Paper Feedback Report (v1)

Prepared by ProofInk.AI (ProofInk.AI can make mistakes. Always double-check responses)

Generated: 8/2/2026, 5:28:31 PM

Paper: Information geometry for approximate Bayesian computation

Feedback Mode: Quick Feedback

Review Version: v1

Processing Time: 70.26 seconds

4/5

Overall Score

Likely outcome (approximate): Likely minor revisions to likely accept — an uncertain interpretation of the score, not a guarantee.

This feedback report is intended to help strengthen your submission before it reaches human reviewers. It highlights potential concerns and opportunities for improvement — not to discourage, but to give you the best chance of success. Not all feedback points carry equal weight. Use your expertise to prioritize.

About This Paper

This paper analyzes the Approximate Bayesian Computation (ABC) algorithm using relative entropy between the true and ABC-simulated posterior distributions. The authors derive a leading-order asymptotic expansion of the relative entropy as the threshold vector $\varepsilon \to 0$, extend this to a vector-valued (elliptical) threshold to enable direction-specific sensitivity analysis, specialize the results to the exponential family to expose explicit dependence on the sample size n, and derive analogous asymptotics for estimator bias and mean rejection rates. Two worked examples (exponential and normal-with-unknown-mean-and-variance) and an accompanying simulation study illustrate that allowing different tolerances per statistic (ellipse vs. ball acceptance regions) can reduce rejection rates for the same overall tolerance.

Highlights: The paper offers a genuinely novel and mathematically careful application of relative entropy/information theory to characterize ABC's dependence on the threshold vector and data size, going beyond the scalar-ε, Euclidean-metric treatments common in prior work (e.g., [3,6]). The derivations (Propositions 2.2, Lemmas 3.1/3.4, Corollaries 3.5/5.2/6.1-6.5) are detailed and internally consistent, and the exponential-family and normal examples plus the simulation study (Tables 1-3, Figures 1-2) provide a concrete, verifiable illustration of the theory's practical implication that asymmetric thresholds can reduce rejection rates.

Executive Summary

This is a solid theoretical contribution that opens a useful information-geometric perspective on ABC threshold/tolerance selection, with rigorous asymptotic derivations and illustrative examples. The main opportunities for improvement concern reconciling the abstract's claim of 'error bounds for finite sample sizes' with the asymptotic (ε→0) nature of the actual results, the restriction of demonstrations to tractable-likelihood exponential-family examples (as the authors themselves note), and a fuller discussion of the observed gap between the asymptotic rejection-rate formula and the simulation results in Table 3.

Tip: For deeper analysis with criterion-by-criterion feedback, gap analysis, and detailed suggestions, consider using the Deep Analysis mode.

Strengths

- Novel theoretical framework: the vectorized threshold ε∈R^q_+ and elliptical acceptance region (2.5) generalize prior scalar-ε analyses (e.g. [3,6]) and yield an explicit, data-driven (τ*-dependent) leading-order relative entropy expansion (2.6)-(2.7), enabling a genuinely new sensitivity-analysis tool for ABC tuning.
- Mathematical rigor and transparency: proofs (Proposition 2.2, Lemma A.1, Lemma 3.4) are worked through in detail with explicit intermediate steps, stated regularity conditions (Condition 2.1, 3.2, 3.3), and appendix proofs of all corollaries, aiding reproducibility of the theoretical results.
- Theory-simulation connection: Section 7 directly operationalizes the theoretical bias/relative-entropy/rejection-rate formulas via a concrete optimization problem (7.1) and compares asymptotic predictions to Monte Carlo results (Table 3, Figure 2), which is a commendable effort to validate the asymptotics empirically.

Weaknesses

- [MAJOR] The abstract states the analysis 'provides error bounds on performance for positive tolerances and finite sample sizes,' but the core results (e.g., Proposition 2.2, Lemma 3.4, eq. (3.4)) are asymptotic expansions valid 'as $|\varepsilon|\to 0$' or 'as $|n\varepsilon|\to 0$' with $O(|\varepsilon|^6)/O(|n\varepsilon|^6)$ remainder terms that are never bounded explicitly with computable constants. This is a mismatch between the claimed finite-sample/finite-tolerance guarantee and what is actually delivered (leading-order asymptotics). Fix: either soften the abstract's language to 'asymptotic expansions for small tolerance' or supply explicit non-asymptotic bounds on the remainder terms (e.g., via Taylor remainder theorems with verifiable Lipschitz/boundedness constants).
- [MAJOR] Both worked examples in Section 6 (exponential rate inference, normal mean/variance inference) are conjugate exponential-family models with fully tractable likelihoods; the authors themselves state on page 17 that 'ABC methods are not necessary for inference' in these cases. This limits the empirical demonstration of the framework's value to settings where ABC would actually be used (intractable likelihoods, non-conjugate/complex simulators). Fix: include at least one example with an intractable or only-simulatable likelihood (e.g., a simple population-genetics or agent-based model referenced in the introduction) to validate practical applicability.
- [MAJOR] Table 3 shows the asymptotic ratio formula $\tilde{U}(\tau^*)$ from Corollary 5.2/eq.(6.9) noticeably overestimating the empirical ratio $\hat{R}\_E/\hat{R}\_B$ across most (tol, n) pairs (e.g., 0.763 vs. 0.641 at (0.25,100); 0.648 vs. 0.537 at (0.25,300)), and the paper only notes qualitatively that 'the asymptotic formula $\tilde{U}$ overestimates the ratio… due to the fact that it is only an asymptotically correct formula' (page 24-25) without quantifying the error or its dependence on n and tol. This leaves the practical reliability of the rejection-rate formula for real (non-infinitesimal) ε unaddressed. Fix: report the relative/absolute discrepancy as a function of $n|\varepsilon|$ to establish a validity regime, or provide a second-order correction term.
- [MAJOR] Lemma 3.4 (and the Laplace-asymptotics argument in its proof, page 14) requires Condition 3.3 — that $g(\theta)=\Sigma\eta_i(\theta)\tau_i-A(\theta)$ has a unique interior non-degenerate maximum — which is a nontrivial restriction excluding multimodal posteriors or boundary-concentrated maxima that can arise in many exponential-family models with informative data or boundary parameter constraints. The generality of the n^4 scaling result and its exceptions (Section 6) is therefore implicitly limited to well-behaved unimodal posterior regimes, but this restriction is not discussed in terms of how it constrains the class of applicable models. Fix: explicitly discuss which common exponential-family models could violate Condition 3.3 (e.g., near-boundary parameters, weakly identified models) and how the theory would need to be adapted.
- [MAJOR] The paper's asymptotic machinery simultaneously requires $|\varepsilon|\to 0$ (for the ε-Taylor expansion in Proposition 2.2/Lemma A.1) and, for the exponential family, $n|\varepsilon|\to 0$ (Lemma 3.4, Corollary 5.2) as the regime of validity, but in practice n is fixed and can be large while ABC still

requires a non-infinitesimal ε for computational feasibility (finite acceptance rate). The text acknowledges this tension only briefly ('It is important to remark here though that in ABC the size of the data n is fixed and one can only change ε,' page 9) without giving concrete guidance (e.g., numerically calibrated ranges of $n|\varepsilon|$ where the theory is quantitatively accurate, beyond the qualitative Table 3 comparison). Fix: provide a table or plot of $n|\varepsilon|$ against approximation accuracy to give practitioners actionable guidance on when the asymptotic formulas are trustworthy.

Major Issues

- Abstract vs. results mismatch on 'error bounds for finite sample sizes': the theoretical results are asymptotic expansions (as $|\varepsilon|\to 0$) with unbounded big-O remainders, not finite-sample non-asymptotic error bounds; see Proposition 2.2 (page 7) and Lemma 3.4 (page 12). This could mislead readers about the strength of the theoretical guarantees; recommend tempering the abstract's claims or adding explicit remainder bounds.
- Empirical validation is confined to tractable-likelihood conjugate exponential-family toy models (Section 6, pages 17-23), a limitation the authors candidly note (page 17) but which nonetheless restricts the demonstrated practical relevance of the framework to genuine ABC use-cases with intractable likelihoods.
- Observed systematic overestimation of the asymptotic mean-rejection-rate ratio Ũ relative to Monte Carlo estimates (Table 3, pages 24-25) is noted but not quantitatively characterized, leaving open how much practitioners should trust the formula outside the $\varepsilon, n \to 0/\infty$ limits.
- Condition 3.3's requirement of a unique non-degenerate interior maximum (page 12) is a strong hidden restriction on the applicability of the n^4-scaling result (Lemma 3.4) and its refinements (Section 6), and its practical scope is not discussed.
- The dual asymptotic regimes ($|\varepsilon|\to 0$ fixed n, vs. $n|\varepsilon|\to 0$ as n→�∞) are used somewhat interchangeably in Sections 3-6 without clearly delineated practical validity ranges for the fixed, finite n and non-infinitesimal ε regime actually used in ABC practice.

Gaps and Inconsistencies

- The paper would benefit from an example with a genuinely intractable likelihood (as ABC is designed for) rather than only conjugate exponential-family cases, to demonstrate practical utility beyond a purely illustrative/pedagogical role (see page 17 acknowledgment).
- A quantitative characterization of the approximation error between the asymptotic rejection-rate formula (6.9) and simulation results (Table 3) is not provided, which would strengthen confidence in

the tool for real (non-infinitesimal) tolerance choices.

- Discussion of which classes of models violate Condition 3.3 (unique interior non-degenerate maximum, page 12) is not provided, leaving unclear how broadly the exponential-family n^4-scaling results (Lemma 3.4, Section 6) generalize.
- The abstract's claim that 'analysis provides error bounds on performance for positive tolerances and finite sample sizes' (page 1) is in tension with the paper's actual results, which are leading-order asymptotic expansions with O(|ε|^k) or O(|nε|^k) remainder terms lacking explicit finite-sample bounds (e.g., Proposition 2.2 page 7, Lemma 3.4 page 12); the asymptotic framing throughout Sections 2-6 is the more accurate description and should be reflected in the abstract's wording.
- The claim that 'error bounds on performance for positive tolerances and finite sample sizes' are provided (Abstract, page 1) is not directly substantiated by explicit non-asymptotic bounds in the body of the paper; supporting evidence would require either deriving explicit constants for the O(|ε|^6)/O(|nε|^6) remainder terms or rephrasing the claim to reflect the asymptotic nature of the results.
- The concluding remark that 'the gain in performance due to the choice of the distance metric seems to stabilize' as n increases (page 26, based on Figure 2) is described as a suggestive pattern from a single simulation setting; broader support (e.g., additional models or repeated simulation runs with confidence intervals) would strengthen this generalization before it is treated as a robust phenomenon warranting future theoretical work.

Executive Action Plan

1. Revise the abstract, introduction, and conclusion to accurately describe the main results as asymptotic expansions in \(|\varepsilon|\to 0\) rather than finite-sample error bounds, because the current wording overstates the theory and creates a credibility gap with the actual proofs.
2. Add a clear limitations/discussion section and, if possible, one more non-conjugate or simulation-based example to test the method beyond tractable exponential-family toy models, because the current empirical evidence is too narrow to support broad claims about practical performance.
3. Reconcile the theoretical asymptotics with the Monte Carlo results by tightening the conditions and/or adding a calibration or remainder analysis for the ratio approximation, because Table 3 shows systematic overestimation and the paper currently does not explain when the asymptotic formula is reliable.