

The Finitely Presented Recognition Dynamics Program

A mathematical architecture for finite presentations,
infinite recognition dynamics, and exact observations

FPRD Lab — comprehensive program overview

Living edition, 13 July 2026

Status boundary. Milestone status: M29 [complete]. This document reports definitions and theorem candidates at their exact public ledger status. Results through M29 have written arguments and executable challenge evidence but remain *informal proofs*. FPRD-D23 and FPRD-D24 are definitions; FPRD-T44, FPRD-T45, FPRD-T50, and FPRD-T51 are *informal-proof*, never mechanically checked. Agreement among AI systems is useful adversarial redundancy, not external peer review. No external novelty conclusion is asserted. Section 8.1 distinguishes imported facts and unpromoted discovery candidates from the M27–M29 claim-ledger results.

Public source: `publications/fprd-program-overview/fprd-program-overview.tex`

Stable edition path: `publications/fprd-program-overview/fprd-program-overview.pdf`

Contents

1	Executive orientation: the program map	1
1.1	The research turn	1
1.2	The architecture in one page	1
1.3	What this program is not	2
2	Ontology and evidence discipline	2
2.1	Four layers that must not collapse	3
2.2	Exact maps and observation-relative abstractions	3
2.3	Evidence ladder	3
2.4	Three forms of independence	4
3	Deterministic algebra and observational completion (M0–M9)	4
3.1	M0–M3: semantics, guarding, and presentation actions	4
3.2	M4–M7: two compactifications of generated free actions	4
3.3	M5, M8, and M9: completion theorems	5
4	Relational and context-free bridges (M10–M17)	5
4.1	A second recognition doctrine	5
4.2	Least relational equations and labeled continuations	6
4.3	CFG adequacy and controlled normalization	6
4.4	Residual minimality and two forms of guarding	6
5	Universal Dyck control and fixed-envelope geometry (M18–M22)	7
5.1	One recursive frame, varying finite control	7
5.2	A fixed envelope and its exact costs	7
5.3	The canonical Dyck residual and exposure calculus	8
6	Hankel, streaming, controlled quotients, and cut complexity (M23–M25)	8
6.1	Exact Hankel rank from saved-stack coordinates	8
6.2	Finite control and canonical future quotients	8
6.3	A general fixed-cut transfer theorem	9
7	M26 residual sewing: the completed Crawl	9
7.1	Established inputs and the assembly theorem	9
7.2	The global plateau mechanism	10

7.3	Why separate fixed cuts are insufficient	10
8	Crawl, Walk, Run: the earned roadmap	11
8.1	Walk opening discovery: access as a second coordinate	12
9	Claim and evidence status	16
10	Primary literature, provenance, and prior-lab boundary	18
10.1	What the sources own	18
10.2	Prior labs and intellectual provenance	19
10.3	Current novelty boundary	19

1 Executive orientation: the program map

1.1 The research turn

The Finitely Presented Recognition Dynamics program begins with a change in the unit of study. A grammar is not identified with its language, and an infinite family of exact continuation states is not treated automatically as a failed finite-state reduction. The central value path is

$$\boxed{\text{finite recognition presentation} \longrightarrow \text{recognition dynamics} \longrightarrow \text{language or suffix observation.}} \quad (\text{V})$$

The first arrow asks what exact state space, actions, topology, finite observations, and boundary behavior a finite presentation generates. The second asks which part of that richer object a chosen observation exposes. Changing an observation can merge states without changing the presentation; changing a presentation can change internal actions without changing the language. The program keeps both phenomena visible.

The motivating precursor is the recursive continuation family $K_n = b^n$ arising from the ordinary PEG rule $A \leftarrow a A b / \text{epsilon}$. Its members are pairwise distinguishable under exact all-input behavior, so no finite exact continuation quotient exists for that specialization. Yet one action generates the whole family, bounded observations are finite, the orbit has a compact failure boundary, and a counter with a uniform readout realizes the recognition behavior. The important distinction is therefore

$$\text{finitely many exact states} \not\Longleftrightarrow \text{finite description of structured infinite behavior.}$$

This example motivates representation rather than forced collapse. The program asks when exact infinitude is generated, observable, reconstructible, and quantitatively useful.

1.2 The architecture in one page

The milestone ladder through M29 has six broad movements. The fifth closes the Crawl stage; the sixth begins the Walk with exact growth and access invariants.

Movement	Mathematical object	Leverage produced
M0–M9	Deterministic suffix behaviors, guarded presentations, generated orbits, compact finite-shadow towers	Exact semantics, presentation actions, uniform readout, and complete infinite-but-finitely-generated examples
M10–M17	Set-valued relational doctrines, least equations, CFG translations, normalization, residual dynamics, and cycle guarding	A second doctrine, exact CFG bridges, presentation-change control, and finite reconstruction schedules
M18–M22	One universal Dyck frame, fixed-envelope recoding, residual transport, generalized-Dyck residuals, and exposure-filtered actions	A separation of recursive geometry from finite control, with exact coordinate costs and graded action maps
M23–M25	Hankel windows, streaming configurations, controlled future quotients, strict layered recognition, and one-way protocols	Exact representation-to-complexity invariants at fixed horizons and cuts
M26	Compatible cut fields, joint residual rectangles, residual sewing, and global plateaus	An informal-proof bridge from nonuniform fixed-length observations to one uniform recognition dynamic, with an exact failure mode for fixed-cut width
M27–M29	Full-table capacity, sparse-marker rectangles, min-plus controlled-Dyck access, and depth-threshold families	Ambient upper and lower calibrations followed by the first finite-presentation-constrained access theorem and a class-wide linear restriction

The ladder is deliberately cumulative. A later milestone imports the exact theorems owned by earlier milestones instead of reclaiming them. M26, for example, does not reprove Myhill–Nerode minimality, derivative completion, or fixed-cut message minimality. It supplies the exact maps that make these three views parts of one two-parameter object while preserving their earlier claim ownership.

1.3 What this program is not

FPRD is not a renamed CFL-versus-PEG project. It does not claim that every grammar has finite-dimensional coordinates or that compactness implies finite-state recognition. It does not identify parser syntax, extensional language, presentation-generated action, and completed dynamic. It also does not treat bounded computation as proof of a universal statement.

The prospective contribution is a recognition-specific synthesis: finite presentations generate exact dynamics; finite observations assemble into canonical towers; actions have measurable observation exposure; and selected towers yield exact algebraic or machine-resource invariants. Whether this synthesis is externally novel in whole or in part remains a literature and peer-review question.

2 Ontology and evidence discipline

2.1 Four layers that must not collapse

The deterministic foundation fixes four different layers.

1. **Semantic algebra.** Exact recognition behaviors and their constructors, quotiented only by exact behavior equality.
2. **Observational completion.** Finite interfaces such as agreement on all inputs through length N , their quotient maps, and compatible inverse threads.
3. **Presentation-generated dynamics.** A distinguished state and named actions extracted from or supplied by one finite presentation.
4. **Coordinate realization.** An injective state encoding together with every named action and a uniform readout that recovers recognition behavior from coordinates and input.

An implication between layers is a theorem obligation. A universal product embedding is not automatically a useful coordinate realization. A state encoding without a uniform readout is incomplete. Equality of languages does not imply equality of labeled continuation actions. Conversely, two models may share a coordinate topology and action monoid while exposing different remaining-suffix observations.

2.2 Exact maps and observation-relative abstractions

For deterministic suffix behavior, the complete observation of a state p is its value $p(w)$ on every input word. These observations jointly separate exact behaviors. A many-to-one map can therefore be meaningful only after it names the interface it preserves: bounded inputs, a test family, reachable inputs, complete-match acceptance, a fixed total length, or another explicit choice.

This principle explains the program's finite quotients. Agreement through depth N is not exact equality; it is an exact quotient *for the declared bounded interface*. Compatible quotients can reconstruct an exact point or a compact closure, but that reconstruction is itself a theorem, not a change of terminology.

2.3 Evidence ladder

The public ledger distinguishes mathematical and evidential force.

Status	Meaning in this program
definition	A stipulated object or relation with explicit parameters; it does not establish adequacy, existence, or novelty.
source-verified	A statement checked against a named primary source at a reviewed location.
proved-in-code	An exact executable construction with tests, without an unbounded mathematical conclusion unless the domain is exhausted.
bounded-experiment	Finite evidence and counterexample search only.
informal-proof	A written unbounded argument supported by review and often by independent executable challenge, but not proof-assistant checked or externally refereed.
conjecture / open	A precise proposed statement or an explicit unsatisfied proof obligation.

The common pattern is not “tests prove the theorem.” Rather, the written argument owns the infinite

statement; executable models materialize its finite consequences; independent oracles and deliberately broken variants challenge the implementation and the proof seams. Source notes establish what is already known and what they do not imply. Claim promotion follows only after these surfaces agree.

2.4 Three forms of independence

The strongest milestone packets try to preserve three separations.

- A direct semantic evaluator and a structural bounded model should not share the code whose agreement they are supposed to test.
- A theorem proof and a finite challenge should meet at explicit interfaces, not silently substitute for one another.
- Prior literature, prior-lab results, and local AI critique remain inputs until a reviewed repository artifact maps them to a claim.

This discipline matters especially at M26. A finite observed rectangle may exhibit an apparent plateau, but later prefixes can add a new row and later futures can split an old one. Only a genuinely global quotient or a separately certified complete finite carrier supports a regularity conclusion.

3 Deterministic algebra and observational completion (M0–M9)

3.1 M0–M3: semantics, guarding, and presentation actions

M0 fixes the ontology, custody rules, observation interfaces, and independent FPRD claim namespace. M1 then models a finite nonrecursive deterministic PEG behavior as a suffix-decreasing partial endofunction of Σ^* . Identity, failure, terminal left-cancellation, parser-order composition, preferential union, and antidomain form the selected ambient algebra. Equality on all inputs through length N is a finite-index congruence, and the ambient behavior algebra is reconstructed from these bounded quotients.

M2 adds a deliberately narrow recursive fragment. Every recursive dependency is preceded by a dedicated consuming prefix, so the semantic operator is a strict contraction in the shortest-distinguishing-input ultrametric. The metric proof gives one fixed point. A separate termination and equation-satisfaction proof identifies it with ordinary operational evaluation. With depth zero equal to failure, depth $N + 1$ is exact through input length N . The stronger claim that every Ford-well-formed PEG is eventually contractive remains only a conjecture.

M3 begins the presentation layer. Unique call labels retain each source, target, syntax path, and ordered right-success frame. Interpreting those frames in the M2 fixed solution yields named partial continuation actions and a graph-generated dynamic. Variable and label renaming transport this structure, and sequence reassociation preserves the scoped action semantics. Calls beneath predicates and complete failure or commitment control are not included.

3.2 M4–M7: two compactifications of generated free actions

M4 reconstructs the motivating family $K_n = b^n$. The exact orbit is pairwise distinguishable; its depth- N shadow has $N + 2$ states, including the tail/failure class; its closure adjoins one failure boundary. The successor action, inverse tower, and uniform prefix-removal readout form a complete recognition object. A bounded rational coordinate realizes the topology and transition, but is explicitly not a recognition realization by itself.

M6 generalizes this family to context rank r . The finite action alphabet generates the free-monoid continuation orbit. Its one-point compactification has

$$1 + \sum_{j=0}^N r^j$$

depth- N states and a uniform free-word readout. Rank one recovers M4; rank two has exponential shadow growth. This does not by itself yield a pushdown or context-free classification.

M7 holds the free action monoid fixed while changing the recognizing family. Direct finite unary-polynomial actions generate a prefix-tree orbit whose closure contains every infinite context word. At rank at least two the boundary is Cantor-like, not one point. M6 and M7 therefore demonstrate that the same abstract action monoid can acquire nonhomeomorphic recognition compactifications because the behaviors and observations differ.

3.3 M5, M8, and M9: completion theorems

M5 supplies constructive finite interpolation. Every legal bounded suffix table is realized at that bound by a finite nonrecursive PEG expression. Consequently the term-generated deterministic algebra is dense and its completion for the selected quotient chain is the full ambient behavior algebra. Density does not mean that every completed point has one finite PEG definition.

M8 moves from the full ambient algebra to one finitely generated orbit. For each well-formed finite unary-polynomial action presentation, the exact orbit closure is compact and invariant. Its depth- N image is exactly the finite bounded orbit, the restriction maps are surjective and action-compatible, and the inverse limit is homeomorphic and action-conjugate to the exact closure. A single thread readout recovers every remaining-suffix observation.

M9 restores finite graph control for the M3 grammar-facing class. Its bounded graph BFS is exact finite reachability for the selected observation bound, not a sample. Surjective graph bonds reconstruct the compact continuation closure with uniform location, continuation, and ordinary-call readouts. The result remains scoped to positive strongly guarded presentations and their right-success projection.

The deterministic core thus reaches a complete value path:

$$\begin{aligned} \text{finite presentation} &\longrightarrow \text{exact generated orbit} \longrightarrow \text{finite compatible shadows} \\ &\longrightarrow \text{compact dynamic with actions and readout.} \end{aligned}$$

4 Relational and context-free bridges (M10–M17)

4.1 A second recognition doctrine

M10 separates the finite-observation architecture from deterministic outcome type. A finite-observation recognition doctrine supplies an exact Hausdorff behavior algebra, finite bounded algebras with compatible restrictions, restriction-natural constructors, inverse reconstruction, and uniform word readout. The compact finite-action orbit theorem then applies at this abstract level.

Set-valued suffix relations instantiate the interface. Relational composition models sequence and commutative union models nondeterministic choice; preferential nonempty choice remains distinct. Deterministic PEG behavior embeds through singleton outcomes. A relational optional-prefix family can share

M7's coordinates, actions, topology, shadows, growth, and complete-match language while returning a different set of suffixes. Thus language observation is provably lossy even when the surrounding dynamic is aligned.

4.2 Least relational equations and labeled continuations

M11 gives finite positive relational presentations a least-fixed-point semantics. Terms use failure, identity, terminals, calls, commutative union, and relational sequence. The induced operator preserves increasing ω -chain joins; iteration from bottom equals finite derivation semantics. At every input bound, the finite equation system computes exactly the corresponding rows of the global least solution and stabilizes within the finite number of membership facts. Nullable and left-recursive cycles are therefore meaningful under this policy without being reclassified as PEG evaluation.

M12 restores presentation actions. Unique call-occurrence labels retain the ordered right-success frames of M11 syntax. Their least-solution meanings induce restriction-natural set-valued continuation actions. Finite reachable graph shadows reconstruct the compact graph-directed orbit closure with uniform location, continuation, and call readouts. A worked optional-frame family converges to a multivalued successful boundary rather than to failure.

4.3 CFG adequacy and controlled normalization

M13 translates each conventional finite context-free grammar compositionally into the positive relational syntax. Terminals become terminal actions, concatenation becomes sequence, alternatives become union, and canonical occurrence labels retain presentation metadata. At every finite parse height, the direct CFG suffix outcomes equal the relational stage. In the least solution, a nonterminal returns exactly the suffixes remaining after derivable terminal prefixes; empty suffix recovers the classical CFG language. The translated presentation inherits M12's compact continuation dynamic.

M14 proves the reverse expressivity bridge. Finite distributive word normal form converts positive relational terms into finite CFG production bodies and preserves the equation operator in every environment. Hence stages, least suffix solutions, and complete-match languages agree exactly. Dynamic structure is subtler: one source call may split into several normalized clones, several origins may merge, or a call may disappear. A finite origin relation proves that each surviving source frame is the union of its clone frames. This is an action-refinement theorem, not raw dynamic isomorphism.

4.4 Residual minimality and two forms of guarding

M15 places the classical finite case inside the same ontology. The residual dynamic of a language L has states $u^{-1}L$, letter derivatives as actions, and empty-word membership as output. A language is regular exactly when this dynamic is finite; every reachable deterministic recognizer for L maps uniquely and surjectively onto it, so it is the unique minimal exact recognizer. Separately, the guarded language equation $X = AX \cup B$ with positive minimum word length has unique solution A^*B , bounded-exact unfoldings, and an exact lift to remaining-suffix relations.

M16 lifts that geometry to an entire Greibach-normal CFG system. Every nonconstant production consumes a terminal before consulting its nonterminal tail. The language and suffix equation vectors are simultaneously contractive, their unique fixed points equal ordinary CFG derivation semantics, and depth $N + 1$ is exact through length N .

M17 observes that guarding can be a cycle property rather than a property of each individual equation.

A min-plus delay matrix records guaranteed input consumption. If the zero-delay dependency graph is acyclic, a power bounded by the number of variables is contractive. This yields a unique fixed point, agreement with the M11 least solution and finite derivations, a computable block-depth reconstruction schedule, and the same labeled M12 actions.

Together M10–M17 establish that the FPRD architecture is not tied to one choice policy. They also make presentation transformation measurable: exact language and suffix equality can coexist with a nontrivial finite refinement of internal actions.

5 Universal Dyck control and fixed-envelope geometry (M18–M22)

5.1 One recursive frame, varying finite control

M18 separates grammar-specific finite control from universal recursive geometry. A generalized-Dyck equation supplies well-nested traces, a finite residual controller selects compatible traces, and a fixed letter-to-letter map decodes output. The finite state-pair product is positively guarded, has one exact suffix solution, and reconstructs every trace and decoded output row through length N after $N + 1$ unfoldings.

For each nonempty output alphabet, the sourced non-erasing Chomsky–Schützenberger representation gives one finite generalized-Dyck alphabet and fixed readout such that varying finite regular control yields precisely the context-free suffix behaviors. FPRD’s added layer is the exact state-pair recognition dynamic and its same-index finite reconstruction. The worked three-state controller decodes the full suffix behavior of $\{a^n b^n : n \geq 0\}$.

5.2 A fixed envelope and its exact costs

M19 replaces a grammar-dependent matched alphabet with a fixed symbolic envelope. Each source bracket pair receives an injective width- m positional code. Envelope symbols carry polarity, current output, mate output, and one digit. Closing blocks reverse both code position and mate fields, so the block map preserves and reflects nesting. A complete trie-image controller accepts exactly the recoded source-control language; partial and malformed blocks are explicitly nonaccepting. Fixed projection and direct source-block readout have the same exact suffix behavior.

The finite-shadow conversion is sharp:

$$\begin{aligned} N \text{ output symbols consult source depth } \lfloor N/m \rfloor, \\ n \text{ source depth guarantees output depth } m(n+1) - 1. \end{aligned}$$

An explicit finite character codec conjugates admitted symbolic instances with the executable M18 product.

M20 identifies what this coordinate change preserves and what it costs. Each source-controller residual appears at a compiled block-boundary state; source letters become width- m macro-actions; residual equality is preserved; and shortest distinguishing lengths multiply exactly by m . Strong controller quotients lift through compilation. The resulting profile separates code capacity, envelope size, shared-prefix allocation, raw and reachable trie control, residual count, and observation lag. These coordinates expose real tradeoffs rather than compressing all representation cost into one state number.

5.3 The canonical Dyck residual and exposure calculus

M21 computes the exact residual object of bare generalized-Dyck recognition. Every live residual is one finite typed opener stack; one additional state is absorbing failure. Openers push, matching closers pop, mismatches and underflow fail, and neutral symbols act identically. At depth N , every stack of height at most N is visible and all deeper stacks share one tail:

$$|X_N| = 1 + \sum_{h=0}^N k^h.$$

A closer can reveal a saved frame, so the correct finite type is graded: $X_{N+1} \rightarrow X_N$, not an arbitrary same-level map. Compatible graded threads recover exactly the finite stacks and failure, with full Boolean and suffix readout.

M22 abstracts this seam. An action has exposure at most e when agreement at depth $N + e$ forces agreement of acted states at depth N . Such an action is equivalent to one shifted-natural family

$$F_N : X_{N+e} \longrightarrow X_N.$$

These families reconstruct exact actions on inverse threads, declared costs add under composition, and the shortest-observation metric records exposure as a Lipschitz bound. On the universal language carrier, derivative by a word u has sharp exposure $|u|$. M8 becomes exposure zero; M21's pops provide a concrete exposure-one action.

6 Hankel, streaming, controlled quotients, and cut complexity (M23–M25)

6.1 Exact Hankel rank from saved-stack coordinates

M23 turns the M21 tower into an algebraic and streaming invariant. In the depth- N bounded Hankel window, every visible stack owns a canonical closing column. These columns form an identity minor, while factorization through the saved-stack coordinate gives the matching upper bound. Over every field and over the Boolean semiring,

$$\text{rank}(B_N) = \sum_{h=0}^N k^h, \quad |X_N| = 1 + \text{rank}(B_N).$$

The extra state is the observable zero/failure row: it changes state diversity without changing rank.

At a common length- N midpoint, neutral padding or the exact neutral-free parity constraint makes all relevant saved stacks simultaneously reachable. Distinct rows force distinct deterministic streaming configurations. A typed-stack recognizer supplies matching asymptotic storage: $\Theta(\log N)$ bits at rank one and $\Theta(N \log k)$ for multiple bracket types. The claim is exact deterministic one-pass recognition at the stated midpoint interface, not a randomized or multipass theorem.

6.2 Finite control and canonical future quotients

M24 adds a total finite controller. The syntactic stack/controller product is not automatically the semantic object. Finite balanced-word saturation first computes every globally reachable low-stack coordinate,

including coordinates reachable through arbitrarily long balanced detours. Recursive future refinement then forms the exact bounded Boolean quotient S_N .

This quotient is the coarsest valid reachable graded summary and the unique reduced minimum up to the declared isomorphism. Surjective bonds and compact nested fibers reconstruct the exact semantic dynamic. Exact-length blocks give literal controlled Hankel cuts and, at each fixed total length, the unique reachable reduced layered recognizer. The at-most and exact views agree by ordered concatenation:

$$B_N = E_0 \parallel E_1 \parallel \cdots \parallel E_N.$$

6.3 A general fixed-cut transfer theorem

M25 abstracts away from Dyck structure. For any Boolean recognition function at one fixed total length, the exact future row after a prefix is simultaneously

canonical strict-layer state = minimum deterministic one-way message $\leq$ every exact one-pass cut configuration	(C)
---	-----

The row construction is terminal among reachable correct strict grade-preserving deterministic layered recognizers. Its cardinality is the exact minimum state count at that grade and its ceiling logarithm is the exact fixed-length message cost. Every deterministic one-pass recognizer factors through complete cut configurations and therefore needs at least as many. M23 supplies closed profiles; M24 supplies controller-sensitive computable profiles.

The crucial boundary is nonuniformity. M25 builds one canonical program for each selected total length. Its state set and transitions may depend on that length, so M25 alone does not construct one finite time-homogeneous recognizer for all lengths. M26 supplies the missing uniform assembly and proves why small separate cuts do not by themselves imply one.

7 M26 residual sewing: the completed Crawl

7.1 Established inputs and the assembly theorem

M26 starts from three established internal inputs: M15 owns residual dynamics and finite residual minimality; M22 owns shifted derivative actions and completion; M25 owns exact fixed-cut minima. The new datum is a literal two-parameter field

$$C_{m,n}(u, v) = \mathbf{1}[uv \in L], \quad u \in \Sigma^m, \quad v \in \Sigma^n.$$

Literal labels matter. Moving a named symbol across an adjacent cut gives

$$C_{m,n+1}(u, av) = C_{m+1,n}(ua, v). \quad (\text{D})$$

Repeatedly applying (D) moves every future symbol to the prefix axis, showing that a compatible field has the unique language $L_C = \{w : C_{|w|,0}(w, \varepsilon) = 1\}$. Anonymous row counts cannot express this law.

For each prefix u , M26 joins all exact future blocks through N :

$$B_N(u) = (\mathbf{1}[uv \in L])_{|v| \leq N}, \quad X[I, N] = \{B_N(u) : |u| \leq I\}.$$

Each M25 cut is an exact-length projection of this joint row, but the joint row also remembers which blocks share one literal prefix. Horizon restriction gives surjective bonds. Prefixing by a has the forced exposed type

$$X[I, N + 1] \longrightarrow X[I + 1, N], \quad B_{N+1}(u) \longmapsto B_N(ua).$$

The reviewed proof identifies the two assembly orders in **Set** as

$$\lim_I \lim_N X[I, N] \cong O_L \hookrightarrow X_L \cong \lim_N \lim_I X[I, N], \quad (\text{S})$$

where $O_L = \{u^{-1}L : u \in \Sigma^*\}$ and $X_L = \overline{O_L} \subseteq 2^{\Sigma^*}$. Topology is added to the first colimit only after identifying it with the literal orbit. The comparison is dense and is onto exactly when the orbit is closed; an infinite orbit need not have a proper completion.

This proof has undergone internal adversarial review; its finite diagrams have executable witnesses, and the repaired independent challenge reconstructs the stored cells and all three naturality families without importing the M26 production modules. The closure audit therefore advances FPRD-T44 exactly to *informal-proof*. Its unbounded conclusions still rest on the written proof and the imported M22 completion theorem, not on the finite test corpus.

7.2 The global plateau mechanism

Let

$$Q_N(L) = \{B_N(u) : u \in \Sigma^*\}, \quad q_L(N) = |Q_N(L)|.$$

The bonds $Q_{N+1} \rightarrow Q_N$ make $q_L(N)$ nondecreasing. The plateau argument is sharper than bounded eventual constancy. If $q_L(N) = q_L(N + 1)$, the finite bond is bijective, hence depth- N equality and depth- $(N + 1)$ equality coincide. If u and u' agree through N , their children ua and $u'a$ also agree through N for every letter a . Iteration makes the relation a right congruence, and epsilon output after an arbitrary continuation identifies it with full residual equality.

At such a *global* plateau, the exposure-one maps collapse to same-level actions by composing with the inverse bond. The resulting machine has start $[\varepsilon]_N$, fixed letter actions, and epsilon readout; the imported M15 theorem identifies it with the minimal residual DFA. The executable supplied- DFA route checks all members of every quotient fiber, source transitions, readout, surjectivity, and product equivalence. Finite membership-window inference remains only a refutation-capable hypothesis route. Even a table through $2q - 1$ is exact only under an independently certified upper bound of q residuals.

7.3 Why separate fixed cuts are insufficient

For any nonempty finite alphabet, consider the length-only language

$$L_{\square} = \{w : |w| \text{ is a perfect square}\},$$

every word in one (m, n) cut has total length $m + n$. Its entire M25 cut matrix is therefore constant: every fixed cut has one row and zero message bits. The uniform residual orbit is nevertheless infinite. If $p < q$, choose a square $s = t^2 \geq p$ with $2t + 1 > q - p$ and append a word of length $s - p$. The length- p prefix reaches s , while the length- q prefix reaches a length strictly between consecutive squares.

The written proof chooses one alphabet letter and embeds these unary residual separators in the full language. Executable evidence uses the unary alphabet as its finite representative: it exhausts 512 Boolean length families, checks 23,040 width-one cuts and zero-bit transfers, verifies 210 explicit residual

separators, and retains finite false-plateau counterexamples. The repaired standard-library challenge executes 44,400 checks and fourteen malformed alternatives. Together with the unbounded proof and closure audit, this advances FPRD-T45 exactly to *informal-proof*; finite observations alone still certify neither regularity nor a global plateau.

8 Crawl, Walk, Run: the earned roadmap

The roadmap is sequential in evidential leverage, not merely chronological. Each stage must produce a reusable invariant, theorem, or obstruction before the next stage earns its place.

Stage	Research work	Earning gate	Leverage if earned
Crawl (earned)	Finish M26's literal sewing theorem, global plateau theorem, square-length separation, independent challenge, and closure audit.	T44/T45 receive complete unbounded proofs, executable seams, adversarial review, accurate ledger promotion, full repository verification, and governed publication.	A trustworthy map from fixed cuts to one uniform residual dynamic; exact boundary between finite evidence and global proof.
Walk (current)	Classify observational growth $G_L(I, N)$ and $q_L(N)$, uniformization cost, exposure collapse, and transformations of presentations on controlled families.	At least one family-level theorem transfers inequalities, realizability constraints, or computable profiles across a transformation; examples must separate coordinates that a scalar invariant misses.	A usable calculus of how semantics, control, coordinate geometry, and observation lag change under compilation or normalization.
Run	Apply the invariants to difficult automata and complexity problems, including controlled recursive families and uniform/nonuniform recognition gaps.	A problem yields a genuinely new upper bound, lower bound, classification, obstruction, or algorithm, with its machine model and source boundary explicit.	Evidence that FPRD is more than an internal reorganization: its objects solve or sharpen external recognition problems.
Later	Test selected broader finitely presented dynamics only after the recognition theory has earned transport.	A precise functor, invariant, or counterexample survives outside language recognition without erasing the readout discipline.	A justified broader theory rather than analogy-driven scope expansion.

Near-term Walk candidates include controlled-Dyck quotient-growth regimes, relations between Hankel rank and future-quotient growth, sharp observation moduli under block recoding, and criteria for when presentation refinement preserves or increases uniformization cost. Symbolic dynamics, state complexity, profinite dynamics, communication complexity, and subword complexity are comparator neighborhoods; analogies enter the program only when they transfer a theorem, inequality, construction, or falsifiable constraint.

8.1 Walk opening discovery: access as a second coordinate

The first Walk reconnaissance has isolated a prospective companion to quotient growth. For a language L , define the *horizon- N prefix-access radius*, provisionally, by

$$a_L(N) = \min\{I : G_L(I, N) = q_L(N)\}. \quad (1)$$

This notation is design input, not a public ledger definition. It measures how far one must search among literal prefixes to witness every class that is visible at future horizon N . Thus $q_L(N)$ measures how many observations exist, while $a_L(N)$ measures how late the last one first appears.

The fixed-state vocabulary is classical. Automata testing calls a word that reaches a state an *access sequence* and a collection reaching every state a *state cover* [13]. At each horizon, the prefix ball through $a_L(N)$ similarly supplies a representative for every bounded row. For a fixed FSM, the maximum shortest transfer length is also called its *accessibility degree* [14]. But $Q_N(L)$ is a graded observation quotient, not automatically a standalone automaton with closed same-level actions: appending a letter consumes future horizon. The candidate coordinate therefore adapts state-access depth to a varying observation tower; it does not claim the classical fixed-automaton notion.

Three candidate basic laws have short direct derivations. The minimum in (1) exists because $Q_N(L)$ is finite and every one of its classes has a literal prefix representative. The sequence $a_L(N)$ is nondecreasing: projection from horizon $N + 1$ onto horizon N sends a prefix-complete set of the former classes onto all of the latter. Finally, if L has a reachable residual DFA with κ states, then $a_L(N) \leq \kappa - 1$ for every N , because a shortest access word to any reachable state has length at most $\kappa - 1$. This regular-language upper bound is the classical accessibility-degree bound, not new FPRD content. If $d = |\Sigma|$, elementary prefix counting also gives

$$q_L(N) \leq \sum_{j=0}^{a_L(N)} d^j;$$

in particular, unary languages satisfy $a_L(N) \geq q_L(N) - 1$, while for $d \geq 2$ inversion gives $a_L(N) \geq \lceil \log_d((d - 1)q_L(N) + 1) \rceil - 1$. These are proof sketches under audit, not promoted FPRD claims.

Together these laws yield a candidate regularity criterion:

$$L \text{ is regular} \iff \sup_N a_L(N) < \infty \iff a_L(N) \text{ is eventually constant.} \quad (2)$$

The forward implication is the $\kappa - 1$ bound and monotonicity. Conversely, a uniform access bound makes $q_L(N)$ uniformly bounded by the prefix-volume inequality, so its nondecreasing integer tower reaches the global plateau of FPRD-T45. For a regular language, once a finite horizon distinguishes every pair of residual states, $a_L(N)$ is exactly the accessibility degree of the minimal residual DFA. Thus the graded profile converges to the classical invariant. An M26 quotient plateau forces an access plateau because the two finite prefix partitions then coincide, but the converse fails: bare Dyck has $a_D(0) = a_D(1) = 1$ while remaining nonregular. An independent finite-DFA challenge matched these laws on 1,920 minimal reachable machines and 15,296 horizon profiles. Equation (2) is a derived candidate dependent on M26's informal proof, not a promoted theorem.

The future axis also recovers a classical testing depth. For a minimal regular DFA, its *distinguishability degree* is the maximum shortest suffix separating a pair of states [14]. Therefore the first horizon with $q_L(N) = \kappa$ is exactly that degree (with value zero for a one-state DFA), and from that horizon onward $a_L(N)$ is the accessibility degree. On regular languages, the two saturation axes of $G_L(I, N)$ thus recover classical reachability and distinguishability depths; the proposed profile extends their joint partial coverage before saturation and beyond the finite-state regime.

Boolean language transformations have a direct candidate count law. For languages over one alphabet and any pointwise Boolean operation Φ ,

$$G_{\Phi(L_1, \dots, L_m)}(I, N) \leq \prod_{j=1}^m G_{L_j}(I, N), \quad q_{\Phi(L_1, \dots, L_m)}(N) \leq \prod_{j=1}^m q_{L_j}(N). \quad (3)$$

Indeed, the output row is obtained by applying Φ pointwise to the tuple of component rows. Complement is stronger: pointwise negation preserves the prefix partitions, so $G_{\bar{L}} = G_L$, $q_{\bar{L}} = q_L$, and $a_{\bar{L}} = a_L$ exactly. There is no analogous access bound for intersection from marginal access alone. For even $T \geq 2$, let $L_T = \{a, a^T\}$ and fix the two-state filter $R_{\text{even}} = \{a^{2j} : j \geq 0\}$. Both have horizon-zero access one, but their intersection $\{a^T\}$ has access T . Even a fixed regular controller can therefore create arbitrary correlation delay. An independent challenge checked 20,000 random finite windows and all 50 even thresholds through 100. These direct laws are unpromoted design results.

On a unary alphabet the proposed coordinate has an exact classical specialization. Given $L \subseteq \{a\}^*$, form the one-sided characteristic word $x_i = \mathbf{1}[a^i \in L]$. Then

$$B_N(a^i) = (x_i, x_{i+1}, \dots, x_{i+N}), \quad q_L(N) = p_x(N+1), \quad a_L(N) = A_x(N+1) - (N+1), \quad (4)$$

where $p_x(n)$ is factor complexity and $A_x(n)$ is the shortest prefix length containing every length- n factor. Burns also names the one-based last-start bound $S_x(n)$ and proves $A_x(n) = S_x(n) + n - 1$ [7]; hence $a_L(N) = S_x(N+1) - 1$ under the zero-based FPRD prefix convention. Appearance complexity is not merely an analogy: it owns the unary specialization. Any new FPRD content must lie in the branching-prefix-tree extension, its non-unary laws, or controlled-family classifications. A separate bounded indexing challenge checked 1,530 eventually-zero characteristic-word and horizon cases, comparing 37,653 residual cells with 37,653 separately constructed factor cells; it challenges the offset in (4), not its unbounded proof.

This bridge immediately sharpens the M26 square example. Firicel records quadratic factor complexity for the coefficient sequence supported on the perfect squares [10]; the harmless distinction at its initial coefficient does not change the asymptotic. Therefore the imported unary result gives

$$q_{L_{\square}}(N) = \Theta(N^2),$$

strictly strengthening M26's internally proved $q_{L_{\square}}(N) \geq N + 2$. The quadratic estimate belongs to classical factor complexity, not to FPRD-T45. A separate elementary square-gap calculation performed for this edition gives the candidate exact access law

$$a_{L_{\square}}(2m) = (m+1)^2 + 1, \quad a_{L_{\square}}(2m+1) = (m+1)^2 + 2. \quad (5)$$

For even horizons the last new length- $(N+1)$ factor is the all-zero factor; for odd horizons it is $0^N 1$. Consecutive-square gaps show both that the displayed factor cannot appear earlier and that every factor containing a one has already appeared by that start. An independent enumeration matched (5) at all 61 horizons through 60; a separate high-range challenge matched all 201 horizons through 200 while comparing 426,263,329 factor cells. Those checks attack the parity thresholds and offsets, not the unbounded gap proof. Thus square lengths provisionally have quadratic observation count and quadratic first-access cost, while their full $G(I, N)$ rectangle remains open.

The count/access separation is already exact on unary singletons. For $L_T = \{a^T\}$,

$$q_{L_T}(N) = \min(T, N) + 2, \quad a_{L_T}(N) = \begin{cases} T, & N < T, \\ T + 1, & N \geq T. \end{cases} \quad (6)$$

At $N = 0$, therefore, every member of this family has exactly two visible classes, but the access radius is the arbitrary parameter T . A bounded reconnaissance checked 60 such profiles; the unbounded argument simply tracks the residuals $\{a^{T-i}\}$ and the empty residual. Formula (6) is a candidate separation theorem, not yet a ledger result.

The strongest controlled nonregular candidate is the generalized-Dyck language $D_{k,\ell}$ with $k \geq 1$ bracket types and any number $\ell \geq 0$ of neutral symbols. Writing $S_k(n) = \sum_{h=0}^n k^h$, the candidate exact rectangle is

$$G_D(I, N) = \begin{cases} 1, & I = 0, \\ 1 + S_k(\min(I, N)), & I \geq 1, \end{cases} \quad q_D(N) = 1 + S_k(N), \quad a_D(N) = \max(1, N). \quad (7)$$

The exceptional $I = 0$ branch matters: the empty prefix has not yet witnessed the failure row. At positive budget, one underflowing closer supplies that row, opener words reach every stack of height at most I , and horizon N keeps stacks through height N while merging deeper stacks with failure. This is a fixed-horizon saturation event, not the global adjacent plateau of FPRD-T45; indeed $q_D(N)$ increases at every horizon.

The independent pre-promotion challenge covered ten declared alphabets, 146 (I, N) grid points, ranks one through three, zero and positive neutral ranks, nonlexicographic symbol orders, both zero axes, mismatch, underflow, and deep-stack collapse. It compared 56,780 whole-word memberships, 71,870 residual-future cells, and 115,397 separately constructed oracle cells without finding a counterexample. Those finite checks challenge the derivation of (7); they do not prove its unbounded quantifiers.

The M24 balanced-detour controller belongs to a sharper controlled family. Let $D_k^{\geq r}$ contain the generalized-Dyck words whose stack height reaches at least $r \geq 1$. A finite controller records that threshold, and neutral symbols do not alter the profile. Before the witness, a stack of height $h < r$ becomes visible only when $N \geq 2r - h$; afterward the ordinary Dyck stack residual remains. FPRD-T51 proves the exact family laws

$$q_{k,r}(N) = 1 + S_k(N) + \sum_{h=\max(0, 2r-N)}^{r-1} k^h, \quad a_{k,r}(N) = \max(2r, N), \quad (8)$$

with an empty final sum when its lower endpoint exceeds $r - 1$. The extra rows are precisely the visible pre-witness stacks. A post-witness stack below height r first appears after reaching the threshold and returning, at access length $2r - h$; this proves the access floor in the display. In particular, $q_{k,r}(0) = 2$ for every threshold while $a_{k,r}(0) = 2r$, so fixed observation count permits arbitrarily late access even within this controlled nonregular scheme. A direct whole-word oracle matched the full $G/q/a$ rectangle across bracket ranks one and two and zero/positive neutral rank: 56 threshold/horizon profiles, 392 grid cells, and 4,375,254 memberships. The written row classification, rather than this finite enumeration, carries the unbounded quantifiers.

The complementary controller branch, whose depth stays below r , is regular and has the exact profile

$$q_{k,<r}(N) = 1 + S_k(\min(N, r - 1)), \quad a_{k,<r}(N) = \max(1, \min(N, r - 1)).$$

It stabilizes at $1 + S_k(r - 1)$ residuals, accessibility degree $\max(1, r - 1)$, and distinguishability degree $r - 1$. The direct oracle checked the threshold and shallow branches jointly, totaling 8,750,508 membership predicates. This pair is an internal controlled-family check of the regularity criterion (2).

At $k = 1, r = 2$, the language is exactly the rank-one Dyck language with the flat words removed,

$D_1 \setminus \{()^j : j \geq 0\}$, which is the existing M24 balanced-detour sentinel. Equation (8) specializes to

$$q_{\text{nest}}(N) = \begin{cases} 2, & N = 0, \\ 3, & N = 1, \\ 4, & N = 2, \\ 6, & N = 3, \\ N + 4, & N \geq 4, \end{cases} \quad a_{\text{nest}}(N) = \max(4, N). \quad (9)$$

For $N \geq 4$ the full rectangle is

$$G_{\text{nest}}(I, N) = 1 + 2 \cdot \mathbf{1}[I \geq 1] + \max(0, \min(I, N) - 1) + \mathbf{1}[I \geq 3] + \mathbf{1}[I \geq 4].$$

The terms are the pre-witness empty row; pre-witness height one and failure; ordinary post-witness stack rows of heights two through N ; and the delayed height-one and empty-stack rows. A production replay and an independent whole-word oracle both match (9); the latter challenged 98 grid cells through horizon ten using 22,351,245 direct membership cells. Finite agreement does not prove the semantic residual classification; the M29 proof supplies that step. This example is the program's smallest exact finite-controller transformation instance: finite control adds two eventual quotient classes and raises the access floor from one to four while preserving linear growth.

FPRD-T50 gives an exact calculus before taking the coarse bound. For controller states p, q , put

$$d_M(p, q) = \min\{|v| : v \in D_{k,\ell}, \delta(p, v) = q\},$$

with infinity when no balanced witness exists. Empty, neutral, concatenation, and matched-wrapping constructors compute these distances by a finite min-plus relaxation. Numerically, d_M is the greatest matrix satisfying the resulting upper-bound inequalities—the all-zero matrix is a smaller solution—but it is the pointwise minimum among generated witness costs.

Let R_Q be the start-reachable states of the underlying controller DFA, ignoring stack legality, and set

$$B_M = \max\{d_M(p, q) : p \in R_Q, d_M(p, q) < \infty\}.$$

Unreachable controller gadgets cannot inflate this constant. If $\rho_M(s, q)$ is the shortest prefix reaching raw stack/controller coordinate (s, q) , the last unmatched opener gives the exact recurrence

$$\rho_M(\epsilon, q) = d_M(q_0, q), \quad \rho_M(so, q) = \min_p (\rho_M(s, p) + 1 + d_M(\delta(p, o), q)).$$

For an M24 bounded-future class $K \in S_N$, let $\alpha_{M,N}(K)$ be the minimum arrival of any raw representative, assigning the failure tail cost one. Then

$$a_{L_M}(N) = \max_{K \in S_N} \alpha_{M,N}(K). \quad (10)$$

The order is essential: minimize away redundant presentation coordinates before maximizing language-level cost.

Finally, M24's unmatched-opener decomposition replaces the $h + 1$ balanced segments around a saved height- h stack by shortest witnesses for the same ordered controller-state pairs. This preserves both the stack and the controller state and gives

$$a_{L_M}(N) \leq \max\{1, (N + 1)B_M + N\}. \quad (11)$$

The class-wide law is pointwise optimal. For every $N \geq 1$, an explicit $(4N + 4)$ -state controller with two bracket pairs and N neutrals has $B_M = 2$ and a uniquely isolated horizon- N row whose first prefix witness

has length $3N + 2$. Bare Dyck attains the $N = 0$ guard. The witness presentation depends on N ; no one fixed presentation is claimed extremal at all horizons. Failure has a one-letter representative, and every raw stack above height N has the all-false horizon row; semantic quotienting can only lower the bound. The constant B_M belongs to the chosen presentation, while a_{LM} remains a language invariant. The import-independent M29 challenge exercises four controllers, 60 raw arrivals, 404,061 literal membership cells, and six executed mutations. Combined with (2), this says that a fixed controlled-Dyck language has bounded access exactly when regular, and otherwise has unbounded but at-most-linear access. The at-most-linear half is M29-internal; the nonregular-unbounded half retains its dependency on the M26/T45 global plateau theorem.

The access analysis now separates six regimes:

Family	Observation $q_L(N)$	count	Access $a_L(N)$	Separation carried
Regular, κ resid- uals	eventually κ		eventually classical degree $\rho_{\text{acc}} \leq \kappa - 1$	Both towers stabilize; bounded access charac- terizes this regime.
Singleton $\{a^T\}$	$\min(T, N) + 2$		T or $T + 1$	Fixed count can have arbitrar- ily late access.
Square lengths	$\Theta(N^2)$		exact parity law (5)	Count and access are both quadratic.
Dyck, $k = 1$	$N + 2$		$\max(1, N)$	Both coordinates are linear.
Dyck, $k \geq 2$	$\Theta(k^N)$		$\max(1, N)$	Exponential count has linear, hence logarithmic-in-count, access.
Depth $\geq r$	$1 + S_k(N) + S_k(r - 1)$ eventually		$\max(2r, N)$	Finite control adds pre- witness rows; $q(0) = 2$ hides access $2r$.

M27–M29 have now promoted three steps: the unrestricted capacity baseline, sparse marker calibration, and the exact access calculus for controlled Dyck presentations. They include a linear presentation bound and exact depth threshold profiles. The square-length parity law, general Boolean transformation observations, and a minimax theory independent of presentation remain discovery results. The next narrow question is not another broad growth census: it is whether optimizing the tropical access certificate over equivalent finite presentations yields a useful language invariant or a constructive procedure for improving presentations.

9 Claim and evidence status

The table below is a compressed ledger view. “Written proof” means an unbounded repository argument at informal-proof, commonly accompanied by a separate executable challenge. It does not mean formal verification or external refereeing.

Milestone	Ledger claims	Evidence-bearing result
M0	FPRD-D01 [definition], FPRD-D02 [definition], FPRD-D03 [definition], FPRD-D04 [definition]	Four-layer ontology, exact-versus-bounded maps, equivalence boundaries, and uniform-readout requirement.
M1	FPRD-T01 [informal-proof], FPRD-T02 [informal-proof]	Deterministic suffix-function algebra and bounded congruence/inverse tower.
M2	FPRD-T04 [informal-proof], FPRD-T05 [informal-proof]	Strong guarding, unique fixed point, operational adequacy, and exact finite unfolding.
M3–M5	FPRD-T06 [informal-proof], FPRD-T07 [informal-proof], FPRD-T03 [informal-proof]	Presentation action extraction, b^n dynamic, and full ambient observational completion.
M6–M9	FPRD-T08 [informal-proof], FPRD-D05 [definition], FPRD-T09 [informal-proof], FPRD-T10 [informal-proof], FPRD-T11 [informal-proof]	Free-context and prefix-tree families, compact orbit representation, and graph-directed guarded lift.
M10–M12	FPRD-D06 [definition], FPRD-T12 [informal-proof], FPRD-T13 [informal-proof], FPRD-D07 [definition], FPRD-T14 [informal-proof], FPRD-T15 [informal-proof], FPRD-D08 [definition], FPRD-T16 [informal-proof], FPRD-T17 [informal-proof]	Relational doctrine, positive least equations, exact bounded solving, and relational presentation dynamics.
M13–M14	FPRD-D09 [definition], FPRD-T18 [informal-proof], FPRD-T19 [informal-proof], FPRD-D10 [definition], FPRD-T20 [informal-proof], FPRD-T21 [informal-proof]	CFG suffix adequacy, inherited compact dynamics, semantic normal form, and finite action refinement.
M15–M17	FPRD-D11 [definition], FPRD-D12 [definition], FPRD-T22 [informal-proof], FPRD-T23 [informal-proof], FPRD-D13 [definition], FPRD-T24 [informal-proof], FPRD-T25 [informal-proof], FPRD-D14 [definition], FPRD-T26 [informal-proof], FPRD-T27 [informal-proof]	Residual minimality, Arden guarding, Greibach system dynamics, and cycle-level delay certificates.
M18–M20	FPRD-D15 [definition], FPRD-T28 [informal-proof], FPRD-T29 [informal-proof], FPRD-D16 [definition], FPRD-T30 [informal-proof], FPRD-T31 [informal-proof], FPRD-D17 [definition], FPRD-T32 [informal-proof], FPRD-T33 [informal-proof]	Universal controlled-Dyck suffix dynamics, fixed-envelope recoding, and exact residual/cost transport.

Milestone	Ledger claims	Evidence-bearing result
M21–M22	FPRD–D18 [definition], FPRD–T34 [informal-proof], FPRD–T35 [informal-proof], FPRD–D19 [definition], FPRD–T36 [informal-proof], FPRD–T37 [informal-proof]	Exact generalized-Dyck residual tower and exposure-filtered action reconstruction.
M23–M25	FPRD–D20 [definition], FPRD–T38 [informal-proof], FPRD–T39 [informal-proof], FPRD–D21 [definition], FPRD–T40 [informal-proof], FPRD–T41 [informal-proof], FPRD–D22 [definition], FPRD–T42 [informal-proof], FPRD–T43 [informal-proof]	Exact Hankel/streaming profiles, controlled future quotients, layered minima, and one-way transfer.
M26	FPRD–D23 [definition], FPRD–T44 [informal-proof], FPRD–T45 [informal-proof]	Compatible uniform cuts, residual sewing, global plateau collapse, conditional finite reconstruction, and the fixed-slice/uniform-recognition separation are closed at informal-proof status.
M27	FPRD–T46 [informal-proof], FPRD–T47 [informal-proof]	Maximal binary bounded-row capacity, simultaneous realization by one decidable cone language, full-shift density, and near-optimal access.
M28	FPRD–T48 [informal-proof], FPRD–T49 [informal-proof]	Exact sparse-marker residuals, prefix/future rectangles, jump-controlled access, and fixed-history access separation.
M29	FPRD–D24 [definition], FPRD–T50 [informal-proof], FPRD–T51 [informal-proof]	Exact min-plus controlled-Dyck access, reachable-source linear certificate, and exact depth-threshold profiles.
Horizon	FPRD–C01 [conjecture]	Eventual contractivity for every Ford-well-formed finite PEG may be false and is not on the active critical path.

10 Primary literature, provenance, and prior-lab boundary

10.1 What the sources own

The literature record prevents classical ingredients from being redescribed as FPRD discoveries. Ford supplies ordinary PEG semantics; McLean supplies partial-function operation vocabulary; guarded iteration and contraction sources calibrate M2 and M17. Tarski supplies the complete-lattice fixed-point background used by M11. Winter, Bonsangue, and Rutten provide coalgebraic CFG comparators. Brzozowski and Wingbrant provide derivatives, residual recognizers, regularity, and minimality. The non-erasing Chomsky–Schützenberger source owns the universal context-free representation used by M18.

Alur and collaborators provide visibly-pushdown language and congruence comparators; Meakin and Sapir provide polycyclic-monoid context; Bryant provides the reduced ordered-function-graph comparator; Yao and Roughgarden provide the communication and streaming settings. These sources calibrate vocabulary, known theorems, and category boundaries. The repository’s proof notes own the specific readout-aware FPRD assembly they claim, at informal-proof only.

Primary PDFs and extracted paper text are intentionally absent from the public repository. Verified ex-

traction notes live under `docs/literature/`; the source documents remain in local custody. A literature analogy is not claim evidence until exact definitions, theorem scope, and nonimplications are recorded.

10.2 Prior labs and intellectual provenance

PEG Lab 1 pursued concrete witness and census routes. Its candidates, palindrome machinery, and experimental code are not imported into FPRD. PEG Lab 2 is the direct technical precursor. Its finite nonrecursive continuation-aware compilation and its $K_n = b^n$ recursive exact-quotient obstruction remain PEG Lab 2 results. FPRD cites them as prior work and independently reconstructs only what its own claims require. The CFL-versus-ordinary-PEG target remains outside this program.

Joshua Gay’s recorded contribution is the research-direction turn: treat the infinite continuation family as a representation problem, use Robson’s globalization work as proof architecture rather than as a literal transferred theorem, raise the unit of study from individual grammars to families of generated dynamics, and seek a theory broader than a richer PEG encoding. Subsequent formal vocabulary, proofs, code, and literature synthesis were AI-assisted and remain subject to ordinary mathematical verification.

Robson’s role is methodological. No claim says FPRD spaces are semialgebraic or that Robson’s theorem applies directly. The transferable question is whether awkward local data can be placed in a faithful global representation where limits, boundaries, actions, and invariants become tractable.

10.3 Current novelty boundary

No novelty claim is made for partial-function algebras, inverse limits, guarded fixed points, derivatives, Myhill–Nerode theory, Dyck languages, Hankel rank, streaming configuration arguments, or deterministic one-way row complexity. The scoped FPRD contribution is their explicit recognition-specific integration with presentation provenance, uniform readout, observation exposure, and controlled transformations.

The closest M26 comparators located for this edition are now explicit. The completed residual orbit, its dense reachable suborbit, and its finite collapse for regular languages coincide with Steinberg’s minimal topological automaton [1], with later Stone-recognition corroboration [2]. The bounded quotient tower is a final-sequence construction; the plateau mechanism is Moore–Nerode refinement in the form studied by Ganty, Gutiérrez, and Valero [3]; the warning that a finite observation rectangle is not a global certificate is the same basic boundary exposed by Angluin’s observation-table method [4]; and observation-relative minimization has a functorial factorization formulation [5]. Hora’s automata topoi already own a total presheaf/section setting [6], so any later FPRD gluing claim must distinguish a genuinely partial-data obstruction rather than relabel that total case. These mechanisms are imports or adaptations. The scoped FPRD contribution is only their readout-aware, presentation-sensitive integration and whatever nontrivial theorems can be proved for the two-parameter growth profile; the sewing construction is not presented as “genuine novelty.”

The first Walk source audit sharpens a different boundary. Follower-set sequences count full right-infinite futures after legal words of one *exact* past length; finite total follower-set collections characterize sofic shifts, and small exact-length counts support separate gap theorems and conjectures [8]. Those results are mandatory comparators, but they vary the past-length axis with future horizon at infinity, whereas $q_L(N)$ varies a finite future horizon over all prefixes. French also proves that sofic follower-set sequences can be eventually periodic without being constant [9], unlike the monotone $q_L(N)$ tower. Neither result can be imported as a horizon-gap theorem.

Similarly, ordinary weighted Hankel theory identifies least weighted-automaton dimension with real Hankel rank [11], and nested Hankel theory characterizes weighted visibly pushdown functions using a rank

bound over well-matched rows and columns [12]. FPRD's G_L and q_L count distinct Boolean rows, not their field-linear span, and the Dyck rows include unbalanced suffixes leading to failure. These are neighboring linear theories, not ownership of the two-parameter Boolean row-count profile.

Operational summary. The program now has a substantial informal-proof architecture through M29. The Crawl stage has closed with exact proof boundaries, executable seams, adversarial review, and a governed claim ledger. The Walk has produced its first finite-presentation theorem: controlled-Dyck access is exactly computable and at most linear. Its next obligation is not more breadth but a presentation-minimax test: determine which tropical access costs can be lowered by changing equivalent finite coordinates and which, if any, survive as genuine language invariants.

References

- [1] Benjamin Steinberg. Topological dynamics and recognition of languages. *arXiv:1306.1468*, 2013. <https://arxiv.org/abs/1306.1468>.
- [2] Jorge Almeida and Ondřej Klíma. Eilenberg correspondence for Stone recognition. *arXiv:2507.00409*, 2025. <https://arxiv.org/abs/2507.00409>.
- [3] Pierre Ganty, Elena Gutiérrez, and Pedro Valero. A congruence-based perspective on automata minimization algorithms. *LIPICs MFCS*, 138:77:1–77:14, 2019. <https://doi.org/10.4230/LIPICs.MFCS.2019.77>.
- [4] Dana Angluin. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2):87–106, 1987. [https://doi.org/10.1016/0890-5401\(87\)90052-6](https://doi.org/10.1016/0890-5401(87)90052-6).
- [5] Thomas Colcombet and Daniela Petrişan. Automata minimization: a functorial approach. *arXiv:1711.03063*, 2017. <https://arxiv.org/abs/1711.03063>.
- [6] Ryuya Hora. Topoi of automata I: four topoi of automata and regular languages. *arXiv:2411.06358*, 2024. <https://arxiv.org/abs/2411.06358>.
- [7] Rob Burns. The appearance function for paper-folding words. *arXiv:2210.14719*, 2022. <https://arxiv.org/abs/2210.14719>.
- [8] Thomas French, Nic Ormes, and Ronnie Pavlov. Subshifts with slowly growing numbers of follower sets. *arXiv:1509.01273*, 2015. <https://arxiv.org/abs/1509.01273>.
- [9] Thomas French. Characterizing follower and extender set sequences. *arXiv:1508.02802*, 2015. <https://arxiv.org/abs/1508.02802>.
- [10] Alina Firicel. Subword complexity and Laurent series with coefficients in a finite field. *arXiv:1001.2548*, 2011. <https://arxiv.org/abs/1001.2548>.
- [11] Borja Balle, Prakash Panangaden, and Doina Precup. A canonical form for weighted automata and applications to approximate minimization. *arXiv:1501.06841*, 2015. <https://arxiv.org/abs/1501.06841>.
- [12] Michael Labai and Johann A. Makowsky. Hankel matrices for weighted visibly pushdown automata. *arXiv:1512.02430*, 2016. <https://arxiv.org/abs/1512.02430>.
- [13] Loes Kruger, Sebastian Junges, and Jurriaan Rot. Small test suites for active automata learning. *TACAS*, pages 109–129, 2024. <https://arxiv.org/abs/2401.12703>.
- [14] Adenilso da Silva Simão, Alexandre Petrenko, and José Carlos Maldonado. Comparing finite state machine test coverage criteria. *IET Software*, 3(2):91–105, 2009. <https://doi.org/10.1049/iet-sen.2008.0018>.

Afterword

FPRD program overview

Methodological and contribution addendum. The paper's mathematical scope and review status remain as stated in the paper.

The Lab's approach

FPRD means **Finitely Presented Recognition Dynamics**. Its starting point is to distinguish a finite description, the possibly infinite process it generates, and what an observer can learn from that process. Equal answers need not come from equal internal structures. Conversely, an infinite state space may have a short, exact description and useful coordinates; failure to compress it into finitely many states need not end the investigation.

A small example: balanced parentheses

The grammar $S \rightarrow (S)S \mid \epsilon$ presents the language of balanced parentheses. To recognize a prefix, keep the number of unmatched opening parentheses, or a failure state if a closing parenthesis arrived too early. An opening parenthesis increments the count. A closing parenthesis decrements a positive count; at zero it enters failure. Failure persists. At the end, accept exactly when the count is zero.

The prefixes $($ and $(($ leave the same future task: any suffix that completes one completes the other. The prefix $(($ leaves a different task: $)$ completes the first two prefixes but not this one. These future tasks are **residuals**. There are infinitely many counts, yet a finite set of update rules generates their dynamics. This is a classical automata example, not a discovery of FPRD Lab. It illustrates the research question: which coordinates preserve all relevant futures, and what does updating or reading those coordinates cost?

From viewpoint to proof

In practice, the Lab fixes the observable and the allowed operations, compares presentations of the same problem, and seeks an exact translation or invariant. A failed shortcut is useful when a counterexample identifies the missing coordinate or hypothesis. Small computations test particular uncertainties; a proof must explain why the result holds beyond those tests. Costs, relation sides, boundary conditions, and the direction of each translation remain explicit.

AI-assisted exploration, source comparison, executable checks, and adversarial reconstruction support this work. Their agreement is not automatically independent evidence: shared assumptions must be exposed. A formal development certifies its stated formal theorem, not every interpretation attached to it. Each paper retains its own evidence and review status.

The approach is most useful when a poor representation hides a reusable construction, a compatibility condition, or a sharp obstruction. Outside recognition theory, it can guide the search for a structural reformulation without producing a literal residual dynamical system. The account below distinguishes that methodological influence from the ingredients of the proof.

Contribution to this publication

This addendum describes the current HTML reading edition. The preceding PDF pages retain their earlier edition; the afterword does not silently revise them.

FPRD's contribution to this report

The overview records the research-direction change that defines the Lab. A recursively generated family of distinct continuations need not be forced into a finite exact quotient. It can instead be studied as a structured infinite object, with named actions and finite observations. Joshua Gay's contribution to that turn, and the earlier PEG Lab work from which it arose, are identified in the report's provenance section.

The report organizes four layers: exact semantic behavior, observations of that behavior, dynamics generated by a finite presentation, and coordinate realizations with a uniform readout. Each comparison between layers is a mathematical obligation. A finite table at one input length, for example, does not establish the existence of one recognizer valid at every length.

This organization led to separate measurements of how many future behaviors can be distinguished, how far into a history their representatives first occur, and what it costs to maintain or retrieve them. The later access questions prevent a small observation count from being mistaken for an efficient online implementation. The report's sewing discussion likewise asks for compatible data across cuts rather than merely many individually correct finite views.

Inherited ingredients and the boundary of the contribution

Derivatives, Myhill-Nerode theory, inverse limits, guarded fixed points, Dyck languages, Hankel methods, and communication-complexity arguments have established sources. The report explicitly identifies overlaps with topological automata, refinement and learning methods, and functorial minimization. Their appearance in a common framework is not a new proof of ownership of those ideas.

Robson's globalization work influenced the research architecture: look for a faithful global representation in which awkward local data become manageable. No theorem asserts that the FPRD spaces are semialgebraic or that Robson's theorem transfers directly. That distinction is an example of crediting an intellectual influence separately from a formal dependency.

The Lab's contribution here is a presentation-sensitive, readout-aware synthesis, together with the particular arguments and quantitative questions whose status is stated in the report. It is an evolving mathematical program, not a universal theorem that every recognition problem has a cheap coordinate realization. The milestone labels organize dependencies; they should not be read as independent verification or as the later claim-maturity scale.

Where to inspect the contribution: the report, Sections 1-2, 6-8, and 10, especially the prior-lab boundary, source ownership, and novelty discussion.