

Reading Nucleic-Acid Backbones with a Loop-Holonomy Instrument

How much does a molecule move when nothing is wrong — and what that number rules out

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Schulz · Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com *Preprint, not peer reviewed.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

A protein or nucleic acid in solution is not a shape. It is a **cloud of shapes** — every model in an NMR ensemble is a legitimate fit to the same data, and the molecule really does occupy all of them at room temperature. We measured how wide that cloud is, on three substrates, and the number turns out to decide what any geometric measurement on those substrates is allowed to claim.

Across 1,172 protein molecules and 24,864 conformers, the frame defect angle moves 35.81° when nothing is wrong. The difference between a knotted protein and an unknotted one, on that same channel, is 3.3°. The feature is one eleventh of the cloud it sits in. Four separate experiments had already been run on that channel; all four were predictable failures, without labels, from a single unattended measurement that takes about a minute.

We report that floor and two more — **0.0640 for RNA [writhe], 0.1886 for protein, ~0.18 for DNA duplex axes** — and note that our registered expectation about which substrate would be noisiest was backwards. Alongside them: a matched positive control (the RydC RNA trefoil, reproduced on 4/4 seeds and confirmed by an independent tool, plus the new result that the knot survives all 24 plausible completions of its unmodelled gap, for a mechanistic reason); a **DNA census of 250 curves returning zero knots**; a citation trap that stops anyone following the literature; a gap artefact that injects 83 units of fabricated writhe at 460× the floor, large enough to invert a sign; two sign-inverting bugs in our own axis construction, both caught by a synthetic of known answer; a symmetry blind spot that is one line of algebra and the difference between AUC 0.98 and 0.51; and a factor-of-two error in our own linking-number code whose residual resolved to the Bishop frame’s own closure holonomy, agreeing with the topological route to 1.8×10^{-13} .

The instrument works. It works at a narrower thing than we hoped, and we can now say exactly how narrow, with a number.

Keywords: holonomy · parallel transport · geometric phase · Berry phase · Wilson loop · winding number · linking number · Burgers vector · topological invariant · noise floor · loop walk · competence surface · nucleic acid topology · writhe · NMR ensemble

WHAT THIS IS, IN PLAIN LANGUAGE

Skip this if you already know what a loop-holonomy instrument is. Nobody does, so probably don’t.

Take a curve — a strand of RNA, the axis of a DNA duplex, the backbone of a protein. Walk along it carrying a little frame, like a compass held at arm’s length, turning it as little as possible at each step. Go all the way around and come back to where you started.

The frame will not be pointing where it was when you left. Nothing rotated it on purpose. The shape of the path did it. That mismatch is the **holonomy**, and it is the only thing this instrument measures. A Foucault pendulum is this. So is light through a coiled fibre, and a cat landing on its feet.

Two very different kinds of answer come back from that walk.

The integer channel counts something that cannot change unless you cut the curve — how many times it threads through itself, which knot it is. A whole number. No amount of gentle jiggling moves it.

The continuous channel measures how much winding piled up. A real number with decimals, and it changes the instant anything moves.

The temptation is enormous, because the continuous channel is fast and the integer channel is slow. Surely the fast one can screen for the slow one.

It cannot, and here is why — this is the whole paper.

The molecule is vibrating. Not metaphorically: an NMR ensemble is a set of conformations the thing actually occupies at room temperature, each one a legitimate answer. So every time you walk the loop you are walking a slightly different curve, and you get a slightly different number. The structure you think you are measuring is more

like an electron's probability cloud than a wire model — it condenses into something that looks like a fixed object only because you averaged, or because someone drew you a single picture of it.

The width of that cloud is the noise floor. And the rule that falls out is brutally simple:

You cannot read a feature narrower than the cloud it sits in.

Not “it is difficult.” Not “you need more samples.” A floor that belongs to the detector can be beaten with a better detector. **A floor that belongs to the molecule cannot, because there is nothing on the other side of it to resolve.**

Our knotted-versus-unknotted difference was 3.3° . The cloud was 35.81° wide. Everything else in this paper follows from that one comparison.




Figure 1. *Left:* twenty-six models of one molecule, all legitimate fits to the same experimental data. *Right:* the same picture as a distribution. The wide band is the width of the cloud — how far the frame defect angle wanders when nothing is wrong. The narrow red pair is the entire knotted-versus-unknotted difference the channel was being asked to resolve. This is not a hard measurement. It is an impossible one, and it is knowable in advance.

THE RESULTS, UP FRONT

#	what we measured	the number	what it rules out
1	Protein conformational floor , 1,172 molecules / 24,864 conformers	35.81° defect angle, 0.1886 writhe	any defect-angle classifier: the knotted/unknotted gap is 3.3°, SNR 0.09
2	RNA floor , 220 molecules / 3,214 conformers	0.0640 writhe , 0.1258 ACN, 14.25° defect	our own registered expectation — RNA is <i>tighter</i> , not noisier
3	DNA duplex floor , crystal replicate series	~0.18	a published “0.19° from canonical” agreement, which is 0.14× this
4	DNA knot census , 250 curves	0 knotted , 0.00%	the premise we started with
5	Gap contamination , 1S72 23S rRNA	83 units of fabricated writhe, 460× the floor	trusting any writhe computed across bridged gaps
6	Open-arc seed floor , one unchanging structure	12.653° across three starting frames	single-seed reporting of any open-arc frame quantity
7	Amphichiral blind spot , 51-curve corpus	AUC 0.9808 chiral vs 0.5128 amphichiral	writhe as a general knot screen — 4 ₁ is common in proteins
8	Positive control , RydC sRNA	trefoil on 4/4 seeds, 24/24 gap completions	the worry that gap modelling changes this particular answer
9	Linking-number code error , ours	returned exactly 2× Lk	anything built on that function — nothing was
10	The residual it left , resolved	0.4877 two ways, agreeing to 1.8×10⁻¹³	the guess that it was numerical slop

WHAT THIS MEANS IF YOU WORK ON —

Structural biology. Before you compare a geometric descriptor between two structures, measure how far it moves across the models *inside* one ensemble. That measurement needs no labels, makes no topology call, and took us about a minute of unattended compute. If your effect is smaller than that spread, the comparison cannot work and you can know it before you run it. And **never compute writhe across a bridged gap** — a straight line drawn through missing residues is not a modelling approximation here, it is a fabricated 83 units.

Anyone doing knot detection. We are not competing with you. KnotProt, Knoto-ID, AlphaKnot and Topoly detect knots better than this instrument does, and where our results overlap yours they agree. What we add is a number none of those tools publish: how much the *molecule* moves, as opposed to how stable the *algorithm’s* vote is. Those are different quantities and only one of them tells you whether a small difference means anything.

Anyone carrying a frame along a curve — polymer physics, cryo-EM flexibility analysis, MD trajectory analysis, DNA nanotechnology. Two things transfer directly. **On an open arc, a frame-defect angle is not a rotation about anything**, so part of your number is your arbitrary choice of where to point; report a multi-seed mean with its spread. And **signed writhe cancels on anything that is its own mirror image**, which is knowable from the definition in one line.

Methodologists. The general form is in the companion field guide (10.5281/zenodo.21811589). This paper is its worked case: what it costs to discover the same fact four times.

SECTION 0 — WHY THIS PAPER LOOKS LIKE THIS

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

We did not arrive here from structural biology. We arrived from a loop walker that had already been driven through a lot of substrates to find out where it broke.

Earthquakes first — rings of GPS stations around a fault, reading the circulation of a strain field and asking whether its failure to close said anything before the ground moved. Then storms, where a rotating system supplies exactly the coherent circulation the instrument reads best, and where the composite radar products mix scales in a way that has to be decomposed before there are honest anchor points at all. Then the sky: we asked the cosmos for geometric

meshes of stars and quasars, built the same way as a ring of GPS stations around a fault, and walked them looking for curvature that might point toward mass nobody can see and for velocities nobody had read that way. Asteroseismology sat on the same list.

Every one of those taught something, and about half taught it by not working. We started calling the exploratory runs **First Light** — a named, disposable project whose only job is to find out whether a new domain can carry the instrument at all, before anyone writes a claim down.

Knots looked like a natural next stop. There is real mathematics under them, and **mirror mathematics** — a knot equal to its own reflection is a cheap, clean test of whether a signed quantity can survive its own symmetry, and it is knowable before any data exists. DNA and RNA looked like a geometric story worth telling on its own terms: two strands winding around each other forever and never once crossing.

We assumed the interesting outcome would be a knot.

It wasn't. **DNA is essentially never knotted** — 250 curves, zero. RNA is essentially never knotted either, and the one published exception reproduced cleanly. The interesting outcome was a question I had not thought to ask: *how do you know whether a geometric measurement is capable of carrying a topological claim at all?*

I found out the expensive way. Four separate times I discovered that a continuous readout could not classify knots, and four separate times I treated it as a finding about holonomy. It was not. **It was one fact, found four ways**, and it was available in one minute at the start.

So the paper reorganised itself around the floor. The results are still here — a matched positive control, a trap that blocks anyone following the literature, an artefact large enough to invert a sign, a census nobody had run — but the spine is a method: measure the floor, rank your observables against it, check for a symmetry blind spot, and only then run the test. **I arrived at that rule by violating it repeatedly and in public**, and the violations are in the paper because they are the argument.

There is a version of this work that reports the positive control and stops. It would be shorter and it would be worse.

One image stayed with me the whole way. The two backbones of a duplex are **Escher's staircase in reverse**: they wind around each other forever and never once cross. That "never cross" is not decoration — it means the two strands form a *link*, with an integer linking number, and chasing that sentence found a factor-of-two error in my own ribbon code that had been silently breaking a validation I thought had passed.

METHODS, BRIEFLY

The backbone trace. A mode-switchable parser reads either C α atoms (protein) or phosphorus atoms (nucleic acid) as an ordered curve. **Gaps are declared, never bridged** — where residues are missing, the trace is broken rather than joined by a straight line. §8.2 is what that rule is for. Eighteen real nucleic-acid traces were extracted on first use.

Writhe. The Klenin-Langowski Gauss double integral (Klenin & Langowski 2000), vectorised in blocks: a 2,740-residue rRNA runs in about three seconds. Validated against three independent scalar implementations — agreement 2.93×10^{-14} on synthetic curves of known answer and 2.22×10^{-15} on real coordinates.

Knot typing. KMT reduction (Koniaris & Muthukumar 1991) followed by the Alexander determinant at $t = -1$, over an ensemble of stochastic closures in the lineage of Millett, Dobay & Stasiak (2005). *Purity* is the fraction of closures returning the modal answer. **The determinant gives an address, never a resident** — 6_1 and $3_1\#3_1$ share determinant 9, so any identification is hedged accordingly.

Duplex axis. DNA topology is axis topology, so duplex screening runs on the helical axis rather than the individual strands, built from base-pair centres with a continuity check and smoothed over one helical repeat. §9 is the record of getting both of those wrong first.

Frames. A Bishop (rotation-minimising) frame is propagated along the curve by quaternion composition. The defect angle is the residual rotation between the first and last normal. §3 is why that number needs a seed spread attached.

Independent referee. `knot_pull` (Jarmolińska et al.) — a different algorithm from a different group — was run on the knotted structures: **3/3 agreement**.

What we did not build. No knot detector. The mature tools listed in the references detect knots better than this instrument does, and where our results overlap theirs they agree.

1. THE FLOOR — AND WHY IT IS THE MOLECULE, NOT THE DETECTOR

Every model of an NMR ensemble is a legitimate fit to the same data, so the within-ensemble spread answers the question directly: *how much does this quantity move when nothing is wrong?*

1,172 molecules. 24,864 conformers. No topology call made.

channel	floor (median sd)	IQR
writhe	0.1886	0.1196–0.3054
ACN	0.3073	0.2012–0.4656
total curvature	0.9138	0.6562–1.2441
defect angle	35.81°	28.47–40.82

Replicated: an independent 389-molecule run gave 0.1913 and 34.82°.

What the floor is made of

This is the part that changes what a floor licenses, and it is worth stating plainly. **An NMR ensemble is not a noisy picture of one static structure.** It is a set of conformations the molecule actually occupies at room temperature. The spread we report is therefore **the molecule moving, not the detector failing.**

The ladder makes the mechanism concrete. The sugar-phosphate backbones are flexible stiles; the stacked base pairs are the rigid rungs, and the stack is what supplies the persistence length. Bend or twist the helix and the load distributes across both — the stiles absorb what the rungs cannot, and the resulting motion is what lets the molecule load and unload without tearing. **The floor is the price of that compliance.**

The consequence is operational. A detector floor can be beaten with a better detector. A substrate floor cannot. When your effect sits under a substrate floor, more averaging is not a strategy; the only moves left are a different observable, a paired difference, or a coarser unit in which the quantity is quantised.

Fenced: this is classical thermal fluctuation. DNA does have phonon modes, but they are not what an NMR ensemble measures, and nothing here is quantum or a resonance.

The consequence we did not test, and should have

If the floor is thermal motion, **the floor is a function of temperature**, and every number in this paper is quoted at one point on a curve nobody has drawn. We report the floor as though it were a constant of the substrate. It is a constant of the substrate *at the temperature the ensemble was collected*, and we did not check how far that matters.

The obvious question is whether cooling narrows the cloud enough to license a classification that room temperature forbids. **The obvious answer is not the right one.** Fast cooling vitrifies — it traps the existing distribution rather than narrowing it — so a snap-frozen ensemble should carry roughly the floor it had when it was warm. Narrowing would require slow cooling into a genuinely lower-energy distribution, which runs into cold denaturation from the other side. Whether there is a usable window between those two is an experiment, not an assumption, and it is answerable from variable-temperature depositions that already exist.

Two further questions follow and we state them without answers. Does the floor track temperature the same way for every channel, or do writhe, ACN and the defect angle have different thermal slopes — in which case the ranking of observables is itself temperature-dependent? And does a biological substrate hold its floor roughly flat across its own operating window, which would be a different and more interesting kind of regulation than mere thermal scaling?

And the floor belongs to a perturbation, not just a substrate

We stated that rule and then only exercised half of it. Every floor here was measured on whatever perturbation the deposition happened to contain: free thermal motion, unconstrained, at one temperature. **Nobody chose that condition. It is just what was lying around.**

The alternative is to specify the bench condition and measure the floor there — hold the molecule at a known temperature *and* a known stress, and report the spread around that held state rather than around no state at all. Two observations make this concrete, and they are two different instruments that get conflated easily.

Temperature is a gas problem, and it is already solved. A gentle stream of dry neutral gas across the sample is precisely how variable-temperature NMR probes work — uniform, tunable, no immersion, no thermal shock. So the temperature axis of this experiment needs no new hardware, and much of the data may already exist in variable-temperature depositions.

Applied bend is not a gas problem. A single nucleic acid in solution will not be pushed into a shape by flowing gas — at that scale viscous drag is negligible against Brownian motion, and the gas is moving the tube, not the molecule. Controlled stress on a single strand is a force-spectroscopy problem: magnetic or optical tweezers, or flow-stretching in a microfluidic channel, all of which already hold single molecules at a specified extension.

Put them together and the proposal becomes tractable: **map the floor as a surface over temperature and applied stress rather than quoting it at one point.** A molecule held at a known bend has fewer places to be, so the spread around that held state should be narrower than the free spread — and if the narrowing is large enough, a comparison that is impossible at rest may become possible on the bench. That would not beat the substrate floor. It would move it, by changing the substrate’s condition, which is the only honest way past a floor that belongs to the molecule.

Two predictions we are willing to be wrong about. **The defect angle should narrow most**, because it is the channel with the least length dependence and the most frame freedom, which makes it the natural candidate for a tuning-fork readout — the width of its distribution reporting how tightly the strand is held. And **the narrowing should degrade with unmodelled gaps**, because a gap is precisely a stretch of chain with no constraint on it, so it is where residual freedom will concentrate. If that holds, the count of gaps becomes a predictor of how much of a measured floor is real motion and how much is modelling latitude — which would connect this section directly to §8.2.

A competence surface is a frame-dependent measurement and we expect it to move. Everything mapped in this paper is mapped at one temperature, and that is a limitation of the map rather than a property of the instrument.

Competence surface

Chain length is the only real axis — Spearman **+0.326** for |writhe|, **+0.592** for ACN. Ensemble size and gap burden are flat (−0.068 to +0.052). The defect angle’s floor is nearly length-independent, uniformly noisy at ~36° whatever the chain, while the others grow 2–3× across the length range.

A trap we walked into: the floor grows with length, so normalising by *n* looks obvious. It makes every channel worse (ACN 0.7222 → 0.5940). The signal scales with length too.

Independent corroboration that this wall is real

A separate group met the same floor from the other side and abandoned the channel. In *Combined force-torque spectroscopy of proteins by means of multiscale molecular simulation* (arXiv 1910.02901) the authors report that “the writhe is rather sensitive to fluctuations in the positions of the atoms... we therefore deem it unsuitable for our analysis.” They hit it and stopped. **We measured it**, which is what turns an unusable channel into a usable one with a stated domain.

2. THREE SUBSTRATES, THREE FLOORS — AND OUR PREDICTION WAS BACKWARDS


A second substrate, measured: **220 RNA molecules, 3,214 conformers.**

channel	RNA floor	protein floor	ratio
writhe	0.0640	0.1886	0.34×
ACN	0.1258	0.3073	0.41×
defect angle	14.25°	35.81°	0.40×

RNA’s floor is 2.4–3× tighter than protein’s, reversing our registered expectation — we predicted RNA would be noisier because it is more flexible. The likely reason is corpus composition: RNA NMR entries are short, heavily base-paired stems and hairpins, while protein NMR entries carry flexible loops and disordered tails.

A crystal-replicate series (10GJ–10GO, six depositions) gives **~0.18** for DNA duplex axes.

Figure 2 — three substrates, three floors, and the prediction was backwards



We predicted RNA would be the noisiest, because it is more flexible. It is the tightest, by a factor of three. One floor does not serve three claims.

Figure 2. Three substrates, three floors, spanning a factor of three. Using one floor for all three claims is an error we made and corrected. The direction of the surprise matters as much as its size: **the substrate we assumed was sloppiest turned out to be the tightest**, which means the assumption is not safe in either direction and the measurement is cheap.

3. THE SECOND FLOOR, WHICH WE ONLY FOUND LATER

The frame defect is a genuine holonomy on a **closed** curve: seed-independent to machine precision, spread **0.000** across twelve random starting normals.

On an open arc it is not. The start and end tangents are different vectors, so the angle between the first and last normal is not a rotation angle about anything, and it depends on where in the normal plane you began.

structure	closed curve	open arc
2EFV	24.309 ± 0.000	97.325 ± 12.653
1UBI	102.733 ± 0.000	51.319 ± 3.100

Figure 4 — the same structure, unchanged, measured three times




Figure 3. The same structure, unchanged, measured three times. Close the loop and three arbitrary starting normals give the identical answer to machine precision — the holonomy belongs to the molecule. Leave the arc open and the same three give 78.635°, 97.325° and 110.220°. Part of the reported value is then the experimenter's arbitrary choice of where to point.

The instrument uses the open arc. So on top of the 35.81° conformational floor sits a **seed floor of roughly 3-13°** on a single, fixed, unchanging structure, from nothing but an arbitrary starting direction.

Consequence for every single-seed value in this program. Six angles that reproduce to $\leq 0.03^\circ$ are each one draw from a distribution $\sim 12^\circ$ wide. They agree to 0.03° **because the same arbitrary seed was used, not because the quantity is that precise.** Open-arc frame quantities must be reported as a multi-seed mean with its

spread.

4. WHY NO CONTINUOUS CHANNEL CLASSIFIES

Seven channels, each pre-registered, on a 7-protein corpus:

channel	AUC
total curvature	0.7500
ACN	0.6667
total winding	0.5833
helix-corrected winding	0.5833
writhe	0.5000
peak window winding	0.5000
defect angle	0.4167

None classify — and the sample size belongs in the same breath as the table. **At $n = 7$, only $AUC = 1.0000$ is significant** against the exact null ($p = 0.0286$). So 0.7500 is not a weak classifier; it is a number with no power behind it, and 0.4167 is not an anti-classifier for the same reason.

The floor explains the defect angle completely. Knotted-vs-unknotted median gap **3.3°** against a floor of **35.8°** — SNR **0.09**. Across 389 ensembles, **98.2%** have an internal spread larger than the entire 2EFV-vs-1UBI gap, and **63.5%** exceed the full knotted-unknotted range. **Two legitimate models of the same molecule differ by more than knotted and unknotted proteins differ from each other.**

Two structures explain the rest. 1TIM — an *unknot* — has the highest total winding in the corpus (25.018): a TIM barrel is enormous real geometric winding with zero topology. And 2CAV, a 5_2 knot at 346 residues, reads **$|\mathbf{Wr}| = 0.8152$** while unknotted 1TIM at 247 reads **21.3969** — a $26\times$ inversion, because signed writhe cancels on large folds. 2CAV's ACN is the highest in the corpus, so the winding is there; the signed sum destroys it.

Helix subtraction was tried and failed. A pure α -helix has a perfectly constant per-step frame rotation (89.63° , sd 0.00000) costing 0.13532 units of winding per residue — a clean $O(n)$ detector needing no secondary-structure assignment. Subtracting it moved the AUC by **zero** (7/12 both ways). Helical winding is not additive over secondary structure.

The rule this yields, and its limit


Ranking observables by effect-per-floor *before* testing predicted classification order perfectly on the corpus where the effect sizes were measured (Spearman = +1.0000). **That validation was partly circular**, and the rule then **failed its first genuinely out-of-sample test**. The floor generalises; the effect size does not.

What survives is weaker and still useful: **a channel whose effect is below its own floor cannot classify, and that is knowable in advance without labels.**

5. THE BLIND SPOT YOU CAN PREDICT WITHOUT RUNNING ANYTHING

Writhe uses *signed* Ω and is antisymmetric under mirroring, so an **amphichiral knot cancels to zero**. Measured: the figure-eight reads $|\mathbf{Wr}| = 0.08$ — invisible — while ACN, the same integral with $|\Omega|$, reads 9.09.

Figure 5 — the blind spot you can predict without running anything



Signed writhe cancels on a knot that is its own mirror image. One line of algebra, knowable before any data exists — and 4_1 is among the most common protein knots, so this is a partial screen.

Figure 4. On a 51-curve label-verified corpus, |writhe| splits sharply by chirality. Restricted to chiral knots at 100% recall it discards 87% of unknots — a **3.62×** work reduction. **4_1 is among the most common protein knots**, so this is a partial screen, and the amphichiral fraction of any target corpus sets what it is worth. That number must be quoted alongside the 3.62×, never separately.

Total curvature was tested as a mirror-symmetric alternative. Fáry-Milnor gives knotted $\implies$ total curvature $> 4\pi$, but the condition is *necessary, not sufficient*: unknotted curves in the corpus already reach 13.46 against a bound of 12.57.

6. THE DNA CENSUS

154 DNA-only entries plus 251 topoisomerase-DNA entries. **250 curves scored (246 after removing four duplicate depositions), 0 knotted, 0.00%.**

Screened on the duplex **helical axis**, not individual strands. DNA topology is axis topology; the backbones wind around it and read as long helices whatever the molecule's real topology. **That correction does not apply to RNA**, which is single-stranded and folds back on itself, so its backbone *is* the curve.

"DNA is not knotted in the PDB" is a result, and it is now a measured fact rather than an assumption.

7. THE POSITIVE CONTROL, AND AN ADDITION TO IT

RydC sRNA (PDB 4V2S, chain Q) is the only published knot in a natural RNA (Niemyska et al., *J Mol Biol* 436:168455, 2024).

Reproduced: determinant 3_1 , trefoil 3_1 , left-handed, on **4/4 seeds**, purity 0.828–0.894. Three writhe engines agree to **5.3×10^{-14}** on the same coordinates. `knot_pull` — a different algorithm from a different group — independently returns 3_1 .

Beyond the published result. The deposited structure has an unmodelled gap (residues 19–22, 12.51 Å against a 5.90 Å median) **inside the knotted core**. We rejection-sampled **24 physically plausible completions**: all P–P bond lengths at 5.90 Å, no steric clash inside 3.0 Å, verified genuinely distinct (per-axis spread ~ 2 Å, max pairwise separation 10.18 Å).

24/24 read trefoil. Every completion returned purity 0.863 to three decimals, which prompted a check rather than a report. **The reason is mechanistic:** KMT reduction deletes both inserted residues every time — 60 points reduce to the same 7 that the deposited 58 reduce to. **The missing loop is topologically inert.**

Handedness, resolved by comparison rather than by lookup. `knot_pull` returns identical signed DT codes for RydC and 2EFV. Signed DT codes encode chirality and 2EFV is documented left-handed, therefore RydC is left-handed. We did not read the answer anywhere; we placed the unknown beside a known and read the difference, which is the comparator move this instrument is actually for.

Knotted core (greedy trim, 101 closures): residues 11–53 of 57. The N-terminal fragment alone (7–19) reads unknot at purity 1.000. The entanglement is arm-against-body.

Prior art note: gap modelling with topological validation is not new — GapRepairer (Dąbrowski-Tumański et al., *Bioinformatics* 34:3300, 2018) does this, from the same group. The contribution here is the application to RydC and the KMT-inertness mechanism, not the method.

8. TWO TRAPS THE FIELD SHOULD KNOW ABOUT

8.1 A citation that cannot be followed


The literature cites “4V2S:G”. The deposited file contains **no chain G** in `auth_asym_id`.

A full scan of all 4,470 `atom_site` records resolves it: `auth_asym_id Q` $\equiv$ `label_asym_id G`. Entity 2 is `RYDC / 'SMALL REGULATORY RNA'`, one copy; Hfq is entity 1 with six copies occupying label chains A-F, so the RNA becomes the seventh, G. **All 57 phosphorus atoms in the entire file are in chain Q.**

The paper cites the *label* chain; every PDB-format tool reads the *auth* chain. Anyone following the citation literally finds nothing and may reasonably conclude the deposition is wrong.

8.2 Gaps that fabricate the signal

1S72 chain 0 (23S rRNA, 2,740 nt) reads **Wr = -91.57 as parsed** against **-8.67 on its longest gapless run**. Twelve bridged gaps inject roughly **83 units of fabricated writhe** — an order of magnitude larger than the signal. 4qln: -3.55 $\rightarrow$ -1.00 across two gaps.



Gaps are rare in the small single-domain proteins the instrument was built on and **routine in real RNA**. “Declare, never bridge” stops being hygiene and becomes load-bearing.

Prior art, and what it leaves us. Gap-filling as a hazard is known: the θ -curve survey (arXiv 1908.05919) models missing residues as straight intervals and then “*removed all structures with artificially long bonds or improbable gap filling.*” The field’s own summary is blunter — there are still no fully automatic protein-knot detectors “*due to the missing residues or chain breaks in the X-ray structures.*” The hazard is named and the standard response is to filter. **Our contribution is the magnitude, not the discovery.**

9. TWO SIGN-INVERTING BUGS IN OUR OWN CODE, BOTH CAUGHT THE SAME WAY

Building the duplex axis is where this work made its most instructive errors, and **both were caught by a synthetic construction of known answer that cost four lines.**

(a) Ordering. The first axis builder sorted base pairs without a continuity check. On a nucleosome the axis teleported across the particle. Simulated on a synthetic superhelix: 15% of centres out of order flips Wr from **-0.4563 to +0.4107** — matching the +0.354 mean the broken run returned.

(b) Wobble. Even correctly ordered, the C1'-C1' midpoint is *not* the helical axis. It orbits the true axis at the DNA's own repeat — roughly **14 right-handed turns against the superhelix's 1.65 left-handed ones**. At 2 Å wobble this flips the sign: truth **-0.4563**, raw midpoint **+0.4239**. Smoothing over one helical repeat recovers **-0.4191** at every amplitude tested.

Corrected result: 2CV5 (136 bp) **-0.510**, 10GJ (129 bp) **-0.461**, against a from-constants prediction of **-0.4563**. Fragments of 29-45 bp correctly read near zero — a fifth of a turn cannot carry writhe.

The twist claim, restated with a scale

The Drew-Dickerson dodecamer reads **+35.81°/bp**, and two independent definitions of twist agree to 0.09°. We originally reported this as landing “0.19° from a canonical ~36°/bp,” which reads as a precision result and is not one.

A twist floor measured across **48 duplex NMR ensembles (449 models)** gives **1.360°/bp** (IQR 0.781-2.038), and duplex twist across those ensembles genuinely ranges 29.96-38.64°/bp. **Our 0.19° agreement is 0.14× the floor** — comfortably inside the conformational noise.

The correct sentence is: *+35.81°/bp, consistent with the canonical value within a measured conformational floor of 1.36°/bp. It sounds weaker and it is stronger, because it carries its own scale.*

10. WHAT DID WORK — BASE-PAIRING DENSITY, AND A CHANNEL SWAP

Pre-registered: base-pairing density predicts a cleaner topological read, tested on 26 experimental RNA chains with partial Spearman controlling for chain length.

outcome	partial p	p
defect angle	-0.0842	0.686
writhe	+0.2448	0.235
ACN	+0.5135	0.0082

Holm-corrected across three outcomes: 0.0246, still below 0.05. The partial exceeds the raw (+0.4343), so chain length was **suppressing** the effect rather than creating it.

The provisional caveat is lifted. This was filed pending an RNA-specific floor, on the stated expectation that RNA would be noisier and the result would soften. Measured: the RNA ACN floor is **0.1258**, not the 0.3073 borrowed from protein. Against an ACN range of 1.207-54.960 across the corpus, that floor is **0.23% of the spread**. The expectation was wrong in the favourable direction and the result stands.

Remaining limitation, declared before the run. The mechanism is *not* the one originally hypothesised. This is helical content — more pairing, more A-form helix, more crossings without sign cancellation — **not topological threading**. So it is a new hypothesis that survived, rather than the old one vindicated.

11. A CORRECTION TO OUR OWN CODE, AND WHERE ITS RESIDUAL WENT

`cwf_ribbon.mutual_linking()` divides by 2π . The Gauss linking integral for two *distinct* curves is $(1/4\pi) \cdot \Sigma \Omega$ — each pair appears once, unlike self-writhe. **It returns exactly twice the linking number.**

Verified on synthetic duplexes of known turn count: as shipped, 2.0000 / 4.0000 / 6.0000 / 8.0000 / 12.0000 for 0.95 / 1.91 / 2.95 / 3.91 / 5.91 turns. Halved, all five are exact integers. `cwf_full.py` validates $Lk = Tw + Wr$ through that function and **reports False on all four test rows as shipped** — so the identity never closed and nothing was built on it.

Halving leaves a residual of ~0.47. **That residual is now resolved, and it is the Bishop frame's own closure holonomy — measured, not conjectured.**

route	Tw (turns)
Lk - Wr (topological)	0.4877
frame parallel-transport holonomy, 175.559°/360° (geometric)	0.4877
agreement	1.8×10^{-13}

The geometric route carries the final Bishop normal across the closing segment and measures its angle to the initial normal. **No linking integral appears in it.** The trefoil therefore decomposes as $Wr = -3.5123$ (continuous), $Tw = -0.4877$ (continuous), $Lk = -4.0000$ (pinned integer): neither continuous term is near a round number and their

sum is exactly an integer. **That is Călugăreanu-White-Fuller demonstrated rather than cited.**

And the half-twist ladder fails pointwise for a reason

k	0	1	2	3	4	5	6	7	8
Lk measured	−4	−3	−3	−2	−2	−1	−1	0	0
$Lk_0 + k/2$	−4.0	−3.5	−3.0	−2.5	−2.0	−1.5	−1.0	−0.5	0.0

Every measured value is an integer to better than 10^{-4} , and **it cannot be otherwise** — the linking number of two disjoint closed curves is an integer by construction, so a half-integer step is a category error rather than a small one. Lk absorbs the added twist in whole units, stepping at odd k and holding at even k, and the mean slope is exactly **+0.5000 per half-twist**. The classical prediction is right on average and cannot be right pointwise. **This is the integer/continuous split on a synthetic curve where the answer is known independently.**

Caveat, measured not assumed: the twist ramp runs to $k \cdot \pi \cdot (n-1)/n$ rather than $k \cdot \pi$, so the ribbon does not exactly close. Closure defect is ≈ 1.96 – 1.99 at even k and ≈ 0.22 – 0.35 at odd k. The integer values are stable regardless; the ramp should be corrected to $i/(n-1)$ before anything further is built on it.

12. LIMITATIONS

Convenience and competition-set sampling throughout. The geometric Watson-Crick detector under-counts (0.514 on 1S72 23S rRNA against a textbook ~ 0.60 , quoted not verified). **n is small nearly everywhere**; knotted RNA barely exists, so any RNA ROC has two real positives. Closure is not free — writhe differences move 6.5×10^{-2} across four closure schemes, comparable to the signal.

1TIM: resolved, and it was never a discrepancy. Six defect angles reproduce to $\leq 0.03^\circ$; 1TIM read 94.99 against a published 108.6. An exhaustive scan of all **35,530** sub-traces ruled out truncation: 28 land within 0.05° of 108.6 against **~ 20 expected by pure counting** (Poisson $p = 0.052$), and **zero** have both endpoints at a natural boundary. A “plausible boundary” filter we applied fired on 5 of 28 against a chance expectation of 5.3 — **it discriminated nothing and was withdrawn.**

The source records resolve it: **the table mixes two statistics.** Five rows are single fixed-seed measurements. **1TIM’s 108.6 is a mean over six random frame seeds, reported as 108.6 ± 38.3 , range [45.8, 150.3].** Given the open-arc seed floor above, that spread is exactly what the quantity does. Our 94.99 is a single draw well inside one standard deviation of the published mean. **Two different statistics of the same quantity, not a disagreement.**

13. WHAT TRANSFERS

1. **Measure the floor first**, on replicates of the right substrate *and* perturbation. One unattended run.
2. **A channel whose effect is below its own floor cannot classify**. Knowable without labels, before any test.
3. **Check for a symmetry blind spot from the definition**. [writhe] is provably blind to amphichiral knots — one line of algebra, and the difference between AUC 0.98 and 0.51.
4. **Never ship a derived geometric quantity without a synthetic construction of known answer beside it**. A synthetic nucleosome caught two sign-inverting bugs in one run and cost four lines.
5. **Floors are substrate- and perturbation-specific**. Protein NMR 0.1886, RNA NMR 0.0640, DNA crystal replicates ~0.18 — and our stated expectation about which would be noisiest was backwards. One floor does not serve three claims.
6. **An unscaled agreement is not a result**. “0.19° from canonical” looked like precision until a floor of 1.36° existed. Quote every agreement against its own noise or do not quote it.
7. **A quantity can be exactly reproducible and still be arbitrary**. Six values reproduced to 0.03° because the same arbitrary seed was used twice, not because the measurement was precise.
8. **Count your candidates before you celebrate a match**. Searching 35,530 sub-traces for one that hits a target will find ~20 by arithmetic alone. Any post-hoc filter meant to rescue such a search must be checked against its own chance rate — ours fired at exactly the rate chance predicted and we nearly reported it as evidence.

HOW TO BREAK THIS

Every floor in this paper is one unattended run on public data. Take any PDB NMR ensemble, compute a geometric descriptor across its models, and read the spread. If our numbers do not reproduce, that is this paper’s problem and we want to know. Breaks get logged in the open with attribution.

The interesting attacks are not arithmetic. Is the within-ensemble spread the right stand-in for thermal motion, or is it partly a refinement artefact? Does the RNA-tighter-than-protein result survive a corpus matched for chain length and disorder? Does a paired design beat the substrate floor where an absolute one cannot? **We think the answers are yes, probably, and yes — and none of those is demonstrated here.**

REFERENCES

House standard: **[VERIFY]** marks an entry located but whose year, volume or pages have not been confirmed against the original in this session. Nothing is cited that was not used.

Instrument lineage - Călugăreanu, G. (1961). Sur les classes d’isotopie des nœuds tridimensionnels et leurs invariants. *Czechoslovak Math. J.* 11, 588–625. **[VERIFY]** - White, J. H. (1969). Self-linking and the Gauss integral in higher dimensions. *Amer. J. Math.* 91, 693–728. **[VERIFY]** - Fuller, F. B. (1971). The writhing number of a space curve. *PNAS* 68, 815–819. **[VERIFY]** - Klenin, K. & Langowski, J. (2000). Computation of writhe in modeling of supercoiled DNA. *Biopolymers* 54, 307–317. - Koniaris, K. & Muthukumar, M. (1991). Self-entanglement in ring polymers. *J. Chem. Phys.* 95, 2873. - Millett, K., Dobay, A. & Stasiak, A. (2005). Linear random knots and their scaling behaviour. *Macromolecules* 38, 601–606. - Bishop, R. L. (1975). There is more than one way to frame a curve. *Amer. Math. Monthly* 82, 246–251. **[VERIFY]**

Bounds and theorems used - Fáry, I. (1949). Sur la courbure totale d’une courbe gauche faisant un nœud. *Bull. Soc. Math. France* 77, 128–138. **[VERIFY]** - Milnor, J. W. (1950). On the total curvature of knots. *Ann. Math.* 52, 248–257. **[VERIFY]**

The prior tools — better at naming a knot than this instrument is - Niemyska, W., Rubach, P., Greń, B. A., et al. (2024). AlphaKnot 2.0. *Nucleic Acids Res.* 52(W1), W187. **[VERIFY]** - Dabrowski-Tumanski, P., Rubach, P., Goundaroulis, D., et al. (2019). KnotProt 2.0. *Nucleic Acids Res.* 47(D1), D367. - Dabrowski-Tumanski, P., Rubach, P., Niemyska, W., Gren, B. A. & Sulkowska, J. I. (2021). Topoly. *Brief. Bioinform.* 22(3), bbaa196. **[VERIFY]** - Dorier, J., Goundaroulis, D., Benedetti, F. & Stasiak, A. (2018). Knoto-ID. *Bioinformatics* 34, 3402. - Jarmolińska, A. I., Gambin, A. & Sulkowska, J. I. (2020). Knot pull — python package for biopolymer smoothing and knot detection. *Bioinformatics* 36(3), 953–955. **[VERIFY]** - Dąbrowski-Tumański, P., et al. (2018). GapRepairer. *Bioinformatics* 34, 3300.

Substrate and corroboration - Niemyska, W., et al. (2024). The first knotted RNA. *J. Mol. Biol.* 436, 168455. **[VERIFY]** — the RydC positive control of §7. - *θ-curves in proteins*, arXiv 1908.05919 — gap-filling named as a hazard and handled by filtering. **[VERIFY]** author attribution. - *Combined force-torque spectroscopy of proteins by*

means of multiscale molecular simulation, arXiv 1910.02901 — an independent group meeting the same fluctuation floor and abandoning the channel. **[VERIFY]** author attribution.

Companion deposits in this program - Schulz, M. (2026). *Loop Walking: A Field Guide to Knowing Whether You Found Anything*. Zenodo 10.5281/zenodo.21811589. - Schulz, M. (2026). *The Tower Is Coordinate-Free*. Zenodo 10.5281/zenodo.21798311. - Schulz, M. (2026). *Reading the Sky Through Loops*. Zenodo 10.5281/zenodo.21331652. - 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.

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

Related deposits (cite all): - *Loop Walking: A Field Guide* — 10.5281/zenodo.21811589 - *The Tower Is Coordinate-Free* — 10.5281/zenodo.21798311 - *Taking a Loop Walk Through the Integer Forest* — 10.5281/zenodo.21449570 - *Walking the Loop: A Quick Guide to Holonomy* — 10.5281/zenodo.21448987 - *Reading the Sky Through Loops* — 10.5281/zenodo.21331652

Prior-art status. A credentialed zbMATH / MathSciNet pass has not been run on this paper. The search covered arXiv, PMC, publisher pages and OSTI. Novelty language is restricted to “not located,” never “new.” Any located hit will be logged in the open with attribution.

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