MCSE denotes Monte Carlo standard error. The full implementation-sensitivity results, including $n = 30$, MCSE values, Bias, RMSE, RMSE(SE), and summaries of the empirical λ distribution, are reported in Supplementary Tables S3–S4.

The implementation sensitivity analyses for AGB are summarized in Table 3. Panel A shows that the empirical performance of AGB held up well across different λ clipping rules. In the heavy-tailed and bimodal settings, coverage produced by the default fixed clipping role $[0.2, 3.0]$ remained close to the nominal level, with interval lengths that were broadly comparable to those obtained under no clipping, wider fixed bounds, and quantile-based clipping. This points to the default clipping rule functioning primarily as a finite-sample stabilization device not as a mechanism that artificially inflates the reported coverage gains.

Panel B shows that broadly similar coverage and interval length behavior was observed when the Gaussian perturbation kernel was replaced with an Epanechnikov kernel. In the heavy-tailed and multimodal settings, coverage close to the nominal level was produced by both kernels. The skewed lognormal setting, by contrast, remained difficult under all clipping rules and both kernels which suggests that the observed under-coverage is driven mainly by percentile interval construction under strong skewness, rather than by any particular implementation choice

3.4. Illustrative Applications to Benchmark Data

To complement the simulation study, the four resampling procedures were applied to two benchmark datasets. These examples are not meant to serve as evidence of coverage performance that is simply not assessable when the true population mean is unknown, as it is in any real data application. What they do offer is a concrete look at how the methods behave in practice. The comparison is therefore restricted to the bootstrap point estimate. The estimated standard error, the 95% percentile confidence interval, and the corresponding interval length.

For the Law School dataset $n = 15$, AGB produced a wider confidence interval than the empirical and Bayesian bootstrap methods, with an interval length similar to that of the globally smoothed bootstrap. This behavior is consistent with the simulation results, where smoothing based methods tended to improve interval calibration at the cost of wider intervals. However, because the true population mean is unknown, these results should be interpreted as an illustration of interval behavior rather than as evidence of improved coverage.

The Old Faithful dataset $n = 272$ is retained only as a large sample multimodal illustration. As expected, the four methods produced broadly similar point estimates and percentile intervals, with AGB giving a slightly wider interval. This example is not used as primary evidence for small sample improvement, but it illustrates that the adaptive generative procedure remains numerically stable in a larger empirical dataset.

Table 4. Illustrative real data comparison of bootstrap percentile intervals ($B = 2000$).

Dataset	n	Method	Estimate	SE	95% Percentile CI	Length
Law School	15	Standard Bootstrap	600.2821	10.3081	[580.8667, 621.6050]	40.7383
Law School	15	Global Smoothed Bootstrap	600.5420	12.1825	[577.8811, 625.0224]	47.1413
Law School	15	Bayesian Bootstrap	600.3029	9.9652	[582.1513, 620.8438]	38.6925
Law School	15	AGB	600.2341	11.9993	[576.8121, 624.2577]	47.4456
Old Faithful	272	Standard Bootstrap	70.9009	0.8220	[69.3199, 72.5408]	3.2210
Old Faithful	272	Global Smoothed Bootstrap	70.8989	0.8243	[69.3006, 72.5119]	3.2113
Old Faithful	272	Bayesian Bootstrap	70.8912	0.8172	[69.3422, 72.4514]	3.1092
Old Faithful	272	AGB	70.9155	0.8601	[69.1956, 72.5924]	3.3968

Note. SE denotes the bootstrap-estimated standard error of the mean. The real data examples are intended as illustrative applications only; coverage probability and RMSE relative to the true mean are not reported because the population mean is unknown.

4. Conclusion

The theoretical results establish the bias and variance behavior of the adaptive density estimation component that underlies AGB, and it is those results that motivate the resampling mechanism itself. The finite-sample inferential behavior coverage probability, interval length and estimation accuracy is then put to the test through the simulation study. In small-sample scenarios, under-coverage is reduced by AGB across a range of irregular settings, though this tends to come at the cost of somewhat wider percentile intervals. The sensitivity analysis further supports the practical stability of the procedure: coverage gains are most pronounced in small samples, which is precisely where the method was designed to matter, and the overall pattern holds reasonably well at larger sample sizes. Taken together, AGB appears to be a workable resampling alternative for limited or distributionally irregular datasets, particularly when interval calibration is the priority and a moderate increase in interval length is an acceptable trad-off.

Appendix

Appendix A: Mathematical Properties of the Adaptive Density Estimator Underlying AGB

The adaptive kernel density estimator used as the generative component of AGB is the focus of this appendix. No direct theoretical guarantees are made here for the coverage probability or bias of a bootstrap estimator of a general statistic. Rather, theoretical motivation is provided for the adaptive density estimation step that is employed in the proposed resampling procedure.

A.1 Assumptions:

Assume, in a neighborhood of x, that:

1. The density f is strictly positive and belongs to C^4.
2. The scaling function $\lambda(\cdot)$ belongs to C^3 and is bounded away from zero and infinity, that is,
$$0 < \inf \lambda \leq \sup \lambda < \infty.$$
3. The kernel K is symmetric, satisfies
$$\int K(u)\, du = 1,$$
and has a finite second moment and a finite L^2-norm:
$$\mu_2(K) = \int u^2 K(u)\, du < \infty, \qquad \int K^2(u)\, du < \infty.$$
4. The bandwidth satisfies
$$h \to 0 \qquad \text{and} \qquad nh \to \infty \qquad \text{as } n \to \infty.$$

In particular, under Abramson's square root law,

$$\lambda(x) = c\, f(x)^{-1/2},$$

Assumptions 1–2 ensure that λ is locally differentiable twice continuously.

A.2 Proof of Theorem 1 (Bias Reduction of the Adaptive Density Estimator)

Consider the adaptive kernel density estimator

$$\hat{f}(x) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{h\lambda(X_i)} K\left(\frac{x - X_i}{h\lambda(X_i)} \right).$$

Its expectation is

$$E[\hat{f}(x)] = \int \frac{1}{h\lambda(y)} K\left(\frac{x - y}{h\lambda(y)} \right) f(y)\, dy.$$

Under the standard regularity conditions for variable-bandwidth kernel density estimators, the bias admits the asymptotic expansion

$$E[\hat{f}(x)] - f(x) = \frac{1}{2} h^2 \mu_2(K) \left[f(x)\lambda(x)^2 \right]'' + O(h^4),$$

where

$$\mu_2(K) = \int u^2 K(u)\, du.$$

Let

$$B(x) = f(x)\lambda(x)^2.$$

Then

$$E[\hat{f}(x)] - f(x) = \frac{1}{2}h^2\mu_2(K)\,B''(x) + O(h^4).$$

Differentiating $B(x)$ twice gives

$$B''(x) = f''(x)\lambda(x)^2 + 4f'(x)\lambda(x)\lambda'(x) + 2f(x)\big(\lambda'(x)\big)^2 + 2f(x)\lambda(x)\lambda''(x).$$

Under Abramson's square root law,

$$\lambda(x) = c\,f(x)^{-1/2},$$

for some positive constant c. Hence

$$\lambda(x)^2 = c^2 f(x)^{-1},$$

so that

$$B(x) = f(x)\lambda(x)^2 = c^2.$$

Therefore,

$$B''(x) = 0.$$

Substituting this into the bias expansion yields

$$E[\hat{f}(x)] - f(x) = O(h^4).$$

Hence, under Abramson's square root law, the $O(h^2)$ bias term vanishes, and the leading bias is reduced to order $O(h^4)$. □

A.3 Proof of Theorem 2 (Asymptotic Variance)

Define

$$Z_i(x) = \frac{1}{h\lambda(X_i)}K\left(\frac{x - X_i}{h\lambda(X_i)}\right), \qquad i = 1, \ldots, n.$$

Then

$$\hat{f}(x) = \frac{1}{n}\sum_{i=1}^{n} Z_i(x),$$

where $Z_1(x), \ldots, Z_n(x)$ are i.i.d. random variables. Therefore,

$$\mathrm{Var}\big(\hat{f}(x)\big) = \frac{1}{n}\,\mathrm{Var}\big(Z_1(x)\big).$$

Now,

$$\mathrm{Var}\big(Z_1(x)\big) = E\big[Z_1(x)^2\big] - \big(E[Z_1(x)]\big)^2.$$

Moreover,

$$E[Z_1(x)] = f(x) + o(1),$$

so that

$$\big(E[Z_1(x)]\big)^2 = O(1).$$

On the other hand, the leading term of $E[Z_1(x)^2]$ is of order $O(h^{-1})$, so

$$\mathrm{Var}\big(\hat{f}(x)\big) = \frac{1}{n}E\big[Z_1(x)^2\big] + o\left(\frac{1}{nh}\right).$$

Next,

$$E\big[Z_1(x)^2\big] = \int \frac{1}{h^2\lambda(y)^2}K^2\left(\frac{x-y}{h\lambda(y)}\right)f(y)\,dy.$$

Using a local expansion around $y = x$, together with the continuity of f and λ, we may write

$$f(y) = f(x) + o(1), \qquad \lambda(y) = \lambda(x) + o(1), \qquad h \to 0.$$

Then, under the change of variable

$$u = \frac{x - y}{h\lambda(x)}, \qquad dy = -h\lambda(x)\, du,$$

it follows that

$$E\big[Z_1(x)^2\big] = \frac{f(x)}{h\lambda(x)} \int K^2(u)\, du + o\left(\frac{1}{h}\right).$$

Let

$$R(K) = \int K^2(u)\, du.$$

Then

$$E\big[Z_1(x)^2\big] = \frac{f(x)}{h\lambda(x)} R(K) + o\left(\frac{1}{h}\right).$$

Therefore,

$$\mathrm{Var}\big(\hat{f}(x)\big) = \frac{f(x)}{nh\lambda(x)} R(K) + o\left(\frac{1}{nh}\right).$$

Since $\lambda(x)$ is locally bounded and bounded away from zero,

$$\mathrm{Var}\big(\hat{f}(x)\big) = O\big((nh)^{-1}\big).$$

This establishes the stated asymptotic variance. □

Appendix B: Computational Implementation and Reproducibility Details

This appendix provides the core computational details for implementing the Adaptive Generative Bootstrap (AGB) algorithm used in this study. To ensure reproducibility without relying on an external repository link, the full simulation scripts, sensitivity analyses, figure-generation code, and real-data analysis scripts are provided as static supplementary code files accompanying this manuscript. The core AGB function was implemented in R, and all numerical results in the simulation study were obtained using $M = 500$ Monte Carlo replications and $B = 2000$ bootstrap resamples per replication.

B.1 Core AGB Algorithm

AGB is built on three interlocking pieces: a pilot density estimate, Abramson-style local bandwidth adaptation, and generative bootstrap resampling. Starting from an observed sample $(X_1, \ldots, X_n)$, a pilot density estimate $\widetilde{f}$ is computed first the rule-of-thumb bandwidth handles this by default. The local scaling factor is then defined as

$$\lambda_i = \left(\frac{\widetilde{f}(X_i)}{g}\right)^{-1/2}$$

where

$$\log g = \frac{1}{n} \sum_{i=1}^{n} \log \widetilde{f}(X_i)$$

The local bandwidth is given by

$$h_i = h_{\mathrm{pilot}} \lambda_i$$

A generative bootstrap observation is then produced as

$$X_{bk}^{*} = X_J + h_J \varepsilon_k$$

where J is sampled uniformly from $\{1, \ldots, n\}$, and ε_k is drawn from a symmetric kernel distribution. In the default implementation the Gaussian disturbances are used. The core algorithm for the Adaptive Generative Bootstrap (AGB) is implemented in R. The static supplementary code files accompanying this manuscript provide the full implementation details.

Listing 1: Core Function of AGB Method

```
agb_core <- function(data,
                     n_boot = 2000,
                     alpha = 0.05,
                     pilot_mult = 1,
                     lambda_lower = 0.2,
                     lambda_upper = 3.0) {
    n <- length(data)
    # 1. Pilot density estimate
  d <- density(data, bw = "nrd0")
  f_x <- approx(d$x, d$y, xout = data, rule = 2)$y
  f_x <- pmax(f_x, 1e-10)
    # 2. Geometric mean of pilot densities
  g <- exp(mean(log(f_x)))
    # 3. Abramson local scaling
  lambda <- (f_x / g)^(-0.5)
    # 4. Finite-sample stabilization
  lambda <- pmin(pmax(lambda, lambda_lower), lambda_upper)
    # 5. Local bandwidths
  h_pilot <- d$bw * pilot_mult
  h_local <- h_pilot * lambda
    # 6. Generative bootstrap samples
  indices <- sample(seq_len(n), n_boot * n, replace = TRUE)
  noise <- rnorm(n_boot * n)
  boot_vals <- data[indices] + h_local[indices] * noise
    # 7. Bootstrap replicate means
  boot_means <- rowMeans(matrix(boot_vals, nrow = n_boot))
    # 8. Percentile confidence interval
  ci <- quantile(boot_means, probs = c(alpha / 2, 1 - alpha / 2))
    return(list(
    estimate = mean(boot_means),
    se = sd(boot_means),
    lower = ci[1],
    upper = ci[2],
    length = ci[2] - ci[1]
  ))
}
```

B.2 Implementation Choices

The R rule-of-thumb bandwidth `bw.nrd0` is used as the pilot bandwidth in the default implementation. The local scaling factor λ is truncated to the interval $[0.2, 3.0]$. This truncation is not derived from the asymptotic Abramson square root law it is applied as a finite-sample stabilization measure to prevent extremely small or large local bandwidths from being produced when the pilot density estimate is unreliable, particularly in small samples.

To determine whether the empirical results are materially affected by this choice, a dedicated sensitivity analysis was carried out. Four clipping strategies were compared: no clipping, fixed clipping at $[0.1, 5.0]$, the default fixed clipping at $[0.2, 3.0]$, and a data-dependent quantile rule based on the empirical 5th and 95th percentiles of λ. The full results are reported in Appendix D.

The Gaussian disturbances are used by default in the generative step, but both Gaussian and Epanechnikov kernels are compared in the kernel sensitivity analysis. To allow a fair comparison, the Epanechnikov perturbations were standardized to unit variance.

B.3 Reproducibility Settings

The following computational settings, $M = 500$, $B = 2000$, and $\alpha = 0.05$, were used in the simulation study. The sample sizes confirmed in the main text were $n = 15$ and $n = 30$, with $n = 15$ used for the main small-sample comparisons and $n = 30$ reported in the supplementary comparative results. The data-generating scenarios included a normal benchmark, a heavy-tailed t_3 distribution, a centered lognormal distribution, a symmetric bimodal mixture, and a symmetric trimodal mixture.

Appendix C: Supplementary Comparative Simulation Results

This appendix reports the expanded comparative simulation results for $n = 30$. These results complement the main small-sample comparison reported in Table 1 for $n = 15$. The same four resampling procedures were compared: the standard empirical bootstrap, the globally smoothed bootstrap, the Bayesian bootstrap, and the proposed AGB method. The same performance measures were used: empirical coverage probability (CP), Monte Carlo standard error of coverage, average confidence interval length, bias, RMSE of the bootstrap-based point estimate, and RMSE of the estimated standard error.

Table S1. Expanded comparative simulation results for $(n = 30,\ M = 500,\ B = 2000)$.

Scenario	Method	CP	MCSE(CP)	Avg. Length	Bias	RMSE	RMSE(SE)
Normal	Standard Bootstrap	0.9200	0.0121	0.6942	-0.0003	0.1884	0.0238
Normal	Global Smoothed Bootstrap	0.9460	0.0101	0.7565	-0.0004	0.1884	0.0282
Normal	Bayesian Bootstrap	0.9100	0.0128	0.6857	0.0000	0.1885	0.0246
Normal	AGB	0.9520	0.0096	0.7660	-0.0002	0.1882	0.0297
Heavy-tailed (t_3)	Standard Bootstrap	0.9280	0.0116	1.1156	-0.0156	0.3018	0.0915
Heavy-tailed (t_3)	Global Smoothed Bootstrap	0.9420	0.0105	1.1760	-0.0158	0.3022	0.0864
Heavy-tailed (t_3)	Bayesian Bootstrap	0.9140	0.0125	1.1079	-0.0161	0.3018	0.0915
Heavy-tailed (t_3)	AGB	0.9360	0.0109	1.1949	-0.0158	0.3017	0.0863
Skewed lognormal	Standard Bootstrap	0.8620	0.0154	1.3723	0.0187	0.4357	0.2115
Skewed lognormal	Global Smoothed Bootstrap	0.8780	0.0146	1.4250	0.0178	0.4355	0.2078
Skewed lognormal	Bayesian Bootstrap	0.8540	0.0158	1.3660	0.0185	0.4358	0.2111
Skewed lognormal	AGB	0.8840	0.0143	1.4412	0.0178	0.4356	0.2072
Bimodal mixture	Standard Bootstrap	0.9440	0.0103	1.5582	0.0373	0.4121	0.0357
Bimodal mixture	Global Smoothed Bootstrap	0.9660	0.0081	1.7177	0.0367	0.4116	0.0493
Bimodal mixture	Bayesian Bootstrap	0.9460	0.0101	1.5308	0.0375	0.4122	0.0370
Bimodal mixture	AGB	0.9700	0.0076	1.7272	0.0372	0.4119	0.0506
Trimodal mixture	Standard Bootstrap	0.9320	0.0113	1.7786	-0.0083	0.4544	0.0404
Trimodal mixture	Global Smoothed Bootstrap	0.9580	0.0090	1.9603	-0.0086	0.4550	0.0572
Trimodal mixture	Bayesian Bootstrap	0.9280	0.0116	1.7486	-0.0085	0.4555	0.0417
Trimodal mixture	AGB	0.9580	0.0090	1.9667	-0.0094	0.4540	0.0579

Table S1 shows that the qualitative patterns observed at $n = 15$ generally persist at $n = 30$, although the differences among methods tend to become smaller as the sample size increases. In regular settings, all methods show broadly comparable RMSE values, while smoothing-based methods tend to produce slightly wider intervals. In the heavy-tailed and multimodal scenarios, AGB generally maintains coverage close to the nominal level, with interval lengths comparable to the globally smoothed bootstrap. In the skewed lognormal setting, under-coverage

remains present across all percentile-based methods, confirming that strong skewness is not fully corrected by adaptive smoothing alone.

Appendix D: Full Sensitivity Analysis Results

The full sensitivity analyses carried out to assess the robustness of AGB to tuning and implementation choices are reported in this appendix. Three components were put under scrutiny: the pilot bandwidth multiplier, the truncation rule applied to the local bandwidth factor λ, and the kernel used in the generative perturbation step.

D.1 Pilot Bandwidth Sensitivity

To assess the sensitivity of AGB to the amount of experimental signal smoothing used in the adaptive density-estimation step, the rule-of-thumb bandwidth was changed by multiplying the pilot bandwidth by factors of 0.5, 1, 2 and 5.

Table S2. Full pilot-bandwidth sensitivity results for AGB at $(n = 15)$ and $(n = 30)$.

Scenario	n	Pilot multiplier	CP	MCSE(CP)	Avg. Length	MCSE(Length)	Bias	RMSE	RMSE(SE)
Normal	15	0.5000	0.9120	0.0127	1.0009	0.0088	0.0150	0.2631	0.0502
Normal	15	1.0000	0.9420	0.0105	1.0787	0.0094	-0.0193	0.2535	0.0569
Normal	15	2.0000	0.9820	0.0059	1.3751	0.0123	0.0052	0.2496	0.1157
Normal	15	5.0000	1.0000	0.0000	2.5889	0.0271	-0.0057	0.2494	0.4303
Normal	30	0.5000	0.9420	0.0105	0.7170	0.0043	-0.0095	0.1857	0.0242
Normal	30	1.0000	0.9620	0.0086	0.7764	0.0046	-0.0033	0.1826	0.0304
Normal	30	2.0000	0.9800	0.0063	0.9558	0.0056	0.0009	0.1892	0.0690
Normal	30	5.0000	1.0000	0.0000	1.7632	0.0132	0.0078	0.1755	0.2774
Heavy-tailed (t_3)	15	0.5000	0.8960	0.0137	1.5156	0.0299	-0.0402	0.4535	0.1925
Heavy-tailed (t_3)	15	1.0000	0.9540	0.0094	1.6591	0.0293	0.0170	0.4491	0.1796
Heavy-tailed (t_3)	15	2.0000	0.9700	0.0076	2.0087	0.0339	-0.0227	0.4490	0.2106
Heavy-tailed (t_3)	15	5.0000	1.0000	0.0000	3.4899	0.0474	0.0207	0.4487	0.5202
Heavy-tailed (t_3)	30	0.5000	0.9460	0.0101	1.1437	0.0208	-0.0337	0.2988	0.1287
Heavy-tailed (t_3)	30	1.0000	0.9560	0.0092	1.2070	0.0193	0.0089	0.3117	0.1187
Heavy-tailed (t_3)	30	2.0000	0.9760	0.0068	1.4022	0.0176	-0.0050	0.3050	0.1106
Heavy-tailed (t_3)	30	5.0000	1.0000	0.0000	2.3809	0.0241	0.0157	0.2983	0.3210
Skewed lognormal	15	0.5000	0.8460	0.0161	1.7164	0.0440	-0.0112	0.5328	0.2957
Skewed lognormal	15	1.0000	0.8300	0.0168	1.7518	0.0434	-0.0489	0.5323	0.2825
Skewed lognormal	15	2.0000	0.9220	0.0120	2.1931	0.0467	0.0200	0.5247	0.2797
Skewed lognormal	15	5.0000	0.9640	0.0083	3.5154	0.0680	-0.0350	0.5244	0.5244
Skewed lognormal	30	0.5000	0.9060	0.0131	1.3603	0.0303	0.0096	0.3941	0.1899
Skewed lognormal	30	1.0000	0.8840	0.0143	1.3739	0.0284	-0.0059	0.4103	0.1763
Skewed lognormal	30	2.0000	0.9300	0.0114	1.6044	0.0304	0.0013	0.3906	0.1840
Skewed lognormal	30	5.0000	0.9860	0.0053	2.5068	0.0376	-0.0012	0.4039	0.3291
Bimodal mixture	15	0.5000	0.9340	0.0111	2.2333	0.0119	-0.0247	0.5713	0.0686
Bimodal mixture	15	1.0000	0.9320	0.0113	2.4330	0.0139	-0.0341	0.5982	0.0919
Bimodal mixture	15	2.0000	0.9820	0.0059	3.2048	0.0187	-0.0205	0.5922	0.2652
Bimodal mixture	15	5.0000	1.0000	0.0000	6.3940	0.0370	-0.0185	0.5358	1.0756
Bimodal mixture	30	0.5000	0.9440	0.0103	1.6027	0.0059	0.0075	0.3973	0.0334
Bimodal mixture	30	1.0000	0.9620	0.0086	1.7263	0.0065	0.0089	0.4074	0.0500
Bimodal mixture	30	2.0000	0.9820	0.0059	2.1578	0.0080	0.0088	0.4268	0.1499
Bimodal mixture	30	5.0000	1.0000	0.0000	4.0268	0.0162	-0.0088	0.4209	0.6284
Trimodal mixture	15	0.5000	0.9220	0.0120	2.5533	0.0152	-0.0455	0.6524	0.0869
Trimodal mixture	15	1.0000	0.9480	0.0099	2.7838	0.0165	-0.0274	0.6715	0.1099
Trimodal mixture	15	2.0000	0.9800	0.0063	3.6392	0.0211	0.0542	0.6533	0.2978

Continued on next page

Table S2 continued from previous page

Scenario	n	Pilot multiplier	CP	MCSE(CP)	Avg. Length	MCSE(Length)	Bias	RMSE	RMSE(SE)
Trimodal mixture	15	5.0000	1.0000	0.0000	7.1377	0.0452	-0.0164	0.6527	1.1942
Trimodal mixture	30	0.5000	0.9460	0.0101	1.8262	0.0060	-0.0014	0.4614	0.0388
Trimodal mixture	30	1.0000	0.9690	0.0079	1.9695	0.0071	-0.0201	0.4432	0.0562
Trimodal mixture	30	2.0000	0.9840	0.0056	2.4389	0.0094	0.0089	0.4694	0.1670
Trimodal mixture	30	5.0000	0.9980	0.0020	4.5088	0.0193	-0.0014	0.4557	0.6957

Table S2 shows a clear and interpretable coverage–length trade-off. Smaller pilot bandwidths generally produce shorter intervals but may lead to under-coverage, especially in heavy-tailed and skewed settings. Larger pilot bandwidths increase empirical coverage by widening the resulting intervals. The multiplier-5 configuration often produces nearly conservative coverage, but at the cost of substantially inflated interval lengths. The default rule-of-thumb multiplier therefore provides a practical compromise between interval calibration and interval width in most examined settings

D.2 Lambda Clipping Sensitivity

The default AGB implementation truncates the local bandwidth factor λ to $[0.2, 3.0]$ as a finite sample stabilization device. To assess whether the results depend on this choice, four truncation rules were compared: no clipping, fixed clipping at $[0.1, 5.0]$, fixed clipping at $[0.2, 3.0]$, and data dependent quantile clipping based on the empirical 5th and 95th percentiles of λ.

Table S3. Full λ-clipping sensitivity results for AGB.

Scenario	n	Clipping rule	CP	MCSE (CP)	Avg. Length	MCSE (Length)	Bias	RMSE	RMSE (SE)	Mean (λ)	SD (λ)
Normal	15	No clipping	0.9500	0.0097	1.0803	0.0095	-0.0006	0.2507	0.0567	1.0279	0.2619
Normal	15	Fixed [0.1, 5.0]	0.9580	0.0090	1.0902	0.0092	0.0019	0.2487	0.0563	1.0278	0.2623
Normal	15	Fixed [0.2, 3.0]	0.9500	0.0097	1.0879	0.0093	0.0046	0.2621	0.0567	1.0274	0.2599
Normal	15	$[Q_{0.05}, Q_{0.95}]$	0.9560	0.0092	1.0837	0.0094	-0.0005	0.2548	0.0571	1.0251	0.2520
Normal	30	No clipping	0.9660	0.0081	0.7719	0.0046	-0.0187	0.1790	0.0297	1.0371	0.3224
Normal	30	Fixed [0.1, 5.0]	0.9740	0.0071	0.7762	0.0046	0.0116	0.1719	0.0304	1.0366	0.3210
Normal	30	Fixed [0.2, 3.0]	0.9560	0.0092	0.7714	0.0044	0.0072	0.1888	0.0289	1.0370	0.3228
Normal	30	$[Q_{0.05}, Q_{0.95}]$	0.9840	0.0056	0.7753	0.0046	-0.0162	0.1680	0.0302	1.0239	0.2732
Heavy-tailed (t_3)	15	No clipping	0.9520	0.0096	1.6615	0.0301	-0.0278	0.4311	0.1837	1.0424	0.3366
Heavy-tailed (t_3)	15	Fixed [0.1, 5.0]	0.9420	0.0105	1.6294	0.0302	0.0250	0.4017	0.1805	1.0414	0.3300
Heavy-tailed (t_3)	15	Fixed [0.2, 3.0]	0.9520	0.0096	1.5882	0.0282	-0.0017	0.4188	0.1744	1.0409	0.3282
Heavy-tailed (t_3)	15	$[Q_{0.05}, Q_{0.95}]$	0.9500	0.0097	1.6404	0.0264	0.0012	0.4299	0.1570	1.0405	0.3317
Heavy-tailed (t_3)	30	No clipping	0.9600	0.0088	1.2058	0.0186	-0.0040	0.3043	0.1121	1.0630	0.4414
Heavy-tailed (t_3)	30	Fixed [0.1, 5.0]	0.9560	0.0092	1.2117	0.0216	-0.0011	0.2955	0.1292	1.0664	0.4553
Heavy-tailed (t_3)	30	Fixed [0.2, 3.0]	0.9460	0.0101	1.1877	0.0185	-0.0057	0.3346	0.1124	1.0638	0.4412
Heavy-tailed (t_3)	30	$[Q_{0.05}, Q_{0.95}]$	0.9520	0.0096	1.2050	0.0167	0.0108	0.3094	0.0989	1.0523	0.4008
Skewed lognormal	15	No clipping	0.8480	0.0161	1.7873	0.0409	-0.0097	0.5094	0.2671	1.0525	0.3876
Skewed lognormal	15	Fixed [0.1, 5.0]	0.8420	0.0163	1.9266	0.0540	0.0379	0.6247	0.3326	1.0527	0.3891
Skewed lognormal	15	Fixed [0.2, 3.0]	0.8520	0.0159	1.8948	0.0626	0.0345	0.6411	0.4104	1.0511	0.3830
Skewed lognormal	15	$[Q_{0.05}, Q_{0.95}]$	0.8620	0.0154	1.8556	0.0483	0.0305	0.5731	0.3053	1.0498	0.3775
Skewed lognormal	30	No clipping	0.9020	0.0133	1.4182	0.0299	0.0201	0.3992	0.1817	1.0834	0.5275
Skewed lognormal	30	Fixed [0.1, 5.0]	0.8980	0.0135	1.4077	0.0340	0.0108	0.4135	0.2192	1.0826	0.5228
Skewed lognormal	30	Fixed [0.2, 3.0]	0.9040	0.0132	1.3810	0.0305	0.0140	0.3934	0.1968	1.0823	0.5203
Skewed lognormal	30	$[Q_{0.05}, Q_{0.95}]$	0.9140	0.0125	1.4262	0.0291	0.0060	0.3840	0.1793	1.0770	0.5039
Bimodal mixture	15	No clipping	0.9540	0.0094	2.4689	0.0144	0.0137	0.5590	0.0988	1.0120	0.1551
Bimodal mixture	15	Fixed [0.1, 5.0]	0.9500	0.0097	2.4709	0.0137	0.0212	0.5763	0.0959	1.0120	0.1569

Continued on next page

Table S3 continued from previous page

Scenario	n	Clipping rule	CP	MCSE (CP)	Avg. Length	MCSE (Length)	Bias	RMSE	RMSE (SE)	Mean (λ)	SD (λ)
Bimodal mixture	15	Fixed [0.2, 3.0]	0.9520	0.0096	2.4804	0.0141	0.0235	0.5828	0.0985	1.0127	0.1618
Bimodal mixture	15	$[Q_{0.05}, Q_{0.95}]$	0.9560	0.0092	2.4334	0.0143	-0.0024	0.5702	0.0940	1.0112	0.1559
Bimodal mixture	30	No clipping	0.9540	0.0094	1.7257	0.0063	0.0038	0.4126	0.0492	1.0130	0.1728
Bimodal mixture	30	Fixed [0.1, 5.0]	0.9600	0.0088	1.7295	0.0064	-0.0142	0.4154	0.0497	1.0124	0.1682
Bimodal mixture	30	Fixed [0.2, 3.0]	0.9740	0.0071	1.7354	0.0067	0.0098	0.3950	0.0519	1.0125	0.1703
Bimodal mixture	30	$[Q_{0.05}, Q_{0.95}]$	0.9340	0.0111	1.7152	0.0067	-0.0227	0.4241	0.0486	1.0050	0.1394
Trimodal mixture	15	No clipping	0.9740	0.0071	2.8199	0.0153	-0.0582	0.6350	0.1080	1.0107	0.1465
Trimodal mixture	15	Fixed [0.1, 5.0]	0.9440	0.0103	2.8105	0.0167	-0.0258	0.6776	0.1135	1.0115	0.1504
Trimodal mixture	15	Fixed [0.2, 3.0]	0.9440	0.0103	2.8032	0.0168	-0.0479	0.6874	0.1119	1.0113	0.1519
Trimodal mixture	15	$[Q_{0.05}, Q_{0.95}]$	0.9640	0.0083	2.8030	0.0161	0.0241	0.6600	0.1099	1.0107	0.1496
Trimodal mixture	30	No clipping	0.9680	0.0079	1.9696	0.0073	0.0265	0.4426	0.0571	1.0078	0.1250
Trimodal mixture	30	Fixed [0.1, 5.0]	0.9640	0.0083	1.9615	0.0069	0.0068	0.4499	0.0540	1.0083	0.1294
Trimodal mixture	30	Fixed [0.2, 3.0]	0.9520	0.0096	1.9634	0.0074	0.0140	0.4976	0.0557	1.0089	0.1336
Trimodal mixture	30	$[Q_{0.05}, Q_{0.95}]$	0.9680	0.0079	1.9662	0.0074	0.0205	0.4613	0.0565	1.0044	0.1176

Broad stability in the empirical behavior of AGB is observed across all examined clipping rules in the normal, heavy-tailed, and multimodal scenarios, as shown in Table S3. The default rule $[0.2, 3.0]$ is found to maintain stable coverage and interval length while keeping extreme local bandwidths in check-all without any material change to the main conclusions. Comparable results are also produced by the quantile-based rule, which offers a data-dependent alternative. In the highly skewed lognormal setting, under-coverage is exhibited by all clipping rules, which points to a limitation tied to percentile interval behavior under strong skewness rather than to the choice of clipping rule.

D.3 Kernel Sensitivity

Gaussian perturbations are used by default in the generative step of AGB. Given that the theoretical formulation is not restricted to a single kernel family, a kernel sensitivity analysis was carried out in which Gaussian and Epanechnikov perturbations were directly compared. To place both on equal footing, the Epanechnikov perturbations were standardized to unit variance.

Table S4. Kernel sensitivity results for Gaussian and Epanechnikov kernels.

Scenario	n	Kernel	CP	MCSE(CP)	Avg. Length	MCSE(Length)	Bias	RMSE	RMSE(SE)
Normal	15	Gaussian	0.9600	0.0088	1.0659	0.0089	-0.0035	0.2510	0.0527
Normal	15	Epanechnikov	0.9320	0.0113	1.0627	0.0094	-0.0096	0.2696	0.0555
Normal	30	Gaussian	0.9620	0.0086	0.7686	0.0047	0.0150	0.1778	0.0299
Normal	30	Epanechnikov	0.9620	0.0086	0.7675	0.0044	-0.0042	0.1877	0.0284
Heavy-tailed (t_3)	15	Gaussian	0.9400	0.0106	1.6532	0.0308	-0.0007	0.4451	0.1872
Heavy-tailed (t_3)	15	Epanechnikov	0.9440	0.0103	1.6973	0.0405	0.0153	0.5007	0.2503
Heavy-tailed (t_3)	30	Gaussian	0.9440	0.0103	1.2182	0.0161	0.0074	0.2982	0.0935
Heavy-tailed (t_3)	30	Epanechnikov	0.9520	0.0096	1.1965	0.0168	0.0171	0.3133	0.0992
Skewed lognormal	15	Gaussian	0.8760	0.0147	1.9382	0.1044	0.0399	0.8736	0.7251
Skewed lognormal	15	Epanechnikov	0.8920	0.0139	1.8606	0.0465	0.0220	0.5414	0.2948
Skewed lognormal	30	Gaussian	0.8820	0.0144	1.3498	0.0228	-0.0181	0.3491	0.1423
Skewed lognormal	30	Epanechnikov	0.8940	0.0138	1.4352	0.0296	0.0271	0.4294	0.1787
Bimodal mixture	15	Gaussian	0.9440	0.0103	2.4614	0.0136	-0.0254	0.5910	0.0947
Bimodal mixture	15	Epanechnikov	0.9540	0.0094	2.4476	0.0138	0.0476	0.5899	0.0938
Bimodal mixture	30	Gaussian	0.9520	0.0096	1.7235	0.0066	-0.0257	0.4189	0.0493

Continued on next page

Table S4 continued from previous page

Scenario	n	Kernel	CP	MCSE(CP)	Avg. Length	MCSE(Length)	Bias	RMSE	RMSE(SE)
Bimodal mixture	30	Epanechnikov	0.9680	0.0079	1.7303	0.0065	-0.0352	0.4135	0.0504
Trimodal mixture	15	Gaussian	0.9580	0.0090	2.7931	0.0155	-0.0188	0.6670	0.1061
Trimodal mixture	15	Epanechnikov	0.9460	0.0101	2.8056	0.0163	-0.0216	0.6465	0.1106
Trimodal mixture	30	Gaussian	0.9520	0.0096	1.9638	0.0075	-0.0498	0.4696	0.0564
Trimodal mixture	30	Epanechnikov	0.9720	0.0074	1.9730	0.0075	0.0068	0.4439	0.0582

Table S4 shows that Gaussian and Epanechnikov kernels produce broadly similar coverage length behavior. In the heavy-tailed and multimodal scenarios, both kernels yield empirical coverage close to the nominal level. In the highly skewed lognormal setting, both kernels continue to under cover, reinforcing the interpretation that this limitation is associated with percentile interval construction under strong skewness rather than with the specific perturbation kernel. The results indicate that the empirical performance of AGB is not driven solely by the Gaussian kernel choice.

REFERENCES

1. Efron, B., *Bootstrap methods: Another look at the jackknife*, The Annals of Statistics, vol. 7, pp. 1–26, 1979.
2. Silverman, B. W., *Density estimation for statistics and data analysis*, Chapman and Hall, London, 1986.
3. Davison, A. C., Hinkley, D. V., *Bootstrap methods and their application*, Cambridge University Press, 1997.
4. Efron, B., Tibshirani, R., *Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy*, Statistical Science, vol. 1, pp. 54–75, 1986.
5. Efron, B., *Better bootstrap confidence intervals*, Journal of the American Statistical Association, vol. 82, pp. 171–185, 1987.
6. Hall, P., *Theoretical comparison of bootstrap confidence intervals*, The Annals of Statistics, vol. 16, pp. 927–953, 1988.
7. DiCiccio, T. J., Efron, B., *Bootstrap confidence intervals*, Statistical Science, vol. 11, pp. 189–228, 1996.
8. Hall, P., *The bootstrap and Edgeworth expansion*, Springer, New York, 1992.
9. Silverman, B. W., Young, G. A., *The bootstrap: To smooth or not to smooth?*, Biometrika, vol. 74, pp. 469–479, 1987.
10. Hall, P., DiCiccio, T. J., Romano, J. P., *On smoothing and the bootstrap*, The Annals of Statistics, vol. 17, pp. 692–704, 1989.
11. De Angelis, D., Young, G. A., *Smoothing the bootstrap*, International Statistical Review, vol. 60, pp. 45–56, 1992.
12. Wang, S., *Optimizing the smoothed bootstrap*, Annals of the Institute of Statistical Mathematics, vol. 47, pp. 65–80, 1995.
13. Sheather, S. J., Jones, M. C., *A reliable data-based bandwidth selection method for kernel density estimation*, Journal of the Royal Statistical Society: Series B, vol. 53, pp. 683–690, 1991.
14. Chen, Y.-C., *A tutorial on kernel density estimation and recent advances*, Biostatistics & Epidemiology, vol. 4, pp. 161–187, 2020.
15. Rubin, D. B., *The Bayesian bootstrap*, The Annals of Statistics, vol. 9, pp. 130–134, 1981.
16. Wu, C. J., *Jackknife, bootstrap and other resampling methods in regression analysis*, The Annals of Statistics, vol. 14, pp. 1261–1295, 1986.
17. Al-Sharadqah, A., Mojirsheibani, M., Pouliot, W., *On the performance of weighted bootstrapped kernel deconvolution density estimators*, Statistical Papers, vol. 61, pp. 1773–1798, 2020.
18. Dong, H., Otsu, T., Taylor, L., *Bandwidth selection for nonparametric regression with errors-in-variables*, Econometric Reviews, vol. 42, pp. 393–419, 2023.
19. Sui, Y., West, M., *Adaptive kernel estimation for sparse functional data*, Biometrika, vol. 110, pp. 675–690, 2023.
20. Abramson, I. S., *On bandwidth variation in kernel estimates: A square root law*, The Annals of Statistics, vol. 10, pp. 1217–1223, 1982.
21. Tang, X., Wang, Y., *TAKDE: Temporal adaptive kernel density estimator for real-time dynamic density estimation*, arXiv preprint arXiv:2203.08317, 2022.
22. Terrell, G. R., Scott, D. W., *Variable kernel density estimation*, The Annals of Statistics, vol. 20, pp. 1236–1265, 1992.
23. Jones, M. C., *Variable kernel density estimates*, Australian Journal of Statistics, vol. 32, pp. 361–371, 1990.
24. Jones, M. C., McKay, I. J., Hu, T.-C., *Variable location and scale kernel density estimation*, Annals of the Institute of Statistical Mathematics, vol. 46, pp. 521–535, 1994.
25. Wand, M. P., Jones, M. C., *Kernel smoothing*, Chapman and Hall, London, 1995.
26. Scott, D. W., *Multivariate density estimation: Theory, practice, and visualization*, John Wiley & Sons, 2015.
27. Efron, B., Tibshirani, R. J., *An introduction to the bootstrap*, Chapman and Hall, New York, 1993.
28. Hall, P., *On the bootstrap and confidence intervals*, The Annals of Statistics, vol. 14, pp. 1431–1452, 1986.
29. Taylor, C. C., *Bootstrap choice of the smoothing parameter in kernel density estimation*, Biometrika, vol. 76, pp. 705–712, 1989.
30. Bouzebda, S., Elhattab, I., Slaoui, Y., Taachouch, N., *Nonparametric recursive kernel type estimators for the moment generating function under censored data*, Statistics, Optimization and Information Computing, vol. 11, no. 2, pp. 196–215, 2023.
31. Slama, S., Slaoui, Y., Fathallah, H., *Statistical inference for multivariate conditional cumulative distribution function estimation by stochastic approximation method*, Statistics, Optimization and Information Computing, vol. 10, no. 2, pp. 789–814.