MATHAUDIT / DETAILED REPORT

Correctness, proof completeness, and novelty audit

Non-asymptotic Estimates for Markov Transition Matrices via Spectral Gap Methods

De Huang, Xiangyuan Li

Paper	arXiv:2408.05963v4
Source version	Submitted 5 August, 2026; v1 submitted 12 August, 2024; announced August 2024.
Audit date	August 18, 2026
Prompt	Standardized prompt
Novelty cutoff	August 12, 2024
Paper license	arXiv non-exclusive distribution license
Scope	Statements, proofs, and novelty only

Important limitation. AI-generated analysis. This report is not a correctness certificate or a substitute for expert scrutiny. A Correct status may coexist with yellow typos or minor formal corrections that do not affect substantive soundness. It means that this audit found no unresolved substantive error under the stated criteria; it is not formal verification and should remain open to expert challenge and correction.

SUMMARY

01	Statements	Correct
02	Proofs	Correct
03	Novelty	No non-novelty findings

01 / STATEMENTS

Statements

Correct

The almost-sure convergence, entrywise tail bound, and dimension-free mean-square bound for the joint-count MLE matrix are correct. The reversible symmetric-counting construction and its matrix tail and mean-square bounds are also correct. The spectral statements for the two path-space chains support exactly the gap substitutions used in those estimates. The standing state-space hypothesis should exclude the trivial one-state case so that the displayed spectral-gap definitions are nonempty.

Theorem 2.1: The MLE joint-count estimates hold

Pages 4–5 and 17–21 · Theorem 2.1 and its proof · arXiv:2408.05963v4

Correct

Writing R_n as the empirical mean of edge indicators on the length-two path chain gives stationary mean $D_\mu P$. The ergodic theorem yields almost-sure convergence. The cited scalar Bernstein inequality applies with variance at most $\mu(u)p(u,v)$, and Theorem 3.3 supplies $\eta_p(P_2) \geq \eta_p(P)/(1+\eta_p(P))$. For the mean-square estimate, stationary covariances reduce to a polynomial in $P_2 - \Pi_2$; the factorization $P_2 = ST$, $P = TS$ reduces its norm to a polynomial in $P - \Pi$, and telescoping gives the claimed $O((n\eta_p(P))^{-1})$ bound. The alternative absolute-gap and reversible bounds follow from the same polynomial estimate.

Theorem 2.2: The reversible symmetric-counting estimates hold

Page 5 and pages 18 and 22–23 · Theorem 2.2 and its proof · arXiv:2408.05963v4

Correct

The orientation coin and one oracle call generate the stated Markov chain on unordered length-two paths. Detailed balance for P makes its invariant edge law reversible, and the symmetrized matrix observable has expectation $D_\mu P$. The path-chain absolute gap is $\eta(P)/2$, so the cited Markov-dependent matrix Bernstein bound gives the displayed operator-norm tail estimate after the constants are rescaled. Nonnegativity of the path-chain spectrum and the stationary covariance calculation give the dimension-free bound $\mathbb{E}\|H_n - D_\mu P\|_F^2 \leq (4-\eta(P))/(n\eta(P))$, with the stated initial-density factor in the nonstationary case.

Theorems 3.3 and 3.4: The path-space spectral identities are correct

Pages 11–16 · Theorems 3.3–3.4 · arXiv:2408.05963v4

Correct

The rectangular factorizations give $P_2 = ST$ and $P = TS$, so the nonzero eigenvalues agree. After weighting by the invariant measures, S has orthonormal columns and T has orthonormal rows; consequently P_2^k and P^{k-1} have the same nonzero singular values. The positive and negative parts of the additive symmetrization identify the reversible spectrum with those of $(P+I)/2$ and $(P-I)/2$, while the conditional-expectation argument yields the claimed IP-gap lower bound. For the unordered path chain, reversing before extension gives $\widetilde{T}\widetilde{S} = (P+I)/2$, proving both spectral-gap identities and uniqueness of the invariant law.

Standing state-space hypothesis: The spectral-gap definitions require at least two states

Pages 1 and 4 · standing hypotheses and Subsection 2.1.2 · arXiv:2408.05963v4

Minor formal correction

The paper permits an arbitrary finite state space but defines $\lambda(P)$ and $\eta_p(P)$ by a supremum or infimum over nonzero vectors in $L^0_{2,\mu}$. If $|\Omega| = 1$, that set is empty and the displayed gaps and subsequent fractions are not finite quantities under the stated definitions. Add the standing assumption $|\Omega| \geq 2$. This is a local range correction: the excluded one-state estimation problem is identically zero-error, and every proof and application in the paper otherwise remains unchanged.

02 / PROOFS

Proofs

Correct

The central proofs are correct and complete, apart from the harmless one-state scope correction. The weighted rectangular-factorization arguments, stationary covariance sums, polynomial norm bounds, and applications of the two external concentration theorems all satisfy their required hypotheses and yield the displayed constants.

Theorem 3.3: The singular-value and gap comparison proof is complete

Pages 12–16 · proof of Theorem 3.3 · arXiv:2408.05963v4

Correct and complete

The identities $S_1^\mathsf{T} S_1 = I$ and $T_1 T_1^\mathsf{T} = I$ follow directly from stationarity and the edge measure. They preserve the nonzero singular values in $S_1(P-\Pi)^{k-1}T_1$, proving the absolute- and pseudo-gap claims. The additive symmetrization is squeezed between its positive and negative Gram components; in the reversible case its square has the positive component as its unique positive square root, giving the asserted spectral decomposition. Finally, conditioning $(I-P_2)h$ on the first coordinate and applying the IP inequality for P gives the stated lower bound with no missing case.

Theorems 4.1–4.2 and the main tail bounds: The cited concentration inequalities are applied correctly

Pages 17–18 · Theorems 4.1–4.2 and proofs of Theorems 2.1(2), 2.2(2) · arXiv:2408.05963v4

Correct and complete

For the scalar edge indicator, centering, the unit bound, the variance proxy, the IP gap, and the initial-density ratio all match the hypotheses of Huang–Li's Bernstein theorem. For symmetric counting, $\widetilde{F} - D_\mu P$ is Hermitian, has norm at most 2, and has matrix second moment at most $2I$; the unordered path chain has the required positive absolute gap. Neeman–Shi–Ward's Bernstein theorem and its nonstationary corollary then give the paper's deliberately weakened constants, including the two-sided factor for the operator norm.

Verifiable source: Huang–Li, Theorem 2.4

Theorem 2.2(2): The matrix concentration input supports the displayed dependence

Page 18 · proof of Theorem 2.2(2) · arXiv:2408.05963v4

Correct and complete

The source theorem uses the contraction parameter $1-\eta_a$ through $(1+\lambda)/(1-\lambda)$ and a linear Bernstein term proportional to $(1-\lambda)^{-1}$. Substituting $\eta_a(\widetilde{P}_2) = \eta(P)/2$, the variance proxy 2, the norm bound 2, and t for the empirical average yields a bound at least as strong as the one printed in Theorem 2.2(2). The source's nonstationary extension supplies exactly the factor $\|\nu/\mu\|_\infty$.

Verifiable source: Neeman–Shi–Ward, matrix Bernstein theorem

Proofs of Theorems 2.1(3) and 2.2(3): The covariance-polynomial estimates are complete

Pages 19–23 · Subsection 4.2 · arXiv:2408.05963v4

Correct and complete

Stationarity turns each mean-square Frobenius error into a sum of lagged scalar covariances. The sum of squared centered-indicator norms is at most one, so an operator norm controls the entire expression without a dimension factor. For MLE, the path-space factorization and the invertibility of $I - P$ on the mean-zero finite-dimensional subspace reduce the polynomial to a telescoping sum bounded by $4/n$. For SCE, all nonconstant eigenvalues of $\widetilde{P}_2$ lie in $[0, 1)$, and evaluating the nonnegative covariance polynomial at the largest one gives the stated constant.

03 / NOVELTY

Novelty

No non-novelty findings

No non-novelty findings.

METHOD

Scope and decision rules

This audit addressed only the correctness of central statements, the correctness and completeness of proofs, and evidence-backed non-novelty. It did not evaluate importance, exposition, style, suitability for publication, or author reputation.

An adverse finding is included only where the report can give a direct calculation, counterexample, formal contradiction, or a verifiable source establishing the precise claim. An inability to verify is kept distinct from a demonstrated defect. Technical compression is not treated as incompleteness when the omitted reasoning is recoverable and correct.

Before any finding was recorded, the argument was closed under immediate, mechanically checkable consequences of facts already established. An observation was suppressed when the printed conclusion followed by elementary logic, direct substitution, weakening, inclusion, or verified monotonicity and no substantive proof obligation or textual correction remained.

Typographical errors and minor formal corrections remain fully visible in yellow with their locations and repairs. When such a defect is unambiguous, local, and harmless to the paper's substantive conclusions, it does not lower the overall statements or proofs status from Correct.

The novelty panel reports only central results for which a verifiable primary source predating the earliest public version that contains the result proves the same statement up to notation, or from which the statement follows directly and immediately. It contains no positive novelty certifications and no entries based merely on an unsuccessful literature search.

All mathematical expressions in this report are compiled by LaTeX. The web presentation independently validates and renders the corresponding expressions with KaTeX.

Sources checked

1. Source paper: arXiv:2408.05963v4
2. Huang–Li, Theorem 2.4
3. Neeman–Shi–Ward, matrix Bernstein theorem

arXiv acknowledgement

Thank you to arXiv for use of its open access interoperability. This service was not reviewed or approved by, nor does it necessarily express or reflect the policies or opinions of, arXiv.

APPENDIX

Standardized prompt

You are a mathematical auditor. Evaluate only the correctness of the paper's central statements, the correctness and completeness of its proofs, and evidence-backed non-novelty. Do not make editorial comments or evaluate importance, exposition, style, or suitability for publication. Keep the report neutral, precise, and professional.

PART 1 — CORRECTNESS OF THE MAIN STATEMENTS

Identify the paper's central statements and assess each one separately using exactly one of these outcomes:

- Correct: the statement has a correct, verifiable proof; or, despite a defect in the printed proof, you can give a fully verified repair or independent proof.
- Incorrect: you can give strong formal evidence that the statement as written is false, preferably a counterexample or a complete contradiction.
- Not able to verify: the supplied proof does not establish the statement, and you can neither supply a fully verified proof nor disprove it.

For every adverse or unverified outcome, restate the exact claim and hypotheses, identify its location in the reviewed version, and give the decisive reasoning. Do not infer that a theorem is false merely because its printed proof is defective.

Assign the overall statements status by the following priority:

1. Contains wrong statements, if at least one central statement is Incorrect.
2. Contains unsupported statements, if no central statement is Incorrect but at least one is Not able to verify.
3. Correct, if every central statement is Correct, apart from any typographical or minor formal findings allowed by Part 4.

PART 2 — CORRECTNESS AND COMPLETENESS OF THE PROOFS

Audit the proofs of the central statements and every material internal or cited input on which they depend. For each proof or proof segment, use the narrowest supported outcome:

- Correct and complete: every substantive step is verifiable and no case or crucial nontrivial step is missing.
- Incorrect as written: a displayed calculation, inference, lemma, or cited application is formally false or inapplicable.
- Incomplete as written: a necessary substantive step or case is missing, even though no explicit contradiction has been proved.
- Not able to verify: the available paper and sources are insufficient to decide whether the step is valid.

Do not flag omitted routine details, compressed arguments, or "analogously" steps when the omitted reasoning is recoverable and correct. For every defect, identify its smallest exact location, explain its downstream dependency, and classify any proposed repair as one of:

- Verified repair: supply enough reasoning to verify that the repair works throughout every affected use.
- Plausible repair only: describe the possible repair but do not treat it as established.
- No repair supplied.

Assign the overall proofs status by the following priority:

1. Contains incorrect or incomplete proofs, if at least one proof is Incorrect as written or Incomplete as written.
2. Contains unverified proofs, if there is no proved incorrectness or incompleteness but at least one material proof is Not able to verify.
3. Correct, if all audited proofs are Correct and complete, apart from any typographical or minor formal findings allowed by Part 4.

PART 3 — NOVELTY OF THE MAIN STATEMENTS

This is a one-sided, evidence-only check. Evaluate only central results. Report a result only when a verifiable primary source, publicly available before the earliest public version of the paper that contains the result, proves the same statement up to a change of notation, or when the result follows directly and immediately from such an earlier theorem.

For each non-novelty finding, provide the earlier source, its public date, the exact theorem or proposition, and an explicit hypothesis-and-conclusion mapping. A later source cannot establish non-novelty. Do not flag a result merely because an easier proof exists or because it follows through a nontrivial argument. Do not flag a result that the paper itself explicitly presents as known, a survey item, or an alternative proof.

Mention only statements that meet this non-novelty threshold. Say nothing about statements for which no such evidence was found: do not certify novelty, list unsuccessful searches, or add statement-level "no prior source found" entries. If the list is empty, the section may say only "No non-novelty findings."

PART 4 — TECHNICAL AND EVIDENTIARY RULES

1. Criticism must never be overstated. Call a statement incorrect, a proof incorrect or incomplete, or a result non-new only when strong formal evidence supports that exact judgment. Otherwise use Not able to verify where that outcome exists, or remain silent in the novelty section.

2. Review a specified paper version. Give exact page, theorem, equation, or proof locations; restate relevant quantifiers and hypotheses; and check domains, endpoints, parameter ranges, constants, and dependencies explicitly. Apply a defect only to the regimes and downstream results that actually depend on it.

3. Immediate-consequence closure: when deciding whether a conclusion is supported, treat the argument as including every immediate formal consequence of facts established earlier in the paper and available at that point. An immediate formal consequence is a local, mechanically checkable inference using only definitions already in force, elementary logic, direct substitution or renaming, weakening of an inequality or set inclusion, or a monotonicity or symmetry property explicit in a definition or already proved.

4. Such an inference must introduce no new lemma, estimate, construction, case analysis, limiting or compactness argument, external result, or new issue of uniformity or parameter dependence. Check the direction of monotonicity and the dependence of every threshold or constant explicitly. Do not exchange pointwise and uniform conclusions, change exceptional sets, interchange quantifiers or limits, or extend an endpoint under this rule unless the required fact was already established.

5. If the claimed conclusion follows after immediate-consequence closure, it is fully supported. Do not report a gap, unsupported statement, minor formal correction, or recommendation to add an explanatory sentence, and do not mention the suppressed item in the report. This suppression does not apply to a genuine typographical error as defined below, because a typo is a literal notation defect that still requires a textual correction. For example, if $P(delta)$ is proved for every sufficiently small $delta > 0$ and P is immediately monotone in the required direction, then the claim that $P(delta)$ holds for every $delta > 0$ requires no correction; the omitted monotonicity step is exposition rather than an audit finding.

6. Immediate-consequence closure permits reconstruction of an omitted inference; it does not permit silent alteration of the printed domain, hypotheses, or quantifiers of an intermediate statement later invoked outside its stated scope. If formal validity requires changing the text of that statement or its stated range, report the appropriate finding even when the repair is easy. Thus a result stated only for $t > 0$ and subsequently applied at $t = 0$ remains a formal scope mismatch unless the endpoint case was separately established before the application.

7. Apply this mandatory gate before recording any adverse, unverified, typographical, or minor formal finding: (a) identify the exact assertion or inference alleged to be defective; (b) reconstruct the strongest conclusion following from the preceding argument after all admissible immediate formal consequences are added; (c) identify the specific defect that remains — a formally demonstrated falsehood, a concrete nontrivial unresolved proof obligation, or a genuine domain, hypothesis, quantifier, or notation mismatch requiring a textual correction; and (d) state the evidence and minimal repair. If no specific defect remains, or if the only possible addition would be explanatory, suppress the finding entirely. Use Not able to verify only when a precise, nontrivial unresolved obligation can be identified.

8. A typographical error is an incorrect symbol, index, exponent, sign, variable name, or notation mismatch whose unique intended correction is mechanically determined from the surrounding text and does not alter the mathematical argument. Report it in yellow as Typo, with the printed text, correction, exact location, and reason it is harmless. It does not lower an overall status.

9. A minor formal correction is a boundary, endpoint, or displayed parameter-range defect that can be repaired consistently in every affected statement by a local change, without a new argument, without invalidating any related result or application, and without creating another substantive defect. Report it in yellow as Minor formal correction, with the same details. It does not lower an overall status.

10. Do not use either minor category if more than one materially different repair is possible; if the repair changes the substance or scope of a central result; if an excluded case is used later; if the correction fails to propagate consistently; or if a new mathematical idea is required. In such cases use the strongest substantive classification actually supported by evidence.

11. Write every mathematical expression as valid, compilable LaTeX. Use ... for inline mathematics and

...

for displayed mathematics, with balanced delimiters and braces. Do not leave raw LaTeX commands outside math delimiters.

12. Separate theorem truth from the correctness of the printed proof. A theorem can be Correct in Part 1 because of a fully verified independent proof while its printed proof is adverse in Part 2. Conversely, a defective proof alone never licenses an Incorrect theorem classification.

Organize the report under Parts 1–3 and apply Part 4 throughout. Do not add editorial commentary, scores, rankings, recommendations, or a list of suppressed observations.

END OF REPORT

Each claim in this prototype report remains open to independent challenge and revision.