

Correctness, proof completeness, and novelty audit

The Pila-Zannier strategy for Drinfeld modules and Drinfeld modular curves

Gal Binyamini, Dmitry Novikov, Francesco Maria Saettone

Paper	arXiv:2607.19952v1
Source version	Submitted 22 July, 2026; announced July 2026.
Audit date	August 15, 2026
Prompt	Standardized prompt
Novelty cutoff	July 22, 2026
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	Contains unsupported statements
02	Proofs	Contains incorrect or incomplete proofs
03	Novelty	No non-novelty findings

This independent report links to the source paper on arXiv; it does not reproduce or redistribute that paper.

01 / STATEMENTS

Statements

Contains unsupported statements

The two Ax-Lindemann conclusions and the Andre-Oort conclusion are supported, the latter after the explicit repair to the cusp cover recorded below. The stated Manin-Mumford theorem for arbitrary Drinfeld modules over C_∞ is not verified by the paper's proof: the preliminary specialization preserves each fixed torsion group scheme but does not preserve a Zariski-dense set of torsion points or lift specialness from the selected fiber. No counterexample to that theorem was found.

Theorem A · Drinfeld-module case (Theorem 3.23): Manin-Mumford over arbitrary C_∞ coefficients

Pages 2 and 31-33 · Section 3.3.1 and Theorem 3.23 · arXiv:2607.19952v1

Not able to verify

The counting argument is carried out only after replacing the given modules and curve by a specialization over a finite extension of F . The preceding paragraph proves that, for each fixed nonzero a , the finite étale group scheme of a -torsion keeps its cardinality in a suitable generic-characteristic fiber. It does not prove that one closed fiber preserves infinitely many torsion incidences on V , that their images remain Zariski dense, or that a torsion-translate conclusion for the fiber lifts to the original curve. These are exactly the implications needed for the sentence 'we may assume' in the proof. The paper notes that the conclusion is within reach of prior Manin-Mumford work, but it also notes a difference for products of distinct modules; no cited theorem is matched here to the full printed scope. The appropriate verdict is therefore inability to verify, not falsity.

Verifiable source: [Full paper, version 1](#)

Theorem A · hyperbolic case (Theorem 3.24): Andre-Oort for a product of Drinfeld modular curves when q is odd

Pages 2 and 34-39 · Theorem 3.24 · arXiv:2607.19952v1

Correct

The printed pigeonhole estimate uses the unnecessarily large center sets from Lemma 2.4 and therefore does not itself force a block. The lemma's construction admits a direct sharper version: for $z = a + b\sqrt{\Delta}$ with $|z|_i \leq R$, the polynomial part P_2 has degree at most $d = \lceil \log_q R \rceil$, while $|b| \leq R|\Delta|^{-1/2}$ implies that the truncated numerator P_1 in P_1/T^c also has degree at most d . Thus only $q^{2d+2} = R^{O(1)}$ centers are needed, independently of the field discriminant. Substituting these sets in (3.12) makes their total cost $D^{O(\varepsilon)}$, so the class-number lower bound $D^{1/2-\varepsilon}$ contradicts the rigid counting upper bound for sufficiently small exponents. The remaining height, local-uniformization, block-counting, and Ax-Lindemann steps then establish the stated conclusion.

Verifiable source: [Full paper, version 1](#)

Theorem B (Theorems 3.1 and 3.11): Ax-Lindemann in the linear and hyperbolic settings

Pages 2 and 19-28 · Sections 3.1 and 3.2 · arXiv:2607.19952v1

Correct

The lattice-counting and additive-polynomial argument in the Drinfeld exponential setting yields a weakly special image, and the stabilizer argument in the j -uniformization setting reduces the correspondence to an ordinary $PGL_2(F)$ conjugation after excluding a Frobenius twist by growth in the unipotent subgroup. No incorrect or unsupported central step was found in these two proofs. The harmless orientation convention for a negative Frobenius exponent can be handled by interchanging the two coordinates and requires no change to the result.

02 / PROOFS

Proofs

Contains incorrect or incomplete proofs

The proof of the Drinfeld-module Manin-Mumford case has an unresolved specialization gap. The Andre-Oort proof uses a cusp cover whose printed cardinality is too large for the claimed pigeonhole contradiction, but the construction itself gives a sharper cover that repairs the proof. No further central proof defect meeting the audit's evidence threshold was found.

Section 3.3.1: Specialization does not transfer the dense torsion problem

Page 31 · paragraph before and first paragraph of Theorem 3.23's proof · arXiv:2607.19952v1

Incomplete as written

After spreading out the coefficients, the text chooses one closed point and observes that every fixed a -torsion scheme remains finite étale of the expected size. It then immediately assumes that the original modules and V are defined over the residue field. This does not follow: preserving a finite group scheme one level at a time neither supplies one fiber carrying a dense collection of the original varying-order torsion points nor proves that specialness detected after specialization lifts to the generic/original fiber. A repair needs a genuine specialization theorem for the pair $(V, D \times D')$ or a separate descent argument covering the full C_∞ case.

Verifiable source: [Full paper, version 1](#)

Lemma 2.4 and Equation (3.12): The printed cusp-cover count cannot yield the asserted contradiction

Pages 10 and 35-39 · Lemma 2.4 and proof of Theorem 3.24 · arXiv:2607.19952v1

Incorrect as written · verified repair

Lemma 2.4 defines $d' = d + c$, with $c = \lceil \deg_T(\Delta)/2 \rceil$, and its displayed center set has cardinality $q^{2d'+2}$. The proof of Theorem 3.24 divides the Galois orbit by the product of two such cardinalities and then says that the result contradicts an H^n block count. That inference is invalid: the two cover costs include discriminant factors of roughly $|\Delta_1||\Delta_2|$ and can dominate the available $D^{1/2-\varepsilon}$ orbit lower bound. The repair is internal to Lemma 2.4. Its estimates give $\deg P_2 \leq d$ and, from $|b| \leq R|\Delta|^{-1/2}$, $\deg P_1 \leq d$ as well. Restricting both numerators to degree at most d preserves the cover and reduces it to q^{2d+2} . With this corrected cardinality, (3.12) has size $D^{1/2-O(\varepsilon)}$ and the block-counting contradiction follows.

Verifiable source: [Full paper, version 1](#)

Remaining proof chain: Functional-transcendence and counting inputs

Pages 7-30 and 31-39 · Sections 2 and 3 · arXiv:2607.19952v1

Correct and complete after the stated repair

Apart from the two issues isolated above, the torsion-height estimates, quadratic-point height bounds, rigid block-counting applications, algebraic-block extraction, and both Ax-Lindemann arguments are coherently matched to their uses. The small omission of first passing to a finite extension with stable reduction before applying Proposition 2.8 is a direct extension-invariance step and does not warrant a separate adverse finding under the prompt's mandatory gate.

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. [Full paper, version 1](#)

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.