

The Loopy Knot Walker, Take Two

What we predicted, what broke, and the knot that beats us on purpose

Matthew “Bear” Schulz · Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com *Preprint, not peer reviewed. CC BY 4.0. Draft v3, 2026-08-05.*

ABSTRACT

This is the second attempt. The first is in Appendix A, and it is there rather than in a separate deposit because a paper whose spine is “here is what we withdrew” wastes a reader’s afternoon. What survived from it is here as results; what did not is here as the reason the rules exist.

Five protein structures and a plan file written before any structure was fetched. It carries **four falsifiable predictions and one deliberate non-prediction** — the plan’s Test 2 asks whether the instrument behaves differently on hyperbolic versus torus knots and says, in its own words, “*We have no prediction. That’s the point.*” **Two of the four failed.** We report both, unsoftened, and the plan file ships in the crate so the registration can be read rather than taken on trust.

The failure that mattered: we registered that writhe must cancel on an amphichiral knot, because a figure-eight is its own mirror image and that is a theorem. **It did not cancel. It came back +42.542 and +7.039** on the two 4_1 structures — both positive, the sign that right-handed α -helices predict. The theorem is about the reduced knot diagram; we measured the open backbone curve. Both are correct and the gap between them is the finding, because it turns the measurement into a clean readout of helix contamination on a substrate where the knot’s own contribution is provably zero.

From that failure comes the result we think transfers furthest. Above a contamination threshold bracketed between $|Wr| = 7$ and 42, **replicate agreement hardens to 100% on an answer the knot type cannot have, while the instrument’s own internal vote gets weaker.** Agreement and certainty move in opposite directions. Nearly every field treats inter-replicate agreement as validation; here it is the failure signature, and the honest internal measure was sagging the whole time.

We also report a structure fetched under the wrong accession — 8BN1 is human ACE2, not the composite knot 8B1N — which produced **the most internally consistent result in the study, at purity up to 1.000 across three runs and six filenames.** The coherence was the symptom, not the safety. A single query to a reference our pipeline did not author corrected it immediately. Re-run on the right structure, the registered prediction confirms: both domains $\det = 3$, full chain $\det = 9$, all four seeds.

And we name the knot that beats us on purpose. The Kinoshita-Terasaka knot has the Alexander polynomial of the unknot. Our instrument reads $\det = 1$ and calls it unknotted at high purity — **exactly as the theory says it must**, on a knot built in 1957 to demonstrate that limit.

This is not a better knot detector and must not be read as one. AlphaKnot would have typed the composite correctly. What is here is a reliability characterisation of a deliberately weak channel, with the failures reported at full volume.

Keywords: holonomy · parallel transport · geometric phase · Berry phase · Wilson loop · winding number · linking number · Burgers vector · topological invariant · noise floor · loop walk · competence surface · knotted proteins · amphichiral · Alexander polynomial

WHAT THIS IS, IN PLAIN LANGUAGE

A protein backbone is a long tangled string. Sometimes it is genuinely knotted — you could not pull it straight without passing an end through a loop — and sometimes it only looks that way.

To find out, you close the string into a loop and compute a number that cannot change unless you cut it. That number is a **knot invariant**, and the one we use is the Alexander determinant. It behaves like an address: trefoil lives at 3, figure-eight at 5, and so on.

Two things you need to know before any result below makes sense.

An address is not a resident. More than one knot can live at the same determinant, so the number tells you where to look and not who is home. Section 5 is what happens when a knot exploits that on purpose.

A backbone is an open arc, not a loop. So it has to be closed artificially, many times, at random, and the answers voted on. The fraction of closures agreeing is called **purity** — and purity is *not* the same thing as being right. Section 2 is about how far apart those two can get.

Everything else in this paper is about where that instrument is competent and where it is not. The general method is in the companion field guide (10.5281/zenodo.21811589); this is the worked case on knots.

THE RESULTS, UP FRONT

#	what we did	the number	what it means
1	Registered that writhe cancels on 4 ₁	+42.542 and +7.039, both positive	the prediction failed; the residue is a clean measurement of helix contamination
2	Pushed 1QMG from 42 to 604 closures	agreement 50% → 100%, purity 0.834 → 0.759	agreement and certainty move apart; the wrong answer got <i>more</i> reproducible
3	Pushed 5VIK the same way	agreement 50% → 0%, purity 0.691 → 0.792	below threshold, more sampling recovers the honest hedge
4	Composite knot, right structure this time	both domains det = 3 , full chain det = 9 , 4/4 seeds	single-chain closure does see a connected sum
5	...and the limit that exposed	det = 9 is shared by 6 ₁ and 3 ₁ #3 ₁ ; the pipeline printed 6 ₁	the determinant is an address, never a resident
6	Writhe across the junction	16.394 + 13.904 = 30.298 vs measured 30.470	additive to 0.56% — the domains contribute independently
7	Competence surface, extended to 6 ₁	purity 0.907 → 0.679 as knots deepen	held, and the composite correctly sits off the curve
8	Mirror sweep, five structures	determinant invariant; writhe negated bit-identically	every seed that committed a handedness flipped
9	The Kinoshita-Terasaka knot	det = 1 , called unknot at high purity	the Alexander limit, met on the knot built to demonstrate it
10	Wrong accession fetched	purity to 1.000 , HIGH tier, 3 runs, 6 filenames	perfect internal consistency on the wrong protein

The first attempt, and everything it got wrong, is **Appendix A** — in this paper rather than a separate deposit.

WHAT THIS MEANS IF YOU WORK ON —

Protein topology. Two things are free and both are worth taking. **Suppress any handedness field for amphichiral knot types** — on 4₁ the knot has no handedness, so whatever gets reported is contamination, and α-helices supply a large positive bias. And **hedge det = 9 as “6₁ or 3₁#3₁”** exactly the way det = 5 is already hedged as “4₁ or 5₁.” Both are one line.

Anyone who validates by running it again. Replicate agreement is the signal nearly every field trusts. Section 2 is a molecular-scale case where agreement climbed to 100% on an impossible answer while the internal vote sagged. If your pipeline reports one confidence, it can’t show you this. **Report both.**

Anyone fetching structures by accession. A checksum proves the bytes arrived. It cannot prove you asked for the right thing, because it is computed on whatever you fetched. Section 8 is what that cost.

Methodologists. The general form is in the field guide. This is the case where every internal check the program had was structurally blind, because a wrong input makes every seed honestly wrong together.

SECTION 0 — WHY WE CAME BACK

(Author’s voice. Provenance, not proof. Skip it if you want the results.)

We came back because the first attempt made a claim that was wrong before the ink was dry, and I would rather say that plainly than let it sit.

Paper 1 closed by claiming the frame defect angle was a fast classifier for whether a protein is knotted. Then the companion nucleic-acid work measured how much that angle moves when nothing at all is wrong — 35.81° — against a knotted-versus-unknotted gap of 3.3° . **The effect was one eleventh of its own noise.** Nobody had to argue with the claim. The floor just walked in and sat on it.

That is not a story about being careless. It is a story about the order I did things in. I built an instrument, pointed it at knots, got numbers, and reasoned about what they meant — and never once asked the cheap question first.

So the return trip had three things the first attempt did not. A corrected instrument, because a factor-of-two error in the linking-number code turned up in the audit. A measured noise floor to check claims against. And **a plan file written before any of the five new structures was fetched**, so that failing would cost something. It registers four falsifiable predictions and one open question that it explicitly refuses to predict.

Two of the four failed. One of them failed because I confused a theorem about a knot with a measurement on a protein, and that failure taught more than the three that worked. The other failed because I fetched the wrong file — one transposed character in an accession — and got the most beautifully consistent wrong answer in the whole program.

That second one is the reason I stopped trusting internal consistency as evidence of anything. Every check we run compares the pipeline to itself. **A wrong input makes every seed honestly wrong together**, and the more of them you run, the more confident the wrongness looks. What caught it was one query to something we did not write.

The first attempt is Appendix A, at the back of this paper rather than in a deposit of its own. What survived from it is reported here as results; what did not is the reason the rules exist.

Knots were a good place to learn this because knots do not negotiate. A knot is a knot; there is no interpretive latitude, no p-value, nothing to spin. When the instrument is wrong about a knot, it is simply wrong, and the mathematics will say so in one line.

METHODS, BRIEFLY

The trace. $C\alpha$ atoms read as an ordered curve. Gaps declared, never bridged.

Knot typing. KMT reduction (Koniaris & Muthukumar 1991) then the Alexander determinant at $t = -1$, over an ensemble of stochastic far-apex closures in the lineage of Millett, Dobay & Stasiak (2005). **Purity** is the fraction of closures returning the modal answer.

Writhe. Klenin-Langowski Gauss double integral, cross-checked against a second independently written implementation: agreement 9.77×10^{-15} on the trefoil.

Mirrors. Generated by negating the x-coordinate of every ATOM record. A mirror image must return the same determinant and exactly negated writhe; both are checked on every structure.

The registration, and where the paper departs from it. The plan file ships in this crate at `03_code/loopy_knot_walker_v2_plan.md`. Read against this paper it differs in three ways, and all three are ours to declare rather than the reader's to find:

- **Its Tests 4 and 5 are the reverse of ours.** The plan numbers mirror mathematics as Test 4 and slipknots as Test 5; earlier drafts of this paper swapped them. The content of each is unchanged; only the labels moved.
- **Its Test 2 registers nothing.** The plan asks whether purity degrades differently on hyperbolic versus torus embeddings at the same crossing number and answers *"We have no prediction. That's the point."* What this paper reports for $4N2X$ — $\det = 9$, purity below shallower knots — was registered later, when the structure was selected, and is scored as a confirmation of the competence surface rather than of the plan.
- **The plan's Test 2 question is still open.** It wanted hyperbolic against torus at *matched* crossing number. 4_1 (hyperbolic, $\det = 5$) and $T(2,5)$ (torus, $\det = 5$) are the pair that would answer it, and that comparison is not run here. It is in Section 10 where it belongs.

Runs of record. Every number below names its run.

run	closures	seeds	total	contents
A	21	2	42	5 originals
B	21	2	42	5 originals + 5 mirrors
C	151	4	604	5 originals, no mirrors

Run A is byte-identical to three other archived copies — one run under four filenames, md5 35e52e57e949a04de994e24252c58791.

Identity. Every accession’s deposited HEADER, TITLE, SOURCE and DBREF records were read and a per-chain C α census taken **before any topology was computed**. That check exists because of Section 8.

1. THE PREDICTION THAT FAILED

Registered, before any structure was fetched: $|Wr| \approx 0$ on both 4₁ structures. The figure-eight knot is amphichiral — identical to its own mirror image — so its writhe must cancel.

That was not a guess. It is a long-known theorem: for an oriented alternating link that is achiral, the writhe of a reduced projection is exactly zero. Cerf (1997), Prop. 9 gives a short proof of it.

accession	N (C α)	registered	measured Wr	det
1QMG	514	≈ 0	+42.542	5
5VIK	296	≈ 0	+7.039	5

The determinant prediction held. **The writhe prediction failed, on both structures, by a wide margin.**




Figure 1. The registered value against what came back. We are reporting the failure rather than restating the registration, because the gap is the result.

Why it failed, and why that is worth more than a pass. The theorem is about the **reduced knot diagram**. We measured the **open-arc writhe of the backbone curve**. Both are correct, and they are about different objects.

What fills the gap is also published: protein α -helices are almost invariably right-handed and therefore contribute a large, **positive** writhe (Panagiotou et al., arXiv 1812.08721). Our two measurements are both positive — the predicted sign — on two structures where **the knot contributes no handedness at all**.

Writhe cancels for the *knot*. It does not cancel for the *backbone carrying the knot*. So the measured $|Wr|$ is a direct readout of helix contamination on a substrate where the knot’s own contribution is known to be zero, which is a cleaner contamination measurement than anything in the first attempt — **and it is only available because the prediction was wrong**.

Corroborated where it can be corrected. One far-apex closure of 2EFV, a knot that *does* have a true handedness, gives closed-curve writhe **−3.31**, the correct left-handed sign, against the contaminated open-arc value of **+2.06**, which reads the wrong way. Same mechanism, observed once where closure recovers the answer and twice where there is nothing left to recover.

2. AGREEMENT AND CERTAINTY MOVE APART

This is the finding we think transfers furthest, and it survives the obvious objection because both numbers are in the table.

The pipeline reports **two independent confidences**. *Purity* is the internal vote — what fraction of closures agree. *Replicate agreement* is whether independently seeded runs reach the same answer. Nearly every field treats the second as validation.

	seeds committed	purity	det vote margin
1QMG @ 42 closures	1/2 (50%)	0.810–0.857	0.714–0.762
1QMG @ 604	4/4 (100%)	0.722–0.795	0.609–0.742
5VIK @ 42 closures	1/2 (50%)	0.667–0.714	0.476–0.571
5VIK @ 604	0/4 (0%)	0.742–0.841	0.556–0.742




Figure 2. On 1QMG the internal vote got *weaker* while replicate agreement rose to 100% — on a handedness a 4₁ knot cannot have. On 5VIK every arrow points the other way: agreement collapsed to zero while purity rose, the instrument becoming internally more certain of its own uncertainty, which is the correct behaviour on an amphichiral knot.

One mechanism, one threshold, two sides. The chirality call tracks the sign of the *contaminated* writhe, not the knot’s handedness. Large contamination (1QMG, |Wr| = 42.5) pegs the majority vote; small contamination (5VIK, |Wr| = 7.0) sits at the margin and washes out to honest indeterminacy with more sampling. The two structures differ in magnitude, not in kind, and they fall on opposite sides of the vote-pegging threshold.

Why this matters beyond two proteins. Run it again with a different seed, a different operator, a different lab — if the answers agree, you believe them. Here **agreement is the failure signature**, and the honest internal measure was telling the truth the whole time by sagging. The instrument never claimed to be more certain. Only the replicates did.

A prediction this generates, not tested here. A 4₁ structure with intermediate |Wr| — roughly 15 to 30 by these two anchors — should be seed-inconsistent at 604 closures rather than converging either way. Independently, recomputing writhe with α -helical residues excluded should collapse |Wr| toward zero on both and break the 1QMG lock at fixed knot type.

3. THE COMPETENCE SURFACE, AND THE COMPOSITE

4N2X (6₁ Stevedore’s knot) is the first hyperbolic knot in the real-protein corpus and the first 6₁ through this pipeline: det = 9 stable across all four seeds, purity 0.656–0.702, handedness consistently right-handed — a chiral knot behaving correctly, in contrast to the two amphichiral structures above.

The composite, on the right structure this time. Run on 8B1N (TrmD-Tm1570), identity verified from the deposited headers before the run:

trace	residues	det	knot	purity (4 seeds)	open-arc Wr
TrmD	1–240	3	trefoil 3 ₁	0.901–0.954	+16.394
Tm1570	241–430	3	trefoil 3 ₁	0.801–0.894	+13.904
full chain	–2–430	9	see below	0.815–0.874	+30.470

Registered det = 3 on each domain: **confirmed 4/4 and 4/4**. Registered det = 9 or det = 1 on the full chain: **det = 9, 4/4**. Single-chain far-apex closure does see the connected sum.

Writhe is nearly additive across the junction: $16.394 + 13.904 = 30.298$ against a measured 30.470 — a residual of **+0.172, or 0.56%**. Two knotted domains contributing their geometry independently, which is what a connected sum whose factors do not interpenetrate should do.




Figure 3. Purity degrades monotonically as a prime knot gets deeper. The composite sits **off** that curve at 0.815–0.874, above 4N2X's 6_1 despite carrying more crossings — because a connected sum of two shallow knots is not the same object as one deep knot, and should not be read as sitting on the same curve.

A defect found, fixed, and re-run — with the receipt. The first extractor did not filter the altLoc column, so nine residues carrying alternate conformations were counted twice and reported as N_CA 245 and 194. The extractor was corrected the same day and the analysis re-run; the run reported above prints **240 and 190**, matching the altloc-filtered header census exactly. Every purity and writhe in the table comes from the corrected run (receipts/composite_8b1n_log_20260801_094511.txt), and all three identity scripts in the crate apply the same one-line filter.

That the determinants were stable across four seeds *and* across a 55-residue difference in domain size is worth stating, because it is the reason the error was survivable: **the knot calls were never sensitive to it**. The counts were wrong; the topology was not.

4. THE SLIPKNOT, AND THE MIRROR SWEEP

2P4Z (slipknot) reads unknot at high purity — 0.905–1.000 at 42 closures, 0.874–0.940 at 604 — stable across every seed. Prediction confirmed. The knot lives in a subchain; single-chain closure on the full backbone correctly finds unknot.

The mirror sweep exists only in run B: 42 closures per structure. Every claim here is at that count and stated as such.

Determinant — invariant. Confirmed on every original/mirror pair across both seeds.

Writhe — antisymmetric. Confirmed to bit-identical negation on all five pairs; +42.54168935833795 against –42.54168935833795, and so on down.

Chirality — reported per seed, not merged. Every seed that committed to a handedness flipped under mirroring — on the chiral knot and on both amphichiral knots alike. On 4N2X the flip is correct, because 6_1 is chiral. On 1QMG and 5VIK it is spurious, because a 4_1 knot has no handedness to report. **Same operation, correct on one, contamination on the other two — and mirroring is what separates them.**

(The first attempt reported one handedness per structure and stated 5VIK was “indeterminate both ways.” The receipts have two seeds and that is false: seed 1 reads right-handed and its mirror reads left-handed, at purity 0.571, confidence tier LOW. Corrected here.)

5. THE KNOT THAT BEATS US ON PURPOSE

The Kinoshita-Terasaka knot (K11n42) and the Conway knot (K11n34) are mutant knots with eleven crossings sharing the same Alexander, Jones and HOMFLY polynomials. **Their Alexander polynomial equals 1 — the same as the unknot.**

Our instrument returns $\det = 1$ on both and classifies them as unknotted at high purity.

This is not an implementation error. Kinoshita and Terasaka constructed this knot in 1957 specifically as an example of a nontrivial eleven-crossing knot with Alexander polynomial 1. The instrument fails exactly where the theory says it must.




Figure 4. The determinant gives a knot an address, not a resident. $\det = 1$ is shared by the unknot and by two knots built to demonstrate that the Alexander polynomial cannot see them. $\det = 9$ is shared by 6_1 and $3_1\#3_1$ — and our pipeline printed 6_1 for a structure that is $3_1\#3_1$. **The number was right; the label was wrong.**

What propagates from this. 4N2X’s identification as 6_1 rests on the determinant alone, and we now hold a second structure at $\det = 9$ that is not 6_1 . **That assignment stands on KnotProt’s independent call, not on ours.** Recommended fix, one line: hedge $\det = 9$ as “ 6_1 or $3_1\#3_1$ ” exactly as $\det = 5$ is hedged.

What the instrument can still say. Neither KT nor Conway reaches $\det = 3$. If a structure reads trefoil cleanly, it is not the KT knot. **The instrument is a reliable negative for KT — it can rule out a suspected identification, and cannot confirm one.** That is narrower than the first attempt claimed (Appendix A). It is also true.

Not yet run: building the K11n42 curve synthetically and confirming $\det = 1$ at high purity directly. This section reasons from the published Alexander polynomial, not from our own measurement, and is labelled accordingly.

6. WHERE THIS SITS RELATIVE TO THE EXISTING TOOLS

Stated plainly, because the reader deserves it before the discussion.

AlphaKnot 2.0 types knots via **Topoly’s** HOMFLY-PT polynomial, reaches **twelve crossings**, and **reports composite knots**. **KnotProt 2.0** reaches eight. Both offer probabilistic and deterministic closure with user-selectable counts — the same stochastic machinery we use, and the lineage is Koniaris & Muthukumar (1991) and Millett, Dobay & Stasiak (2005).

AlphaKnot would have typed 8B1N correctly. Our pipeline reads the Alexander determinant at $t = -1$, which cannot separate 6_1 from $3_1\#3_1$ at $\det = 9$. HOMFLY-PT can. **This paper does not present a better knot detector and must not be read as one.**

What it presents is a reliability characterisation of a deliberately weak channel: noise floors measured on the *substrate* rather than the algorithm, a competence surface, mirror controls, and pre-registered predictions reported including the ones that failed. Those are the parts these tools do not publish.

7. THE TOWER BELOW THE KNOTS

Knot determinants are integers, so a knot has an address in the integer tower of the companion paper. $T(2,q)$ torus knots have $\det = q$, sitting at prime rungs with $\Omega = 1$ each. The 4_1 figure-eight has $\det = 5$ — the same rung as the cinquefoil $T(2,5)$. Different knots, same address.

Both sides of the connected-sum identity were measured here, on the same molecule. Two domains at $\det = 3$, $\text{rung } \log 3 = 1.0986$, $\Omega = 1$ each. Their connected sum at $\det = 3 \times 3 = 9$, $\text{rung } \log 9 = 2.1972 = 2 \times \log 3$, $\Omega = 2$. The full chain measures $\det = 9$. **Multiplication in determinant space is addition in log space**, and the residual is $0.000\text{e}+00$ because determinant multiplicativity is a theorem rather than an approximation.

Set that beside the writhe residual across the same junction — $+0.172$, or 0.56% — and the paper’s two channels are visible on one object. **The integer channel’s addition is exact and algebraic. The continuous channel’s is approximate and analytic, and its leftover measures the geometric coupling at the junction.**

8. WHAT THE MISIDENTIFICATION COST

The composite-knot result in the first attempt was **the most confident result in the study**: $\det = 1$ across every seed at both closure counts, across the mirror, in three independent runs archived under six filenames. Confidence tier HIGH. Purity up to 1.000. Perfect internal reproducibility.

It was human ACE2. We requested **8BN1** believing it to be the TrmD-Tm1570 composite knot. **The composite is 8B1N** — one transposed character.

	8B1N	8BN1
TITLE	TrmD-Tm1570 from <i>C. nitroreducens</i>	ACE2 with bicyclic peptide inhibitor
SOURCE	<i>C. nitroreducens</i> , gene TRMD	<i>Homo sapiens</i> , gene ACE2
chain A Cα	433	595

Every number the pipeline recorded matches **8BN1** exactly and matches **8B1N** on none of them. **The misidentification is proven from the deposited headers, not inferred.**

Why no internal check could have found it. Every check we run is a form of disagreement-vetting — seed sweeps, mirror controls, purity tiers, cross-run stability. Each is structurally blind to a wrong input, because **a wrong input makes every seed honestly wrong together**. The pipeline had no representational slot for “this is not the protein you named.” The coherence of the result was the symptom, not the safety.

What caught it was a single query to an exterior reference — the RCSB entry title. It corrected the error immediately and completely.

A compounding detail worth carrying. ACE2 is numbered from residue 19, so a domain cut written as “residues 1–240” silently became “residues 19–240” and 181 Cα were never analysed. **Nothing in this corpus starts at residue 1** — 1QMG begins at 82, 4N2X at 4, 2P4Z at 10, 5VIK at 19. A coverage assertion costs one line and would have caught it.

The rule this buys, at the cost of one withdrawn section: every accession must carry an out-of-band identity assertion, checked against the deposited file before any topology is computed. Print HEADER, TITLE and a per-chain census. Assert that any subset you carve covers the range you claim.

This is the section the first attempt got wrong, and the one this paper withdraws. See Appendix A.

The literature is itself split on this accession — at least one paper prints **8BN1** for TrmD-Tm1570 while RCSB and the force-spectroscopy literature print **8B1N**. We inherited a propagated typo and did not check it. Our own research plan flagged the ambiguity and the check was never run.

9. WHAT ELSE THE AUDIT CORRECTED IN THE FIRST ATTEMPT

Every item below is a correction to the unpublished first paper. Appendix A says what that paper was, what survived it, and why it is not deposited.

A factor of two in the linking-number code. `cwf_ribbon.mutual_linking()` divided by 2π ; the Gauss linking integral for two *distinct* closed curves normalises by 4π . It returned exactly twice the true linking number. Validated on five curves of known answer — planar circle, Hopf link, unlink, (2,4) torus link, Solomon’s seal — where 4π returns the true value in all five and 2π returns double in every linked case.

And the residual it left is real geometry. The corrected trefoil gives $Lk - Wr = -0.4877$, one half-turn to within 1.2%. Measuring the Bishop frame's parallel-transport holonomy directly — carrying the final normal across the closing segment, with no linking integral anywhere in the calculation — gives 175.559° , or **0.4877 of a turn. The two routes agree to 1.8×10^{-13} .**

So the trefoil decomposes as $Wr = -3.5123$, $Tw = -0.4877$, $Lk = -4.0000$. Neither continuous term is near a round number; their sum is exactly an integer. **That is Călugăreanu-White-Fuller demonstrated rather than cited.**

Mixed statistics in the first attempt's Table 6. Five of seven defect-angle rows were single fixed-seed measurements; one was a multi-seed mean ($108.6 \pm 38.3^\circ$). Different quantities. All seven now reported as multi-seed means with spreads.

The defect-angle classifier claim: withdrawn. Floor 35.8° , gap 3.3° , SNR = 0.09. The defect angle remains a real geometric quantity with a measured domain — it correctly characterises holonomy on closed synthetic curves — but it cannot classify open-backbone knottedness. **The corrected framing is a geometric probe with a measured noise floor, not a classifier.**

10. OPEN QUESTIONS

- Does a KT-like knot — nontrivial but Alexander-trivial — exist in any protein, undetected because every standard instrument shares this blind spot?
- Does KMT reduction behave differently on hyperbolic versus torus embeddings at the same determinant? 4_1 and $T(2,5)$ both read $\det = 5$ with different topology.
- Could a different invariant (Jones, knot Floer homology) detect what Alexander cannot, on a real structure?
- The intermediate- $|Wr|$ prediction of Section 2, and the helix-exclusion control that should break the 1QMG lock at fixed knot type.
- The K11n42 curve, built synthetically and measured rather than reasoned about.

HOW TO BREAK THIS

Every structure here is a public accession and every tool in the lineage has a published implementation. Fetch 1QMG, compute the open-arc writhe, and see whether you get +42.542. Run the mirror and check the negation is exact.

The interesting attacks are not arithmetic. **Is the vote-pegging threshold real, or an artefact of two structures that differ in size as well as contamination?** Does the agreement-versus-purity divergence appear on chiral knots too, where there is a right answer to converge on? Would excluding α -helical residues actually collapse $|Wr|$, or is the contamination distributed more widely than the helices?

If a number here does not reproduce, that is this paper's problem and we want to know. Breaks get logged in the open with attribution.

APPENDIX A — THE FIRST ATTEMPT

This is a summary of unpublished work, not a reproduction of it. It is here so the corrections in Sections 5, 8 and 9 have something to point at, and so nobody has to find a separate deposit to learn what we got wrong.

What it was. “The Knot Feels No Time” applied the loop-holonomy instrument to protein backbone topology and recovered the catalogued knot type and chirality of a minimal test pair — 2EFV, a left-handed trefoil, and 1UBI, an unknotted control — from $C\alpha$ coordinates alone. Corpus: 2EFV, 1UBI, 1J85, 1K3R, 2CAV, 1TIM. That recovery was never the contribution; it was the entry fee.

What it claimed, in three parts.

1. **The competence surface** — purity degrades with crossing number independent of arc truncation, and beyond a threshold complexity additional closures within one seed cannot rescue a systematically wrong verdict, while additional independent seeds can.
2. **The chirality mechanism, demonstrated rather than asserted** — a controlled local-wobble perturbation carries raw writhe through zero while a reduced-diagram invariant holds the correct handedness throughout, isolating secondary-structure contamination as the specific cause.
3. **The self-linking number under a quaternion-propagated Bishop frame**, with its closure-defect angle offered as a fast classifier for knottedness.

What survived, and is reported in this paper as results. Points 1 and 2, both strengthened: the competence surface now spans 0_1 , 3_1 , 4_1 , 5_2 and 6_1 rather than torus knots alone, and the contamination mechanism has a measured magnitude threshold and a controlled mirror-pair experiment behind it. The corpus results stand. The Knoto-ID cross-check stands at 3/3.

What did not survive.

- **The classifier claim (point 3) is withdrawn.** Its floor was measured after publication of the draft at 35.81° against a 3.3° effect.
- **All Bishop-frame linking numbers were doubled,** by a factor-of-two normalisation error found in the audit.
- **The composite-knot section measured the wrong protein.**
- **One table mixed single-seed and multi-seed statistics** in the same column.

Why it is not deposited separately. It never made a claim in a live DOI, so there is nothing to retract publicly. Depositing it now would put a paper into the record whose headline result is already withdrawn, and send anyone who found it on a chase through two documents to reach one answer. **The honest thing is one paragraph and an appendix, which is what this is.**

The first attempt was not wasted. Every rule in the companion field guide traces to something it got wrong, and rules without their violations are not believable.

REFERENCES

On the tags below. Every entry marked **Verified 2026-08-05** was checked against the publisher record on that date. **Every entry in this list has now been checked** — the nine that carried a working tag, and the seven that carried no tag at all and were therefore the more dangerous half. An untagged, unchecked reference is the failure mode this list exists to prevent. **A citation quoted from memory is a claim, and this corpus does not make claims it has not checked.**

One error that check caught, recorded rather than quietly fixed: an earlier draft of this list gave AlphaKnot 2.0 the author line belonging to AlphaKnot 1.0. Both are now cited separately and correctly. It is the same class of error as Section 8 — a name attached to the wrong object — found the same way, by asking a source we did not author.

The knots

- Kinoshita, S. & Terasaka, H. (1957). On unions of knots. *Osaka Math. J.* 9(2), 131–153. **Verified 2026-08-05.**
- Piccirillo, L. (2020). The Conway knot is not slice. *Ann. of Math. (2)* 191(2), 581–591. doi:10.4007/annals.2020.191.2.5. **Verified 2026-08-05.**
- Thurston, W. P. (1982). Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)* 6(3), 357–381. doi:10.1090/S0273-0979-1982-15003-0. **Verified 2026-08-05.**
- Cerf, C. (1997). Nullification writhe and chirality of alternating links. *J. Knot Theory Ramifications* 6, 621–632. doi:10.1142/S0218216597000340. Preprint: arXiv q-alg/9612005. **Prop. 9** — if an oriented alternating link is achiral, the writhe of a reduced projection is zero. Cerf records the result as already long known and supplies a short proof; her own contribution is the nullification writhe, which this paper does not use. **Verified 2026-08-05**, proposition read against the source rather than cited from memory.

Instrument lineage

- Koniaris, K. & Muthukumar, M. (1991). Self-entanglement in ring polymers. *J. Chem. Phys.* 95(4), 2873–2881. doi:10.1063/1.460889. **Verified 2026-08-05.**
- Millett, K., Dobay, A. & Stasiak, A. (2005). Linear random knots and their scaling behavior. *Macromolecules* 38, 601–606. **Verified 2026-08-05** — title is US spelling in the original.
- Klenin, K. & Langowski, J. (2000). Computation of writhe in modeling of supercoiled DNA. *Biopolymers* 54(5), 307–317. doi:10.1002/1097-0282(20001015)54:5<307::AID-BIP20>3.0.CO;2-Y. **Verified 2026-08-05.**
- Călugăreanu, G. (1961). Sur les classes d'isotopie des nœuds tridimensionnels et leurs invariants. *Czechoslovak Math. J.* 11(4), 588–625. doi:10.21136/CMJ.1961.100486. **Verified 2026-08-05** — title supplied; the entry previously carried none.
- White, J. H. (1969). Self-linking and the Gauss integral in higher dimensions. *Amer. J. Math.* 91, 693–728. **Verified 2026-08-05** — title supplied.
- Fuller, F. B. (1971). The writhing number of a space curve. *PNAS* 68(4), 815–819. doi:10.1073/pnas.68.4.815. **Verified 2026-08-05** — title supplied.

- Bishop, R. L. (1975). There is more than one way to frame a curve. *Amer. Math. Monthly* 82(3), 246–251. doi:10.1080/00029890.1975.11993807. **Verified 2026-08-05.**

The prior tools — better at naming a knot than this instrument is

- Rubach, P., Sikora, M., Jarmolinska, A. I., Perlinska, A. P. & Sulkowska, J. I. (2024). AlphaKnot 2.0: a web server for the visualization of proteins' knotting and a database of knotted AlphaFold-predicted models. *Nucleic Acids Res.* 52(W1), W187–W193. doi:10.1093/nar/gkae443. **Verified 2026-08-05.**
- Niemyska, W., Rubach, P., Greń, B. A., Nguyen, M. L., Garstka, W., Bruno da Silva, F., Rawdon, E. J. & Sulkowska, J. I. (2022). AlphaKnot: server to analyze entanglement in structures predicted by AlphaFold methods. *Nucleic Acids Res.* 50(W1), W44–W50. doi:10.1093/nar/gkac388. **Verified 2026-08-05.** The earlier server — cited separately because an earlier draft of this list gave its author line to the 2024 paper.
- Dabrowski-Tumanski, P., Rubach, P., Goundaroulis, D., Dorier, J., Sułkowski, P., Millett, K. C., Rawdon, E. J., Stasiak, A. & Sulkowska, J. I. (2019). KnotProt 2.0: a database of proteins with knots and other entangled structures. *Nucleic Acids Res.* 47(D1), D367–D375. **Verified 2026-08-05.**
- Dabrowski-Tumanski, P., Rubach, P., Niemyska, W., Greń, B. A. & Sulkowska, J. I. (2021). Topoly: Python package to analyze topology of polymers. *Brief. Bioinform.* 22(3), bbaa196. doi:10.1093/bib/bbaa196. **Verified 2026-08-05** — author line was already correct; title supplied.
- Dorier, J., Goundaroulis, D., Benedetti, F. & Stasiak, A. (2018). Knoto-ID: a tool to study the entanglement of open protein chains using the concept of knotoids. *Bioinformatics* 34(19), 3402–3404. doi:10.1093/bioinformatics/bty365. **Verified 2026-08-05.**

Mechanism and cross-checks

- Panagiotou, E. & Plaxco, K. W. (2020). A topological study of protein folding kinetics. *Topology and Geometry of Biopolymers*, AMS Contemp. Math. Ser. 746, 223–233; preprint arXiv 1812.08721 — right-handed helices contribute large positive writhe. **Verified 2026-08-05** — two authors, not *et al.*; the preprint was published.
- Dietler, G., Kusner, R., Kusner, W., Rawdon, E. & Szymczak, P. (2020). Chirality for crooked curves. arXiv 2004.10338 — polygonal chirality; the middle eigenvalue of Q_R for a ropelength-critical 4_1 configuration is 5×10^{-16} , detecting amphichirality. **A comparison method we have not run. Verified 2026-08-05** — preprint, not journal-published as of this date; authors listed alphabetically in the original.
- Rusková, R. & Račko, D. (2021). Channels with helical modulation display stereospecific sensitivity for chiral superstructures. *Polymers* 13(21), 3726. doi:10.3390/polym13213726; PMCID PMC8588256 — helical channels show no stereoselective effect on achiral knots. **Verified 2026-08-05.**

Companion deposits

- Schulz, M. (2026). *Reading Nucleic-Acid Backbones with a Loop-Holonomy Instrument*. Zenodo 10.5281/zenodo.21815320.
- Schulz, M. (2026). *Loop Walking: A Field Guide*. Zenodo 10.5281/zenodo.21811589.
- Schulz, M. (2026). *The Tower Is Coordinate-Free*. Zenodo 10.5281/zenodo.21798311.
- Schulz, M. (2026). *Taking a Loop Walk Through the Integer Forest*. Zenodo 10.5281/zenodo.21449570.
- Schulz, M. (2026). *Walking the Loop: A Quick Guide to Holonomy*. Zenodo 10.5281/zenodo.21448987.
- Schulz, M. (2026). *Reading the Sky Through Loops*. Zenodo 10.5281/zenodo.21331652.

DEPOSIT METADATA

Author Matthew “Bear” Schulz · Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com **License** CC BY 4.0 · **Type** Preprint · **Date** 2026-08-05

Corpus identity, verified from deposited headers before any topology claim

acc	deposited title	chain A Cα	res span
1QMG	acetohydroxyacid isomeroreductase, <i>S. oleracea</i>	514	82-595
5VIK	miRFP703 near-infrared fluorescent protein, <i>R. palustris</i>	296	19-314
4N2X	DL-2-haloacid dehalogenase, <i>Methylobacterium</i>	298	4-301
2P4Z	ferredoxin-like metallo-β-lactamase superfamily, <i>T. tengcongensis</i>	275	10-284
8B1N	TrmD-Tm1570, <i>C. nitroreducens</i>	433	−2-430
8BN1	ACE2 + peptide inhibitor, <i>H. sapiens</i> — not used	595	20-614

Prior-art status. A credentialed zbMATH / MathSciNet pass has not been run. The search covered arXiv, PMC, publisher pages and OSTI. Novelty language is restricted to “not located,” never “new.”

Reproducibility. Every number in this paper appears in a tool output; nothing lives only in prose.