MATHAUDIT / DETAILED REPORT

Correctness, proof completeness, and novelty audit

Verifiable Regularity Criterion for Conditional Expectation Operators and Conditional Mean Embeddings with Applications to Nonparametric Regression, Bayesian Inverse Problems, and Koopman Operators

Maximiliano Hertel, Ilja Klebanov, Manuel Schaller, Karl Worthmann

Paper	arXiv:2608.06155v1
Source version	Submitted 6 August, 2026; announced August 2026.
Audit date	August 18, 2026
Prompt	Standardized prompt
Novelty cutoff	August 6, 2026
Paper license	CC BY-NC-ND 4.0
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	Contains unsupported statements
02	Proofs	Contains incorrect or incomplete proofs
03	Novelty	No non-novelty findings

01 / STATEMENTS

Statements

Contains unsupported statements

The Hilbert–Schmidt regularity criterion, the abstract projection-error bounds, and the three application theorems are verified. The general learning-rate theorem is also verified under its explicit covariance-eigenvalue hypothesis. Its final claim that the Sobolev setting automatically gives the exponent $p = d/(2\ell)$ is not established for an arbitrary input law P_X: the proof substitutes approximation numbers for the Lebesgue embedding into an operator whose codomain is $L^2(P_X)$.

Theorem 3.3 and Corollary 3.8: Regular conditional densities yield Hilbert–Schmidt CEOs and CMEs

Pages 7–11 · Proposition 3.1, Theorem 3.3, and Corollary 3.8 · arXiv:2608.06155v1

Correct

If $p_{X|Y}(\cdot \mid y)$ belongs to $\mathcal{H}_{k_X}$ for almost every y and its squared RKHS norm is integrable, the Bochner kernel $y \mapsto p_{X|Y}(\cdot \mid y)$ lies in $L^2(P_Y; \mathcal{H}_{k_X})$. Proposition 3.1 therefore gives a Hilbert–Schmidt integral operator. The conditional-expectation identity verifies that this operator is the CEO, and applying it to $k_Y(\cdot, y)$ gives the stated CME representation. The boundedness, compactness, and trace-class consequences follow from standard Hilbert–Schmidt operator identities.

Theorem 3.12 and the general part of Theorem 3.16: Projection and learning rates under an assumed eigenvalue bound

Pages 12–17 · Theorems 3.12 and 3.16 · arXiv:2608.06155v1

Correct

The range condition $\operatorname{ran}(U) \subseteq \operatorname{ran}(C_{XX}^{\alpha/2})$ gives the required factorization through the fractional covariance power. The projection residual is then controlled by the corresponding power-function or interpolation estimate. Under the separately stated decay $\lambda_j(C_{XX}) \leq c j^{-1/p}$, the effective-dimension calculation and the sampling estimate yield the displayed rate after the chosen regularization balance. These conclusions do not depend on the automatic Sobolev-exponent clause discussed separately below.

Theorem 3.16, final Sobolev clause: The exponent $p = d/(2\ell)$ is not automatic for an arbitrary input distribution

Pages 16–17 · final paragraph of Theorem 3.16 and its proof · arXiv:2608.06155v1

Not able to verify

The theorem assumes only that $\mathcal{H}_{k_X} \simeq H^\ell(\mathcal{X})$ on a bounded Lipschitz domain and then states that the covariance eigenvalues satisfy $\lambda_j \asymp j^{-2\ell/d}$. But $C_{XX} = E_{P_X}^* E_{P_X}$, where $E_{P_X} : H^\ell(\mathcal{X}) \to L^2(P_X)$ is the embedding determined by the arbitrary probability law P_X. The proof cites the classical approximation-number rate for $H^\ell(\mathcal{X}) \to L^2(\mathcal{X})$ with Lebesgue measure and does not supply a comparison between these two codomains. Consequently the claimed automatic exponent, and hence the unconditional specialization of the learning rate, is not verified under the printed assumptions. A verified sufficient repair for the upper eigenvalue bound actually used in the rate is to assume P_X has an essentially bounded density with respect to Lebesgue measure; a two-sided $\asymp$ claim additionally needs an appropriate lower comparison. No counterexample to the learning-rate conclusion itself is asserted.

Theorems 4.2, 4.5, and 4.11: Regression, Bayesian inverse, and elliptic-diffusion verification results

Pages 18–32 · Section 4 · arXiv:2608.06155v1

Correct

In the regression and Bayesian settings, the stated differentiability and integrability hypotheses permit the required Sobolev estimates for the conditional density. For the uniformly elliptic diffusion, the cited transition-density regularity and the compact-state-space bounds give the required square-integrable Sobolev norm. Substitution into Theorem 3.3 then yields the advertised Hilbert–Schmidt CEO and CME conclusions in each application.

02 / PROOFS

Proofs

Contains incorrect or incomplete proofs

The core conditional-density, projection, and application proofs are correct and complete. The final Sobolev-rate specialization contains an unresolved change from $L^2(P_X)$ to Lebesgue L^2, so that portion of Theorem 3.16 is incomplete. The finite-sample projection formulas also print ordinary inverses where the already-defined Moore–Penrose pseudoinverse is required; this is a harmless notation typo.

Proposition 3.1, Theorem 3.3, and Corollary 3.8: Hilbert–Schmidt kernel and conditional-expectation argument

Pages 7–11 · Section 3.1 · arXiv:2608.06155v1

Correct and complete

The paper verifies strong measurability and square integrability of the Hilbert-space-valued density kernel, identifies its integral operator with conditional expectation by testing against functions of Y, and then evaluates the operator on kernel sections. The norm and trace identities used downstream are valid for a Hilbert–Schmidt operator.

Theorem 3.12: Fractional-range factorization and Galerkin error estimate

Pages 12–15 · Theorem 3.12 and proof · arXiv:2608.06155v1

Correct and complete

The Douglas-type range factorization, spectral calculus for C_{XX}, and projection-residual estimate are applied with compatible domains. The separate regularity regimes in the theorem match the exponents used in the interpolation bounds, and the finite-rank estimator inherits the displayed operator-norm control.

Theorem 3.16, Sobolev specialization: Lebesgue approximation numbers are used for an $L^2(P_X)$ embedding

Pages 16–17 · last step of the proof of Theorem 3.16 · arXiv:2608.06155v1

Incomplete as written

For the covariance operator, the relevant singular values are those of $E_{P_X} : H^\ell(\mathcal{X}) \to L^2(P_X)$. The proof instead invokes $a_j(H^\ell(\mathcal{X}) \hookrightarrow L^2(\mathcal{X})) \asymp j^{-\ell/d}$ and immediately concludes $\lambda_j(C_{XX}) \asymp j^{-2\ell/d}$. Without a measure-comparison hypothesis this does not follow. Downstream dependency: only the final automatic choice $p = d/(2\ell)$ and its specialized rate; the theorem under an assumed covariance-eigenvalue bound remains verified. Repair classification: verified sufficient repair for the needed upper bound. Add $dP_X/dx \in L^\infty(\mathcal{X})$, so $\|f\|_{L^2(P_X)} \le \|dP_X/dx\|_\infty^{1/2} \|f\|_{L^2(dx)}$ and the approximation-number upper rate transfers. Require a corresponding lower density bound if retaining the printed two-sided equivalence.

Notation 3.10 and finite-sample formulas: Gram-matrix inverses should be pseudoinverses

Pages 12–14 · Notation 3.10, Table 1, and the displayed empirical projection formulas · arXiv:2608.06155v1

Typo

The points are not assumed distinct and the kernels are not assumed strictly positive definite, so G_{XX} and G_{YY} need not be invertible. The paper has already defined the Moore–Penrose pseudoinverse. Replace G_{XX}^{-1} and G_{YY}^{-1} in these projection formulas by $G_{XX}^{\dagger}$ and $G_{YY}^{\dagger}$. This is the standard orthogonal-projection formula and leaves every abstract result unchanged.

Section 4: Verification of the regularity hypotheses in the three applications

Pages 18–32 · proofs of Theorems 4.2, 4.5, and 4.11 · arXiv:2608.06155v1

Correct and complete

Each application derives the claimed conditional-density regularity from its explicit model assumptions, checks the needed integrability uniformly over the conditioning variable, and then invokes the abstract criterion with matching RKHS/Sobolev norms. No missing case or unsupported implication was found in these proof chains.

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:2608.06155v1

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.