Gauge Freedom Field Notes

Shared science, insights, and news from the Gauge Freedom Lab.

Seed-Window Algebraic Instability in Dementia-Associated miRNAs: A Pilot Study

Field Notes / Preprint (Not Peer-Reviewed)

Marcelo Maciel Amaral

Gauge Freedom, Inc., Los Angeles, CA, USA

marcelo@gaugefreedom.com

DOI: 10.65323/gf-lab.2026.001 • **Posted:** April 2026 • **License:** CC BY 4.0

Abstract

We introduce a presentation-first algebraic profiler for short RNA windows and use it to obtain a worked pilot example across a small set of human miRNAs. Each overlapping 6-mer, 7-mer, or 8-mer window in the canonical seed neighborhood is represented as a finitely presented group, and its $SL(2,\mathbb{C})$ character variety is probed through a sliced Gröbner-basis computation that resolves a free-like class, a singular versus non-singular classification, and (where resolved) the singular-surface family. Per-window outcomes are assigned to one of three tiers — fully resolved, projection-skipped (basis computed, low-dimensional elimination bypassed because of complexity), or timeout-preliminary — so that compute-limited windows do not silently distort downstream summaries. Adjacent-window comparisons yield a provisional instability score that decomposes into free-class transitions, singular/non-singular flips, and singularity-type changes. We illustrate the framework on nine human miRNAs: seven dementia-associated cases (`hsa-miR-124-3p`, `hsa-miR-29a-3p`, `hsa-miR-29c-3p`, `hsa-miR-92a-3p`, `hsa-miR-486-5p`, `hsa-let-7i-5p`, `hsa-miR-15a-5p`) and two preliminary controls (`hsa-let-7a-5p`, `hsa-miR-16-5p`). In this pilot, the resolved-singular signal is concentrated in a single anchor case, `hsa-miR-124-3p`, which retains the previously published `AAGGCA`/`AAGGCAC` contrast as a regression fixture and embeds it within an independently computed local neighborhood. The miR-29 family shows matched free-class variation across both family members but no resolved singular/non-singular flips, while `hsa-let-7a-5p`, `hsa-let-7i-5p`, and `hsa-miR-92a-3p` are flat across the recorded dimensions. We do not interpret these scores as a disease-versus-control separation: the cohort is too small to support such an inference, and disease-labeled and control-labeled cases score identically in several comparisons. The contribution is therefore methodological — a reproducible measurement framework with auditable failure modes — and the dementia cohort is illustrative rather than evidential.

Keywords— microRNA, presentation-first profiler, finitely presented groups, $SL(2,\mathbb{C})$ character variety,

Gröbner basis, sequence instability, methods pilot

1 Introduction

MicroRNAs (miRNAs) are short noncoding RNAs whose mature sequences regulate gene expression in large part through short seed regions that shape target recognition and regulatory specificity [Kozomara et al., 2019]. Dementia-associated miRNAs have therefore attracted attention as potential regulatory mediators and biomarker candidates in neurodegeneration, including Alzheimer's disease [Kou et al., 2020]. The question addressed here is not whether one isolated motif exhibits one isolated algebraic signature, but whether disease-linked miRNAs tend to lie near *local algebraic transition points* when one scans overlapping windows across the seed neighborhood.

The algebraic framework used here builds on earlier work in which short DNA/RNA motifs were represented by finitely presented groups and compared against free-group baselines through subgroup-count signatures [Planat et al., 2022a]. That program was subsequently extended to $SL(2,\mathbb{C})$ character varieties and Gröbner-basis analysis, where isolated singular algebraic surfaces emerged as an additional signal beyond simple free-group proximity [Planat et al., 2022b, 2023]. On the mathematical side, the present implementation also draws on presentation-to-character-variety constructions for finitely presented groups and on trace-coordinate descriptions of low-rank $SL(2,\mathbb{C})$ moduli spaces [Ashley et al., 2018, Goldman, 2009, Cantat, 2009].

The present study returns to the original seed-window idea and asks whether overlapping 6-mer, 7-mer, and 8-mer windows around the canonical seed neighborhood show local changes in free-like class, singular versus non-singular status, or singularity type. Rather than assigning a single global algebraic label to an entire mature miRNA, we ask whether neighboring seed windows can move between algebraic regimes under one-base shifts or extensions. The aim is methodological and exploratory: this is a pilot computational-genomics study of seed-window algebraic instability, not a diagnostic classifier and not a mechanistic model of dementia biology. The overall workflow is summarized in Figure 1. Each mature miRNA is decomposed into overlapping seed-neighborhood windows, each window is assigned a presentation-first algebraic profile, and adjacent windows are then compared to derive a local instability score. Projection-skipped and timeout-preliminary outcomes are retained explicitly rather than being treated as negative findings.

2 Methods

2.1 Study design

We analyzed the pilot cohort listed in Table 1. Mature miRNA sequences were curated from miRBase [Kozomara et al., 2019]. The dementia-focused pilot cohort consisted of `hsa-miR-124-3p`, `hsa-miR-29a-3p`, `hsa-miR-29c-3p`, `hsa-miR-92a-3p`, `hsa-miR-486-5p`, `hsa-let-7i-5p`, and `hsa-miR-15a-5p`. Two preliminary comparison controls, `hsa-let-7a-5p` and `hsa-miR-16-5p`, were included to provide an initial low-singularity reference context. The cohort was selected as a concept-driven pilot and was not intended as a statistically powered case-control panel.

2.2 Seed-neighborhood window definition

The analysis was restricted to short windows in the canonical seed neighborhood rather than full-length mature-sequence scanning. For each mature miRNA, we evaluated overlapping 6-mer, 7-mer, and 8-mer windows centered on the canonical seed region (positions 2–8 in the mature sequence), with a local start-shift

allowance of one nucleotide on either side. Operationally, this yields three windows of each length and therefore nine seed-neighborhood windows per mature miRNA in the current implementation.

2.3 Algebraic profiling pipeline

For each window, we constructed a finitely presented group from the nucleotide word and evaluated its algebraic profile with the current presentation-first profiler. The resulting pipeline combines the free-group baseline criterion developed in the earlier DNA/RNA group-theory work [Planat et al., 2022a] with the character-variety and singular-surface perspective developed in the subsequent algebraic-surface studies [Planat et al., 2022b, 2023]. Four windows from the earlier framework are retained as published regression fixtures used for consistency checks rather than recomputed by the current run: `UCCUAC`, `UCCUACA`, `AAGGCA`, and `AAGGCAC` [Planat et al., 2023]. When the current pipeline encounters any of these four motifs, the previously published classification (free-like class, singular/non-singular status, and singularity type, where applicable) is returned directly; the surrounding seed-neighborhood windows around these fixtures are computed independently by the pipeline. Rows produced from fixtures are tagged `analysis_status = published-regression` in the per-window data and can be reproduced or audited separately via the regression-test entrypoint of the profiler (see Section 2.6).

For each window, we recorded the window sequence, the implied free-group baseline F_r, the observed free-like class, singular versus non-singular status when resolved, singularity type when available, and whether the window differed algebraically from adjacent seed-neighborhood windows.

2.4 Resolution tiers

Symbolic computation was routed through a `Singular` subprocess when possible, with a `SymPy` fallback. Because low-dimensional elimination can become substantially more expensive than basis construction itself, each window was assigned to one of three outcome tiers:

1. **Resolved (`benchmark-sliced`).** Full Gröbner-basis computation and the corresponding variable-elimination projection both completed successfully.
2. **Projection-skipped (`projection-skipped-due-to-complexity`).** The trace-algebra ideal and sliced basis were computed, but elimination was bypassed because the elimination-variable count exceeded the complexity cutoff. In these cases, the free-like class remains informative, whereas singular versus non-singular classification is left unknown.
3. **Timeout-preliminary.** The per-window symbolic ceiling of 200 seconds was reached before basis computation completed.

Projection-skipped windows were therefore treated as known in the free-class dimension and unknown in the singularity dimension. Timeout-preliminary windows were treated as unresolved and were excluded from transition counting in any dimension that remained unknown.

2.5 Instability summary

For each miRNA, we computed a provisional instability score as the sum of three transition counts across adjacent seed-neighborhood windows: *free-class transitions* (any change in the recorded free-like class between two adjacent windows, including same-rank shifts such as $F_3 \to F_2$ and shifts between free-like and non-free-like classes such as $F_2 \to \pi_2$), *singular/non-singular flips*, and *singularity-type changes*. A transition

contributed to a given count only when both neighboring windows had known state in the relevant dimension; projection-skipped and timeout-preliminary windows were treated as unknown in any dimension they did not resolve, and missing-state sentinels (including any post-roundtrip `NaN` values that may appear in tabular outputs) were canonicalized before comparison so that an unknown state never counted as a transition. For the singularity-type count, "no singularity" was treated as a single canonical value, so that two adjacent non-singular windows did not contribute a type change. The resulting score is exploratory and descriptive rather than calibrated, and it is intended to summarize local algebraic transition structure rather than to function as a diagnostic or predictive classifier.

2.6 Execution parameters and reproducibility

The pilot data underlying the tables and figures in this paper were produced in two passes. The first pass evaluated the full seed-neighborhood cohort with the default per-window symbolic ceiling of the profiler. The second pass re-evaluated the windows that did not resolve under the first pass, using an extended 200-second per-window symbolic ceiling routed through the `Singular` backend with a `SymPy` fallback and the projection cutoff `MAX_ELIMINATION_VARIABLES = 6`. The exact second-pass invocation, captured in `run_prompt16_extended.sh`, was:

```
python3 -u Sliding_Window_Algebraic_Profiler.py \
    --window-list-csv data/timeout_windows_prompt16.csv \
    --output-prefix data/pilot_run_prompt16_extended \
    --window-timeout-seconds 200
```

The merged per-window outputs (`data/pilot_run_prompt16.csv` and the corresponding `.json`), the cohort-level summary (`data/pilot_run_prompt16_summary.csv` and `.md`), and the manuscript figures and tables are all produced from this combined data set; the regeneration script for figures and Markdown tables is `generate_manuscript_figures.py`. The repository `README.md` lists the pinned Python dependencies and the full reproduction command set for both passes.

3 Results

3.1 Cohort-level overview

The pilot analysis yielded three distinct outcome classes across the seed-neighborhood windows: fully resolved windows, projection-skipped windows, and true timeout-preliminary windows. Fully resolved windows completed both the sliced Gröbner-basis computation and the corresponding low-dimensional projection. Projection-skipped windows retained a computable higher-dimensional trace-algebra object and a resolved free-like class, but low-dimensional elimination was bypassed because of symbolic complexity. Timeout-preliminary windows remained unresolved within the current per-window compute budget. This three-tier outcome structure is central to the interpretation of the pilot because it distinguishes partially informative symbolic outcomes from genuinely unresolved cases.

Across the cohort, the pilot recovered one resolved-singular anchor case, a family-level free-class variation signal, and a flat low-variation comparison group. `hsa-miR-124-3p` is the clearest anchor and the only miRNA in the cohort with a resolved singular window in the present pass; it is also the highest-scoring case, retaining the previously published singular/non-singular seed-extension contrast. The miR-29 family (`hsa-miR-29a-3p` and `hsa-miR-29c-3p`) shows matched seed-neighborhood behavior characterized by free-class variation across neighboring windows, with no resolved singular/non-singular flips and no singularity-type

Table 1: Pilot cohort used in the seed-window study. Mature sequences were curated from miRBase and grouped as dementia-focused pilot cases or preliminary controls.

miRNA	Mature sequence	Cohort - Role	Source
hsa-miR-124-3p	UAAGGCACGCGGUGAAUGCCAA	dementia - pilot-anchor	miRBase MIMAT0000422
hsa-miR-29a-3p	UAGCACCAUCUGAAAUCGGUUA	dementia - pilot-cohort	miRBase MIMAT0000086
hsa-miR-29c-3p	UAGCACCAUUUGAAAUCGGUUA	dementia - pilot-cohort	miRBase MIMAT0000681
hsa-miR-92a-3p	UAUUGCACUUGUCCCGGCCUGU	dementia - pilot-cohort	miRBase MIMAT0000092
hsa-miR-486-5p	UCCUGUACUGAGCUGCCCCGAG	dementia - pilot-cohort	miRBase MIMAT0002177
hsa-let-7i-5p	UGAGGUAGUAGUUUGUGCUGUU	dementia - pilot-cohort	miRBase MIMAT0000415
hsa-miR-15a-5p	UAGCAGCACAUAAUGGUUUGUG	dementia - pilot-cohort	miRBase MIMAT0000068
hsa-miR-16-5p	UAGCAGCACGUAAAUAUUGGCG	control - preliminary	miRBase MIMAT0000069
hsa-let-7a-5p	UGAGGUAGUAGGUUGUAUAGUU	control - preliminary	miRBase MIMAT0000062

changes; the singularity dimension is partially unresolved because several 8-mer windows are projection-skipped. hsa-let-7a-5p, hsa-let-7i-5p, and hsa-miR-92a-3p contain no resolved singular windows and no free-class variation, giving a flat low-variation comparison group. Table 3 summarizes the cohort-level instability profiles, and Figure 3 decomposes the provisional scores into their component transition counts together with the unresolved-window burden.

3.2 hsa-miR-124-3p as the anchor case

hsa-miR-124-3p is the strongest and most interpretable case in the pilot, and the only miRNA in the cohort with a resolved singular window in the present pass. As shown in Table 2 and Figure 2, the 6-mer AAGGCA carries the published singular classification, whereas the 7-mer extension AAGGCAC carries the corresponding non-singular classification. These two windows are retained from the earlier framework as published regression fixtures (Section 2.3); the surrounding windows extend the known pair into a broader local transition profile across the seed neighborhood rather than leaving it as an isolated two-window observation.

This result is important for two reasons. First, the fixture-anchored contrast is embedded in an independently computed seed-window neighborhood, so the local profile around AAGGCA/AAGGCAC is the product of the current pipeline rather than a single legacy observation. Second, the key signal is genuinely local: a one-nucleotide extension coincides with a change in algebraic class, and that local change is embedded in a broader neighborhood of nearby windows with differing algebraic behavior. In practical terms, hsa-miR-124-3p is the clearest evidence in the present cohort that a dementia-linked miRNA can sit near a seed-window algebraic transition point rather than within a single fixed class.

3.3 The miR-29 family as a free-class variation signal

hsa-miR-29a-3p and hsa-miR-29c-3p show matched seed-neighborhood behavior in the free-class dimension. Both family members register four free-class transitions across adjacent windows, with no resolved singular/non-singular flips and no singularity-type changes (Table 3). The shared structure across the two family members is the main reason this pair is reported as a family-level observation, rather than the absolute size of the score.

The present evidence should be described with precision. The miR-29 signal in the current pass is entirely a free-class signal: within the resolved windows, every miR-29 window is non-singular, and the alternations between F_2 and F_3 across adjacent seed-neighborhood windows are what the score reflects. Several 8-mer windows in both family members reached projection-skipped status, meaning that the higher-dimensional

trace-algebra object was still computed but the low-dimensional elimination needed for singularity classification was bypassed. The singularity dimension is therefore unknown for those windows, not negative. The miR-29 family is consequently best summarized as showing matched free-class variation across neighboring windows, with the singularity dimension partially unresolved rather than confirmed flat.

3.4 Single-transition cases

`hsa-miR-486-5p`, `hsa-miR-15a-5p`, and the preliminary control `hsa-miR-16-5p` each contain exactly one free-class transition across their seed neighborhoods, with no resolved singular/non-singular flips and no singularity-type changes. The shared characterization is therefore a single localized free-class change rather than a transition-rich profile.

`hsa-miR-15a-5p` and `hsa-miR-16-5p` additionally remain the most computationally difficult pair in the current pass. Their neighborhoods contain the residual true timeout-preliminary windows and additional projection-skipped windows, so part of their singularity dimension remains genuinely unresolved. `hsa-miR-486-5p` likewise contains one projection-skipped 8-mer. These three miRNAs are best described as partially characterized: a single observed free-class transition each, with limited singularity-side coverage.

3.5 Low-variation comparison group: `hsa-let-7a-5p`, `hsa-let-7i-5p`, and `hsa-miR-92a-3p`

`hsa-let-7a-5p`, `hsa-let-7i-5p`, and `hsa-miR-92a-3p` form the flat low-variation comparison group in the current pilot. None of these three sequences exhibited a resolved singular window in the present pass, and none registered any free-class transitions across adjacent seed-neighborhood windows. The two let-7 family members are fully resolved across all nine windows and remain in a single free-like class (F_2) throughout; `hsa-miR-92a-3p` is resolved or projection-skipped depending on window, and the resolved windows are all F_3 non-singular. Their provisional scores are therefore zero by construction in this pass.

This group serves as a natural comparison context for the anchor and free-class-variation cases above. The present framework records multiple dimensions of local algebraic change, and the comparison group shows that some seed neighborhoods, both dementia-linked (`miR-92a-3p`) and preliminary control (`let-7a-5p`), can be flat across the recorded dimensions in the present pass. The distinction between `hsa-miR-124-3p`, the miR-29 free-class-variation pair, and this flat comparison group is the main qualitative differentiation that the pilot supports at this scale.

3.6 Interpretation of projection-skipped windows

A notable feature of the current pilot is the appearance of projection-skipped windows as a recurrent middle category between fully resolved and true timeout-preliminary outcomes. These windows are not failures in the same sense as full timeouts. In each such case, the trace-algebra ideal and sliced basis were computed, and the free-like class remains informative. What is missing is only the final low-dimensional elimination needed for singular versus non-singular classification in the current representation.

This distinction has two consequences for interpretation. First, the free-class dimension is the better-resolved dimension across the cohort: every miR-29 free-class transition reported above sits against at least one projection-skipped 8-mer in the singularity dimension, so the free-class signal is not contaminated by missing singularity information. Second, any singularity-based component of the score for a miRNA that contains projection-skipped windows is a lower bound rather than a complete count. The pilot therefore captures both biological variation and symbolic burden, and the latter should be understood as part of the present computational boundary rather than as negative biological evidence.

3.7 Summary of the pilot signal

Taken together, the results support a shift in emphasis from single-motif classification to local transition structure, while showing that at this cohort size and compute budget the signal is concentrated rather than broad. The strongest case is `hsa-miR-124-3p`, which is the only miRNA in the cohort with a resolved singular window and the highest-scoring miRNA overall; its anchor pair is retained from the earlier framework as a published regression fixture and is contextualized within a broader local neighborhood that the present pipeline computes independently. The miR-29 family shows matched free-class variation across neighboring windows but no resolved singularity-based transitions, and the let-7 pair together with `hsa-miR-92a-3p` form a flat low-variation comparison group. The qualitative differentiation among these three groups, rather than the absolute size of any single score, is the central finding of the pilot.

4 Discussion

The present pilot supports three main observations. First, the seed-window algebraic instability framework is operationally workable and produces a reproducible cohort-level summary, distinguishing the resolved-singular anchor case (`hsa-miR-124-3p`) from the miR-29 family, which varies in the free-class dimension only, and from a flat low-variation group (let-7a, let-7i, and miR-92a) with no resolved singular windows and no free-class transitions. Second, in this cohort the resolved singularity signal is concentrated in a single anchor case rather than recurring across multiple sequences, and the family-level signal in miR-29 is best read as matched free-class variation rather than as a transition-rich profile. Whether this concentration reflects an underlying biological feature, the limited cohort size, or the conservative effect of projection-skipped windows on the singularity dimension cannot be decided at this scale. Third, the three-tier outcome system is scientifically useful in its own right: projection-skipped windows preserve information about the higher-dimensional trace-algebra object and free-like class even when low-dimensional elimination is impractical, whereas timeout-preliminary windows remain genuinely unresolved. Treating these categories separately produces a more faithful description of algebraic complexity across the cohort.

These observations should be interpreted as exploratory rather than diagnostic. The cohort is small, the instability score is descriptive rather than calibrated, only one resolved singular window is observed across the cohort in the present pass, and the study does not attempt to connect the algebraic signal to RNA structure, target-network biology, or clinical prediction. Even so, the present results show that the seed-window viewpoint is operationally workable as an extension of the earlier algebraic-morphology framework, and that local transition structure within the seed neighborhood is a tractable quantity to record alongside any single isolated motif classification.

5 Limitations

The cohort is small and concept-driven rather than statistically powered. Three windows remain genuine timeouts, and twelve windows remain projection-skipped, so singularity-based instability is still partially underestimated. The provisional instability score uses equal weights and should be read as an exploratory summary rather than a clinical or biological classifier. The present work also does not attempt a mechanistic bridge to RNA folding, target-network biology, or matched large-scale case-control expression studies.

6 Conclusion

This pilot study shows that seed-window algebraic instability can be measured in dementia-associated miRNAs within a reproducible presentation-first workflow. `hsa-miR-124-3p` is the clearest case and the only miRNA in the cohort with a resolved singular window; its singular/non-singular seed-extension contrast is retained from the earlier framework as a published regression fixture and is embedded in a broader local neighborhood that the present pipeline computes independently. Beyond this anchor case, the pilot shows matched free-class variation in the miR-29 family but no resolved singularity-based transitions, while `hsa-let-7a-5p`, `hsa-let-7i-5p`, and `hsa-miR-92a-3p` form a flat low-variation comparison group across the recorded dimensions.

The main contribution of the present study is therefore methodological and empirical. It establishes a workable framework for evaluating overlapping seed-neighborhood windows, distinguishes resolved from partially resolved symbolic outcomes, and shows that neighboring seed windows can differ sharply in algebraic class in at least one cohort member. At the same time, the study remains exploratory. The cohort is small, the provisional instability score is descriptive rather than predictive, only one resolved singular window is observed across the cohort in the present pass, and some windows remain projection-skipped or unresolved. Larger follow-up studies with expanded controls and stronger symbolic backends will be needed to determine how broadly seed-window algebraic instability generalizes across dementia-linked miRNAs and related disease contexts.

Acknowledgments

An early internal version of the AI research tool Quastra was used for code assistance, formatting, and language refinement. All scientific decisions, analyses, and interpretations were performed and verified by the author.

References

Ana Kozomara, Maria Birgaoanu, and Sam Griffiths-Jones. miRBase: from microRNA sequences to function. *Nucleic Acids Research*, 47(D1):D155–D162, 2019. doi: 10.1093/nar/gky1141.

Xianjuan Kou, Dandan Chen, and Ning Chen. The regulation of microRNAs in Alzheimer's disease. *Frontiers in Neurology*, 11:288, 2020. doi: 10.3389/fneur.2020.00288.

Michel Planat, Marcelo M. Amaral, Fang Fang, David Chester, Raymond Aschheim, and Klee Irwin. Group theory of syntactical freedom in DNA transcription and genome decoding. *Current Issues in Molecular Biology*, 44(4):1417–1433, 2022a. doi: 10.3390/cimb44040095.

Michel Planat, Marcelo M. Amaral, Fang Fang, David Chester, Raymond Aschheim, and Klee Irwin. DNA sequence and structure under the prism of group theory and algebraic surfaces. *International Journal of Molecular Sciences*, 23(21):13290, 2022b. doi: 10.3390/ijms232113290.

Michel Planat, Marcelo M. Amaral, and Klee Irwin. Algebraic morphology of DNA–RNA transcription and regulation. *Symmetry*, 15(3):770, 2023. doi: 10.3390/sym15030770.

Caleb Ashley, Jean-Philippe Burelle, and Sean Lawton. Rank 1 character varieties of finitely presented groups. *Geometriae Dedicata*, 192(1):1–19, 2018. doi: 10.1007/s10711-017-0281-6.

William M. Goldman. Trace coordinates on Fricke spaces of some simple hyperbolic surfaces. In Athanase Papadopoulos, editor, *Handbook of Teichmüller Theory, Volume II*, volume 13 of *IRMA Lectures in Mathematics and Theoretical Physics*, pages 611–684. European Mathematical Society, Zürich, 2009.

Serge Cantat. Bers and Hénon, Painlevé and Schrödinger. *Duke Mathematical Journal*, 149(3):411–460, 2009. doi: 10.1215/00127094-2009-042.

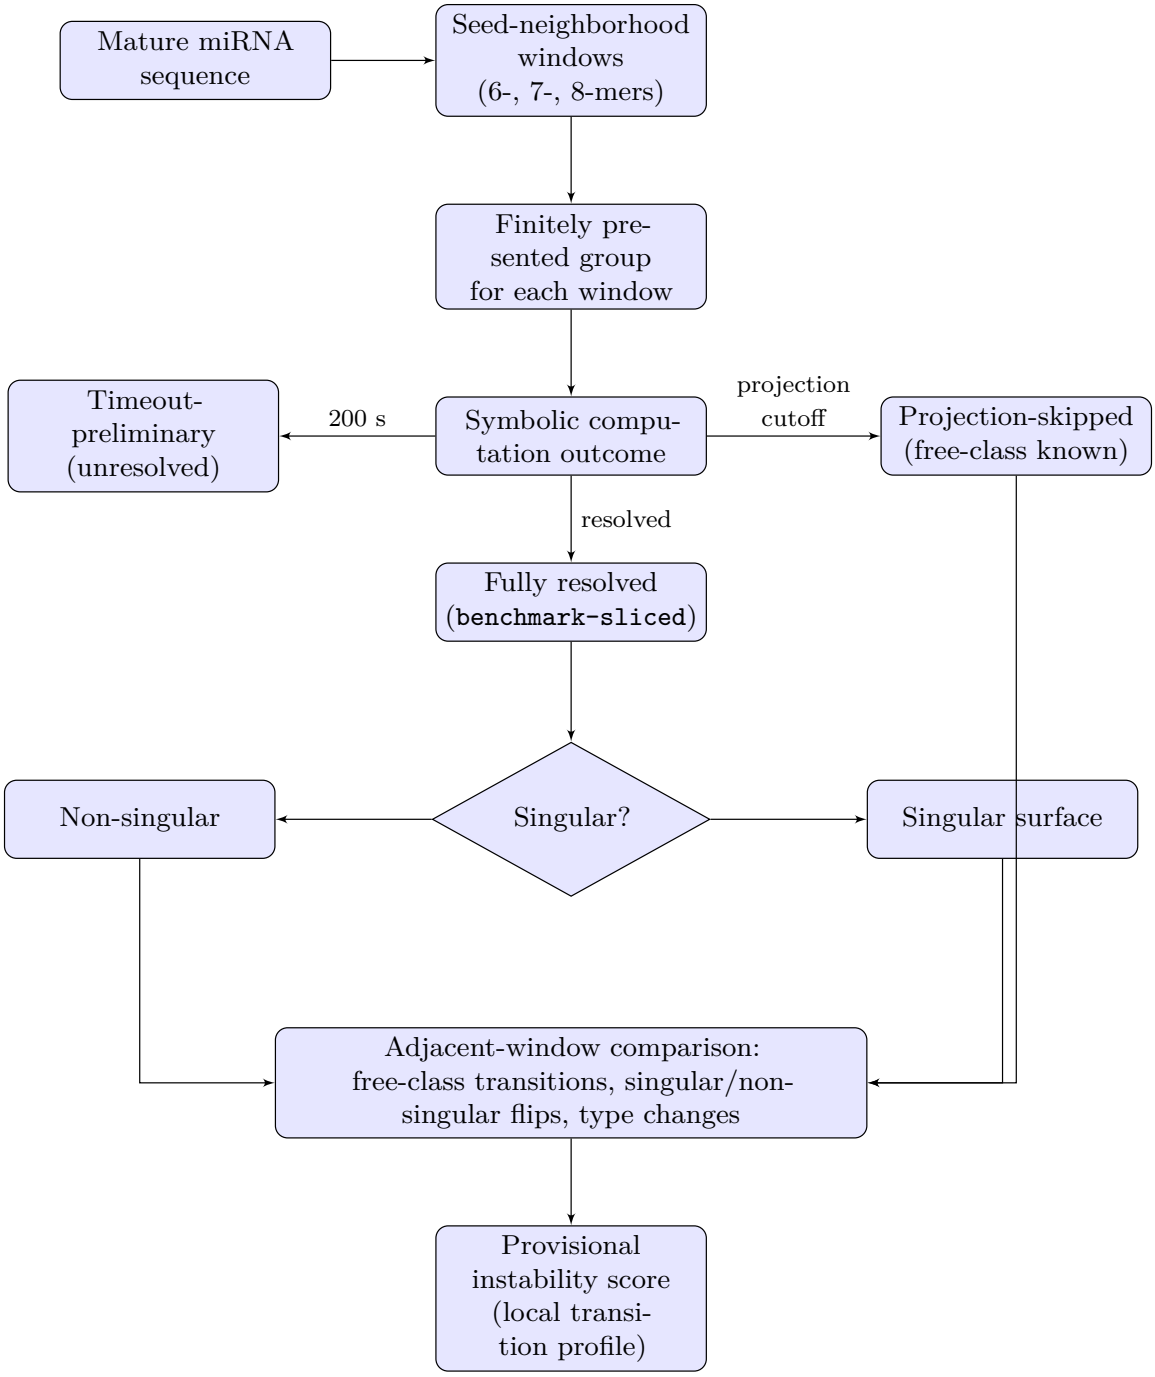

Figure 1: Seed-window algebraic instability workflow. Overlapping seed-neighborhood windows are profiled individually and then compared across adjacent windows to obtain a local instability score.

Table 2: Representative seed-neighborhood window-level algebraic results for the anchor case `hsa-miR-124-3p` and the highest-scoring miR-29 family member `hsa-miR-29a-3p`. Published regression fixtures are marked by a dagger.

miRNA	Motif	Len	Free class	Singular	Singularity type	Tier
miR-124-3p	`UAAGGC`	6	F_3	No	–	Resolved
miR-124-3p	`UAAGGCA`	7	F_3	No	–	Resolved
miR-124-3p	`UAAGGCAC`	8	F_3	No	–	Resolved
miR-124-3p	`AAGGCA`	6	F_2	Yes	`A1 f_(2,{})`	Resolved †
miR-124-3p	`AAGGCAC`	7	π_2	No	–	Resolved †
miR-124-3p	`AAGGCACG`	8	F_2	No	–	Resolved
miR-124-3p	`AGGCAC`	6	F_2	No	–	Resolved
miR-124-3p	`AGGCACG`	7	F_2	No	–	Resolved
miR-124-3p	`AGGCACGC`	8	F_2	No	–	Resolved
miR-29a-3p	`UAGCAC`	6	–	–	–	Timeout
miR-29a-3p	`UAGCACC`	7	F_3	No	–	Resolved
miR-29a-3p	`UAGCACCA`	8	F_3	n/a	n/a (unprojected)	Projection-skipped
miR-29a-3p	`AGCACC`	6	F_2	No	–	Resolved
miR-29a-3p	`AGCACCA`	7	F_2	No	–	Resolved
miR-29a-3p	`AGCACCAU`	8	F_3	n/a	n/a (unprojected)	Projection-skipped
miR-29a-3p	`GCACCA`	6	F_2	No	–	Resolved
miR-29a-3p	`GCACCAU`	7	F_3	No	–	Resolved
miR-29a-3p	`GCACCAUC`	8	F_3	No	–	Resolved

Table 3: Per-miRNA instability summary. Free-class transitions, singular/non-singular flips, and singularity-type changes are counted only over adjacent window pairs where both windows have a known classification in the relevant dimension. Projection-skipped windows contribute to free-class counts but not to singularity-based counts.

miRNA	Windows	Free-class trans.	Sing./NS flips	Type changes	Total trans.	Score
miR-124-3p	9	3	2	2	7	7.0
miR-29a-3p	9	4	0	0	4	4.0
miR-29c-3p	9	4	0	0	4	4.0
miR-15a-5p	9	1	0	0	1	1.0
miR-16-5p	9	1	0	0	1	1.0
miR-486-5p	9	1	0	0	1	1.0
miR-92a-3p	9	0	0	0	0	0.0
let-7a-5p	9	0	0	0	0	0.0
let-7i-5p	9	0	0	0	0	0.0

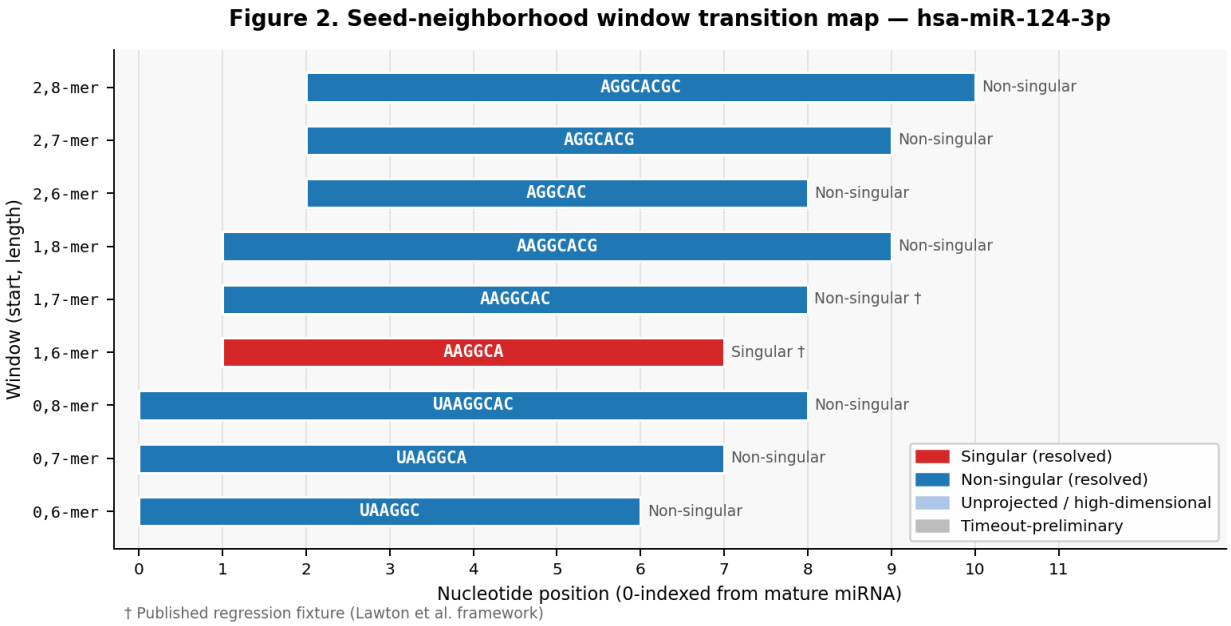

Figure 2: Seed-neighborhood transition map for `hsa-miR-124-3p`. The singular anchor window `AAGGCA` and its non-singular extension `AAGGCAC` reproduce the published regression contrast, while the surrounding seed-neighborhood windows show the broader local transition context in which that contrast occurs. Colors indicate resolved singular and non-singular windows; partially resolved and unresolved categories are indicated separately when present.

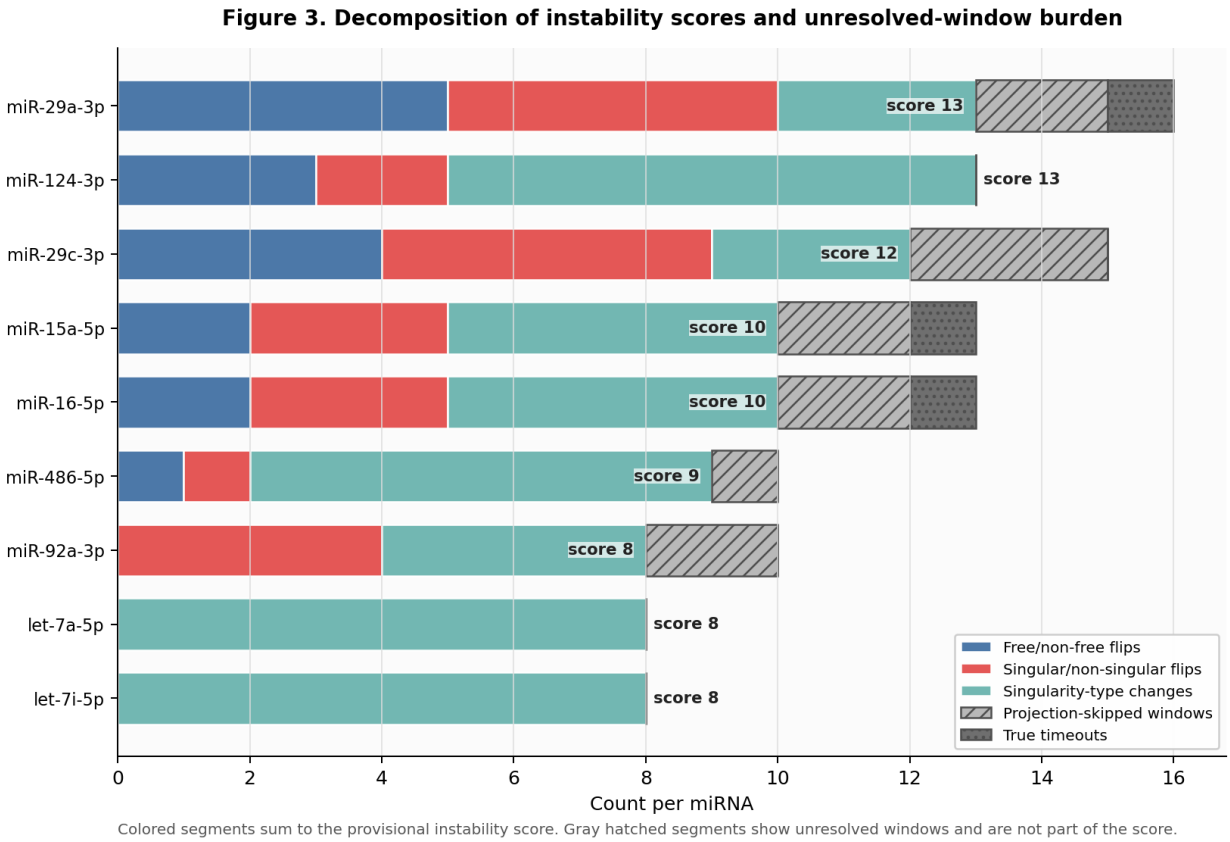

Figure 3: Decomposition of the provisional instability profile for each miRNA in the pilot cohort. Colored segments show the three components that sum to the provisional score: free-class transitions, singular/non-singular flips, and singularity-type changes across adjacent seed-neighborhood windows. Hatched gray segments indicate projection-skipped windows and true timeouts, which represent unresolved symbolic burden but are not included in the score itself.