

From Observable States to Acquirable Predictive Models

Brian Theory

Revised technical report, 7 October 2026

Abstract

A predictive representation can be defined from known dynamics yet remain impossible to acquire through the available observations. We examine this distinction for a finite hidden state on $P = \{a\} + M + \{z\}$, with M an antichain, and a supplied library of oriented binary queries. A query is admissible when its pre-answer order-complex inclusion is a homotopy equivalence. This is a formal constraint on sensing; no physical safety guarantee is assumed.

The structural result, consolidated from earlier research notes, characterizes every reportable target and gives a finest attainable transcript partition R . For known deterministic actions and output, let E_* identify states with the same output after every action word. Fixed-state acquisition of an exact predictive label is possible precisely when $R \subseteq E_*$. Standard congruence constructions then describe feasible predictive representations and the finest feasible coarsening of the original output labels. A reported class need not be one exact transcript leaf.

When queries must alternate with actions, finite-horizon support dynamic programming gives a different feasibility test. We state its common-protocol quantifier, query-first phase, branch-dependent actions and stopping rules explicitly. A reset can give current predictive certainty without revealing the initial state; a complementary invariant trap shows that fixed-state feasibility need not survive mandatory alternation. The report supplies self-contained proofs, worked examples and bounded computational corroboration. It establishes neither external priority for the specialized classification nor a new general minimization, planning or synchronization method.

1. The acquisition problem

Suppose the output and transition tables of a finite system are known, but its actual state is hidden. Computing which states have the same future behavior is an offline problem. Obtaining that behavior class from the observations available at runtime is a separate task. This report is for readers of finite-state systems and adaptive testing who want an exact example of that separation.

Our sensing constraint makes the orientation of a supplied query matter. On the three-state chain $a < x < z$, the tests 110 and 001 carry the same unrestricted binary information, but only the latter preserves the stipulated order-complex inclusion on $\{a, z\}$. Computing the complement of an answer already obtained cannot make an inadmissible physical query executable. Example 8 develops this running example with an output and a known transition.

Throughout, $P = \{a\} + M + \{z\}$ is the ordinal sum of two distinct endpoints and a finite antichain, possibly empty: $a < x < z$ for each $x \in M$. Every nonempty support has the order

induced from this specified P . Height counts strict inequalities, so the height is at most two. We do not claim a theorem for every poset of that height. The proofs start with this order and the supplied truth tables; a background observation model is motivation, not an additional hidden premise. Appendix A explains positive observation topology and complete signatures.

The central structural result is the endpoint-sensitive classification in Theorem 2 and its finest attainable partition in Corollary 3. These results are reproduced from the dated notes [1,2], with their local admissibility proof included. Ordinary predictive refinement and generated congruence then give Propositions 4 and 6 and the compatibility criterion. Corollary 7 is a short output-quotient consequence. Proposition 10 specializes standard support planning to an explicit alternating protocol. The purpose is a consolidated technical report with a complete operational contract. External priority of the specialized classification remains unresolved; writing out the later consequences does not establish a separate research increment.

We distinguish four objects: a reported label may combine several terminal transcripts; an exact leaf records the whole transcript; a current predictive class describes the state after any executed actions; an origin is the state before the experiment. Neither changing the requested label nor changing the state should be mistaken for learning a previously inaccessible distinction. Section 8 places these tasks beside equivalence-class determination, automata equivalence, contingent planning and adaptive homing. Appendix B records which results already occur in the source notes and which source identities remain unconfirmed.

2. Fixed-state experiments and local admissibility

Let Q be an explicitly supplied finite set of deterministic maps $q : P \rightarrow \{0,1\}$. Fixed-state histories contain query-answer pairs, and their support is

$$S_h = \{x \in P : q(x) = d \text{ for every recorded } (q, d)\}.$$

There is no noise, query-dependent state change, shared consumption across counterfactual branches, or additional history-dependent restriction. All legality tests precede answer conditioning. Following the formal model [1], a query on nonempty support S refines the order by

$$x \leq_q y \iff x \leq_P y \text{ and } q(x) \leq q(y).$$

It is legal if the inclusion of order complexes $\Delta((P|S)_q) \hookrightarrow \Delta(P|S)$ is a homotopy equivalence. The order complex has chains as simplices. Legality concerns this inclusion, not an isomorphism of finite posets. We retain the source abbreviation RF for this formal admissibility rule. The term “repair-safe” in the notes does not establish physical safety. The test is applied before the answer: conditioning itself may change topology. For example, on $\{a\} + \{x, y\} + \{z\}$, query 0110 is legal, but its answer-one support is the disconnected antichain $\{x, y\}$, whereas the original order complex is contractible.

After an answer, all remaining states have the same value of the query just asked. Therefore that query imposes no extra order restriction within its answer support. The same holds for every previously recorded query. This is branch cancellation: along a fixed-state branch, the current order is precisely the original induced order on its support. Repeated queries are constant there and add no information. Constants are legal, but can be removed from useful fixed-state decision trees.

Lemma 1 (local RF admissibility; [2, P1-L1]). If both endpoints are present, legality is equivalent to $q(a) \leq q(z)$. If only a is present, legality means $q(a) = 0$ or the restriction is constant one. If only z is present, legality means $q(z) = 1$ or the restriction is constant zero. If

neither endpoint is present, every query is legal. These rules include singleton supports and the two-endpoint support without middle states.

Here is the complete local argument. With both endpoints, the only triangles are axz . Endpoint values 01 delete no simplex. For endpoint values 00, each middle state valued one loses its edge xz and triangle axz . That edge is a free face of its unique triangle, so deleting these pairs gives elementary collapses onto the refined complex. Endpoint values 11 give the dual collapses, deleting ax for middle states valued zero. These explicit collapses prove homotopy equivalence of the inclusion in the safe cases.

For endpoint values 10, let $k = |S \cap M|$. There are $k + 1$ deleted edges, including az , and k deleted triangles. In relative chains over $\mathbb{F}_2$, each triangle boundary is $az + e_x$, where e_x is its unique deleted spoke. The columns are independent because the spokes are distinct. There are no relative vertices, so relative first homology has dimension one. A homotopy equivalence would induce zero relative homology, which is impossible. This uses homology only to obstruct equivalence, not to infer equivalence from acyclicity.

On a one-endpoint support, the order complex is a star. Any descending spoke is deleted without a triangle and disconnects the complex; otherwise nothing is deleted. On a middle-only antichain there are no edges to delete. This proves all four cases. In this particular family, legality is also equivalent to preserving the number of connected components. With endpoint word 10, the refined complex consists of two disjoint nonempty stars; with one endpoint, a deleted descending spoke isolates its middle vertex. The legal cases have the explicit collapses or unchanged complexes above. This component criterion is a consequence of these cases, not a criterion for homotopy equivalence of general complexes.

Define $Q_s = \{q \in Q : q(a) \leq q(z)\}$. An observation outside Q_s has endpoint word 10. It cannot be both useful and legal on any support containing an endpoint: with both it is illegal, and with one it is legal only when constant. Such a query can nevertheless be useful on a middle-only branch. This asymmetry is the mechanism behind the classification criterion. Membership in Q_s means legality on the full carrier and on supports containing both endpoints. It does not ensure legality on every smaller support: 010 is legal on $a < x < z$ but illegal on $\{x, z\}$. Each support must pass its own test.

3. Reportable targets and the finest transcript partition

Theorem 2 (target classification; [2, P1-T1]). Starting with support P , for every finite-valued target $c : P \rightarrow C$, a finite legal supplied-query tree reports $c(x)$ for every hidden state if and only if two conditions hold: full Q separates every differently labelled pair of middle states, and Q_s separates every differently labelled endpoint-state pair. Equal endpoint labels are permitted. In particular, the theorem does not require an ascending endpoint query when the target does not need one.

Necessity follows by following transcripts. Two states with identical full signatures receive identical answers to every supplied query. For an endpoint and another state with different required labels, consider their first separating query in a successful tree. Before that answer, both states remain in the support. The query is useful and legal there, so the preceding lemma puts it in Q_s . Separation by an unavailable or inadmissible query cannot substitute for this first legal separator.

For sufficiency, divide Q_s into $Q_=_$, the queries with equal endpoint values, and $Q_\uparrow$, those with endpoint values 01. While the branch containing both endpoints has a nonconstant $Q_=_$ -query, ask it. The answer opposite to the common endpoint value detaches a middle-only branch; the other answer retains the endpoints. On a middle antichain every query is legal. Full-signature

classes there are target-homogeneous by the first condition, so querying separating members of Q classifies every detached branch.

After these peels, all remaining middle states agree with both endpoints on every $Q_{=}$ -query. If $Q_{\uparrow}$ is empty, the second condition forces this entire core to share one target label. If the core is already homogeneous, stop regardless of $Q_{\uparrow}$. Otherwise ask any member of $Q_{\uparrow}$. Its zero child is a lower star with a , and its one child is an upper star with z .

On the lower star, a state labelled differently from a has a Q_s -separator from a . It cannot be in $Q_{=}$, after the completed peels, so it belongs to $Q_{\uparrow}$, has value zero at a , and value one at that state. It is legal throughout the lower star and detaches a middle-only branch. Repeating removes every incorrectly labelled state. The upper star is dual, using queries with value one at z and zero at the state being detached. All detached branches are classified as above. Each informative query strictly shrinks both children, so this finite construction terminates.

Corollary 3 (acquisition equivalence and exact leaves; [2, P1-T2]). The pair conditions compress into a single equivalence. Let K partition states by their Q_s -signatures. Keep each K -class containing an endpoint intact, and refine each endpoint-free K -class by full Q -signatures. Call the resulting equivalence R . Equivalently, take the equivalence closure of endpoint-state pairs sharing their root-admissible signatures and middle-middle pairs sharing their full signatures. Every generator lies in one K -class. If that class contains an endpoint e , its endpoint-state generators join every member to e ; hence the entire class stays together, including middle states with distinct full signatures. If it contains no endpoint, only full-signature equality generators occur, and no endpoint path can enter it. Its equivalence classes are therefore exactly the full-signature fibers. This proves that the two descriptions of R agree.

The theorem becomes

$$c \text{ is RF-reportable} \iff R \subseteq \ker(c).$$

Here and throughout, inclusion means inclusion of equivalent pairs: every R -class lies inside one target fiber. This convention avoids reversing the comparison between finer partitions and smaller equivalence relations.

There is also a legal tree whose exact transcript leaves are precisely the R -classes. Apply the classification theorem to the R -class label. Every legal tree must keep each generating pair, and hence each R -class, together; otherwise its own leaf labels would violate necessity. The constructed classifier must also separate different R -classes. Thus R is the finest attainable transcript partition. Its coarsenings can all be reported, but need not be exact transcript partitions. A reduced fixed-state tree has no constant internal query, so every internal node has two nonempty children. An exact R tree consequently has $|P/R|$ leaves and $|P/R| - 1$ internal nodes. Its depth is at most $|P/R| - 1$; this is an existence bound, not an optimal-depth algorithm.

The formula for R concerns acquisition starting from all of P . It cannot in general be restricted to a fresh smaller support: for $Q = \{101\}$ on $a < x < z$, global $R = \{a, z\} \mid \{x\}$, but on fresh $\{a, x\}$ the only query is descending and illegal, although the restricted relation is discrete. A new starting support is a new admissibility problem.

4. Predictive equivalence and feasible quotients

Fix a finite action set U , known deterministic maps $T_u : P \rightarrow P$, and a finite-valued output b . In fixed-state acquisition these tables define the prediction problem; actions are not executed during sensing. For a word $w = u_1 \cdots u_k$, set $T_w = T_{u_k} \circ \cdots \circ T_{u_1}$; the empty word acts as the identity. The predictive equivalence is

$$xE_*y \iff b(T_w x) = b(T_w y) \text{ for every finite } w.$$

No value of b is observed for free. If its measurement is needed, an appropriate query must actually belong to the library and be legal.

Proposition 4 (standard predictive refinement; [2, P1-T3; 4]). Compute this equivalence by ordinary finite refinement:

$$E_0 = \ker(b),$$

$$E_{k+1} = E_k \cap \bigcap_{u \in U} (T_u \times T_u)^{-1}(E_k).$$

Induction shows that E_k means agreement after every word of length at most k . The induction step includes the empty word through E_k , and longer words through their first action and remaining suffix. Each strict refinement increases the number of classes, so there are at most $|P| - |b(P)|$ strict refinements. Once consecutive relations agree, the recurrence remains fixed, giving E_* .

A transition congruence is an equivalence F satisfying $xFy \Rightarrow T_u x F T_u y$ for every action. The fixed relation E_* is such a congruence and lies inside $\ker(b)$. Every congruence inside $\ker(b)$ lies inside every E_k , by induction, hence inside E_* . Therefore E_* is the largest such relation, or the coarsest exact predictive partition. Output and transition maps on its classes are well-defined. Knowing the current class suffices to update it and predict outputs indefinitely without additional observations.

Corollary 5 (fixed-state compatibility; [2, P1-T3]). RF-admissible acquisition reports E_* exactly when $R \subseteq E_*$, by classification. More generally, suppose some RF-reportable transition congruence $F \subseteq \ker(b)$ is acquired. Reportability implies $R \subseteq F$, while predictive maximality implies $F \subseteq E_*$. Hence compatibility is necessary. Conversely, when compatibility holds, E_* itself is a reportable predictive representation.

This also explains a failed repair strategy. If xRy but not xE_*y , every finer predictive equivalence still separates these states. No additional subdivision of the desired labels can make that blocked pair observable. Failure can be certified by the pair together with a finite word w for which the two outputs differ. Such a word exists by the characterization of the refinement sequence; tracing a separating refinement step supplies one.

Proposition 6 (generated transition congruence; [2; 10]). Start from acquisition rather than output:

$$G_0 = R,$$

$$G_{k+1} = \text{Eq} \left(G_k \cup \bigcup_{u \in U} (T_u \times T_u)(G_k) \right).$$

Here Eq denotes equivalence closure. The increasing sequence stabilizes at the least transition congruence G containing R . Indeed a fixed point is forward invariant, and every congruence containing R contains every iterate. There are at most $|P/R| - 1$ strict class mergers.

Compatibility is equivalently $G \subseteq \ker(b)$. If $R \subseteq E_*$, forward invariance of E_* contains every closure iterate. Conversely, if $G \subseteq \ker(b)$, maximality puts $G \subseteq E_*$. When compatible, the reportable predictive equivalences are precisely the transition congruences satisfying $G \subseteq F \subseteq E_*$. Thus G is the finest and E_* the coarsest reported representation in this class. Arbitrary intermediate equivalences need not be congruences. These are standard refinement and closure operations used with P1's acquisition obstruction, not new automata algorithms.

Corollary 7 (finest feasible output coarsening). If exact prediction of the original output is infeasible, relabeling inaccessible states cannot restore it. A different repair is available only if the prediction requirement itself may be weakened. Let $B = b(P)$, let G be the least transition congruence containing R , and let D be the least equivalence on B containing every pair $(b(x), b(y))$

with xGy . Write $\pi : B \rightarrow B/D$ for the quotient map and define $b' = \pi \circ b$. Then $G \subseteq \ker(b')$ by construction. Since G is a transition congruence, it is contained in the predictive equivalence E'_* for b' . Therefore $R \subseteq G \subseteq E'_*$: the coarsened future outputs are safely acquirable and predictable.

This is the finest feasible coarsening of the specified output labels, by the following elementary quotient argument. Let $f : B \rightarrow C$ be any relabeling whose future outputs are safely acquirable. Its predictive equivalence E_f is a transition congruence containing R , so minimality gives $G \subseteq E_f$. The empty action word gives $f(b(x)) = f(b(y))$ whenever xGy . Every generating pair of D thus lies in $\ker(f)$, and equivalence closure gives $D \subseteq \ker(f)$. Consequently f factors through π . Conversely, a relabeling factoring through π is constant on every G -class, so the same congruence argument proves feasibility. This statement concerns output-label coarsening, not optimal protocol depth. It discards the original output distinctions merged by D , and does not establish exact prediction of that unchanged output. Changing the query library is a separate intervention: it can change R , whereas this construction fixes the library and transitions.

5. A running example: orientation and timing

Example 8 (three-state obstruction and controls; [2, P1-C1]). This is a smallest obstruction under the conditional assumptions stated below. On $a < x < z$, write binary tables in that order and take

$$Q = \{010, 110\}, \quad b = 010, \quad T = (x, x, z).$$

The map is order preserving. The second supplied query is already the pullback $b \circ T = 110$; this is not an example where the needed truth table was omitted. The base-output partition is $\{a, z\} \mid \{x\}$, but a and z have different outputs after one action. Consequently E_* is discrete.

Only 010 belongs to Q_s , so $R = \{a, z\} \mid \{x\}$. The base query is legal initially, with endpoint values 00. Its zero answer leaves $\{a, z\}$, where the other query is the descending test 10, hence illegal. Asking that other query first is also illegal. Repetition yields no new information. Thus RF-admissible acquisition of the predictive class fails, although the base output is safely obtained in one query. Ignoring legality, ask 010, then 110 on the zero branch: ordinary identification succeeds in depth two.

Actually supplying the complement 001 in place of 110 repairs this example. Ask 010, then 001 on $\{a, z\}$; the latter is ascending and legal. The repaired tree identifies the state in depth two. The two libraries induce the same unrestricted Boolean information, but not the same permissible physical experiments. Computationally flipping an answer from 110 would first require performing that illegal query; it is not a substitute for supplying 001.

The conditional minimality argument is elementary. With two states, a nonconstant base output is already injective, so it cannot have a strict predictive refinement. A constant output remains constant after every word. At least three states are therefore required. If the base is nonconstant and safely obtainable, while ordinary acquisition can obtain its strict predictive refinement, one binary query is insufficient: every obtainable nonconstant base must distinguish its two fibers, and repetitions cannot split either fiber. At least two distinct supplied binary query maps are necessary under those conditions. This example attains that library-size lower bound and the three-state lower bound; its exhibited worst-case query depth is two. This says nothing about other sensing or control contracts.

Two timing controls are also inherited from [2]. In the monotone example above, passive observations of b at times zero and one produce histories 01, 11, 00 from a, x, z , respectively. They retrospectively distinguish initial states only after the second observation. Under the moving-state rule used below, however, the initial zero branch $\{a, z\}$ maps to $\{x, z\}$. On that image support,

b is descending and illegal. The passive-history success therefore does not implement a safe two-sample experiment.

For the swap control, use $Q = \{010, 100\}$, keep $b = 010$, and let T fix a while swapping x, z . Again Q_s contains only the base query, R joins a, z , and the predictive equivalence is discrete, so fixed-state acquisition fails. But query b , execute T , and query b again. The zero branch becomes $\{a, x\}$, where the base query is ascending; the one branch becomes a singleton. Histories 00, 10, 01 distinguish initial states, and therefore also current states because this action is bijective. The report is after two queries and one action, not at time zero. These controls motivate the changed interaction model of Section 7.

The unsuccessful approaches are now explicit: ignore orientation, subdivide blocked labels, treat an unsupplied complement as available, or postpone an observation without checking its new support. Each changes or violates an assumption rather than overcoming the fixed-state theorem.

6. A reported class need not be one transcript leaf

Example 9 (no coarsest exact-leaf refinement; [2, P1-C2]). Let $P = \{a\} + \{x, y\} + \{z\}$, with tables ordered as (a, x, y, z) , identity dynamics, and desired partition

$$\mathcal{B} = \{a, x, z\} \mid \{y\}.$$

Supply 0001, 0011, 0111, 0101. Every root query has endpoint word 01, so every root query is legal. Their full signatures distinguish all four states, giving discrete R . Identity dynamics makes $E_* = \mathcal{B}$ when b is the class label of $\mathcal{B}$. Compatibility therefore guarantees reportability.

One legal tree asks 0001, then 0011 on $\{a, x, y\}$. Its exact leaves are

$$\Pi_1 = \{a, x\} \mid \{y\} \mid \{z\}.$$

The second test is legal on the lower star because its value at a is zero. Another legal tree asks 0111, then 0101 on $\{x, y, z\}$, producing

$$\Pi_2 = \{a\} \mid \{y\} \mid \{x, z\}.$$

The second test is legal on the upper star because its value at z is one. Both trees report $\mathcal{B}$ by assigning the same reported label to multiple leaves.

No tree has exactly $\mathcal{B}$ as its transcript partition. The target is nonconstant, so a successful tree must ask a root query. Every supplied query splits $\{a, x, z\}$; once states have different recorded answers, later queries cannot merge their transcripts. Moreover, any common coarsening of Π_1, Π_2 that still refines $\mathcal{B}$ must join a with x and x with z , hence must equal $\mathcal{B}$. Since $\mathcal{B}$ is unattainable as exact leaves, there is no coarsest attainable exact-leaf refinement of this target.

The distinction is substantive rather than terminological. A report may discard acquired information, whereas an exact transcript records it. The finest attainable transcript partition R still exists, and the coarsest predictive reported partition still exists here. Neither fact provides a coarsest feasible transcript refinement of every target. Any implementability statement must specify which object it minimizes.

7. Alternating observations and actions

We now change the acquisition contract explicitly. The same finite carrier, supplied queries, output, and known deterministic actions are used, but actions may change the hidden state

during acquisition. A persistent history records every chosen query, observed bit, and chosen action. Queries read the current pre-action state without changing it. The next action may depend on that answer. After the action, the order is reset to the original order induced on the image support. Actions need not be injective or order preserving; no additional action-safety condition is imposed. These are modeling assumptions, not consequences of [1]. The syntax is strict: no action-only start and no consecutive queries without an intervening action. An identity action, if supplied, can embed a fixed-state query tree in this syntax; no identity action is assumed.

At a pre-query support S , a legal query q gives the nonempty answer support $S_d = \{x \in S : q(x) = d\}$. If action u is then chosen, the next current support is $T_u(S_d)$, with duplicate images identified. The RF test is applied on S before learning d , not separately on hypothetical favorable outcomes. There is no free output observation, complemented query, pullback oracle, or transition learning. A constant query is legal and is allowed to precede an action even though it yields no information.

Reporting targets the current E_* -class, computed from the same known future dynamics and output. A support is terminal, denoted $\mathcal{H}(S)$, when it is contained in a single E_* -class. After a correct report, all subsequent output predictions must be produced by quotient updates without further sensing. This target does not require identifying the initial state or recovering distinctions that actions have erased.

Proposition 10 (finite-horizon support planning). For each nonempty support S and integer $h \geq 0$, let $W_h(S)$ mean that there exists one deterministic legal history-dependent protocol, starting at a pre-query checkpoint with current support S , that reports the correct current E_* -class for every possible state in S . On each executed branch it uses at most h actions and at most $h + 1$ queries. It alternates query then action and may stop before a query, immediately after its answer, or immediately after an action. Then

$$W_0(S) = \mathcal{H}(S) \vee \\ \exists q \in Q \text{ legal on } S \quad \forall d \in \{0, 1\} \text{ with } S_d \neq \emptyset : \mathcal{H}(S_d),$$

and, for $h \geq 1$,

$$W_h(S) = \mathcal{H}(S) \vee \\ \exists q \in Q \text{ legal on } S \quad \forall d \in \{0, 1\} \text{ with } S_d \neq \emptyset : \\ [\mathcal{H}(S_d) \vee \exists u \in U W_{h-1}(T_u(S_d))].$$

This is finite-horizon AND-OR planning on supports [9], with RF deciding which query choices are enabled. The formula permits different actions, or a decision to stop, on different answer branches. All existentially chosen queries belong to Q . Empty answer branches impose no obligation.

Proof proceeds by induction on the action allowance. With none available, a protocol either reports immediately or asks one legal query and reports after every possible answer. This is exactly the first equation. With positive allowance, an immediate report again requires homogeneity. Otherwise the first query must be legal on the full current support. After its answer, the protocol either reports a homogeneous class or chooses an action and continues with one fewer action and at most one fewer query available. Known determinism gives exactly the image support appearing in the equation. Inductive necessity follows by considering every possible answer.

Conversely, choose witnesses to the displayed recurrence. A homogeneous branch receives its unique predictive label. On every other branch execute its witness action and attach the finite protocol supplied by induction. Legality holds before the first answer by construction and thereafter by the inductive protocols on the reset image orders. The horizon decreases after

each action; stopping after the final query is already covered by W_0 . This produces a finite legal experiment with the stated bounds and proves sufficiency.

Non-injective actions cause no exception. Every image state has at least one predecessor in the conditioned support, and every such predecessor can generate the recorded branch. Coalescing predecessors therefore preserves the exact set of possible current states. At pre-query checkpoints, the current support and remaining action allowance determine whether a successful continuation exists: the available queries, actions, induced order and predictive target are then identical. A controller can maintain this support from its recorded history and choose branch-specific witnesses; this is not a memoryless strategy on the hidden physical state. At a post-answer checkpoint, the next permissible instruction is instead an action or a report. This sufficiency depends on the absence of consumption constraints or extra history-dependent syntax. It would not justify discarding memory in a different model.

A separate relation $\mathcal{I}$ tracks initial reconstruction. Start with $\mathcal{I} = \{(x, x) : x \in P\}$. A query keeps pairs whose second coordinate has the received value; an action replaces (x_0, x) by $(x_0, T_u x)$. Induction identifies $\mathcal{I}$ with precisely the initial-current pairs consistent with the history. Its second projection gives current support, while its first projection gives possible initial states. This explains exactly what information a support-only description omits.

Example 11 (reset, current certainty and origins). Consider the following explicitly refuted transfer claim: moving-state prediction is possible only if the original fixed-state inclusion $R \subseteq E_*$ holds. This is a claim tested here, not an assertion attributed to a publication. On $a < x < z$, supply just $q = 000$, let $b = 010$, and supply one action T sending every state to a . Every nonempty action word ends at a , so $E_* = \ker(b) = \{a, z\} \mid \{x\}$. But the constant library gives universal R , violating the fixed-state criterion.

Ask the constant query once, receive zero, execute the reset once, and report immediately after that action. The query is legal, the image support is $\{a\}$, and the current predictive class is known. This uses exactly one query and one action, with report time one measured in executed actions and no post-action query. Every initial state produces the same entire record. The pair relation is now $\{(a, a), (x, a), (z, a)\}$, so initial identification fails. Indeed no later experiment can distinguish those origins: all subsequent states and query answers coincide under the same history-dependent choices. The output is nonconstant on the original carrier, making the separation nontrivial. This proves the negative transfer result while explaining its mechanism: control can acquire current predictive certainty by changing the state, rather than by learning its past.

Example 12 (fixed-feasible but alternating-impossible). Conversely, on $a < x < z$, take $Q = \{010, 011\}$, $b = 010$, and the sole action $T = (x, x, z)$. Both queries lie in Q_s and their signatures separate all states; a, z differ after one action and the other pairs differ already under b . Thus $R = E_*$ is discrete, and fixed-state acquisition succeeds.

No alternating protocol succeeds at any finite horizon. If the first query is 010, its zero branch $\{a, z\}$ is not terminal and the required action takes it to $\{x, z\}$. If the first query is 011, its one branch is already $\{x, z\}$, is not terminal, and the required action fixes it. On this support, 010 is illegal, 011 is constant, and T is the identity. The support contains two predictive classes forever. If no action allowance remains, the nonterminal branch fails immediately; with any positive allowance it remains trapped. Hence $W_h(P)$ is false for every finite h . This invariant argument, not a finite extrapolation, proves the claim. Together with Example 11 it shows that neither direction of fixed-state compatibility transfers under the specified mandatory alternation.

Even winning supports need not be downward closed. With $Q = \{010\}$, $b = 010$, and no actions on the three-state chain, $W_0(P)$ holds but $W_0(\{x, z\})$ fails because its sole informative query is illegal. Support planning still applies; an optimization that assumes every subset of a winning support is winning would not be justified.

8. Relation to earlier work

The terminal target in the fixed-state problem is equivalence-class determination: an adaptive test policy may stop when all consistent hypotheses belong to one required class [7, §3, expanded version p4]. Decision-region determination generalizes this condition to overlapping regions [8, §2, Eq.(2), p432]. Thus reporting a label need not identify the hidden state or make the terminal support equal the whole label fiber. The additional problem here is that the enabled physical query depends on the entire current support through an oriented topological inclusion. The RF classification and the exact realization of its signature partition address that constraint; the expected-cost algorithms of [7,8] are not imported.

The predictive equivalence is ordinary finite deterministic behavioral equivalence. Hopcroft's introductory account describes repeated output partition refinement by successor blocks [4, p1], and the formal word-equivalence definition appears on p2. Here we use the routine finite-output (Moore) version; the formal automaton in [4] has a binary final-state output. Treating each action as a unary operation also makes G a standard generated congruence [10, II§5]. The output-label result is the elementary quotient property applied after this closure. These constructions organize the consequences of the acquisition obstruction; they are not new minimization or closure algorithms.

The moving-state recursion is finite-horizon AND–OR dynamic programming on exact belief supports, with RF legality as its enabled-query predicate. Bonet and Geffner give belief images, observation conditioning and a worst-case Bellman formulation [9, §2, Eqs.(1),(5),(6)]. The present contract fixes a separate query-first phase, answer-dependent actions, a remaining horizon and the current predictive-class goal. Support-dependent admissibility need not survive taking subsets, so hereditary-support optimizations require additional proof. The source also explicitly distinguishes setting a known current value from recovering its initial value [9, p54, footnote3].

Adaptive homing determines the state after an experiment; the finite adaptive-test definition in [6, §2, p75] makes this timing explicit. The target here is a current predictive class, with separate pre-action queries and RF constraints. The constant reset example is synchronization: one action maps all states to a common state [11, §1]. Its value here is to expose an invalid transfer of the fixed-state criterion, not to introduce synchronization.

The order-complex and elementary-collapse methods are standard [12, §2]; the particular binary RF classification is specialized to the two-ended antichain family. The positive-topology convention uses upward sets, whereas [12] uses the dual finite-space order. Finally, computational mechanics groups stochastic histories by conditional future laws [5, Def5]. Its predictive-sufficiency and minimal-complexity statements [5, Thms1–2] concern a different stochastic and almost-sure setting; they neither identify the present controlled quotient with causal states nor establish lawful acquisition of it.

The fixed-state theorems and examples are reproduced from the matching dated P1 note [2], with the local RF argument inherited there from the RF foundation. The manuscript consolidates those arguments and adds an elementary output quotient and an explicit moving-state specialization. These are provenance statements about the inspected snapshots, not established publication histories. The broader priority and significance of the RF classification remain open.

9. Reproducibility, limits and conclusions

The structural content is the classification of admissible fixed-state information on the specified two-ended antichain poset. Corollary 3 proves that all reportable targets are coarsenings of an attainable partition R . Standard predictive equivalence describes which distinctions matter for

every future action word; comparing the two relations gives the exact acquisition criterion. A report may discard transcript distinctions, so it is essential to specify whether a desired partition is a label or an exact leaf partition.

The accompanying verification material contains the independent audit’s retained checker, actual environment, stdout and results. That prior run reported 5,332 direct-complex legality cases for all supports and binary queries at 2–6 states; 984,352 distinct full-support library/target classification cases through four states; and 754,080 bounded moving-policy comparisons at 2–3 states, all query libraries, zero or one arbitrary action, every output partition and support, and horizons 0–2. These are separate metrics with different ranges, not an additive total. Larger moving checks were seeded samples. The legality path uses actual collapses or nonzero relative homology; the moving comparator explicitly tracks histories, phases, budgets and initial-current pairs independently of the support recurrence. Full ranges and implementation limits are recorded in the accompanying verification report. These retained audit results are not relabeled as a new full revision run. Targeted revision checks and their actual logs are recorded separately. Finite checks corroborate the arguments and do not prove the unbounded theorems.

Known actions change the state whose class is sought. Synchronization can remove current uncertainty while preserving uncertainty about every origin. Mandatory query/action alternation can also obstruct an experiment that succeeds with a fixed state. Proposition 10 handles both effects by checking legal queries before answers, selecting actions on each answer branch, and using the stipulated induced order on each image support. Its correctness is ordinary finite-horizon support reasoning; its operational content depends on this precise contract.

The report does not justify RF as a physical safety condition. Within this poset family its local homotopy test reduces to preservation of component number, and neither answer conditioning nor actions are required to preserve topology. Noise, unknown transitions, consumed queries, restricted memory, action admissibility and sensing costs would require changed models and proofs. No general-poset classification, optimal-depth formula, new automata minimization bound or general infinite-horizon characterization is established.

The consolidated account makes the fixed-state theorem, standard consequences and timing counterexamples independently checkable. The source notes already contain most of the structural mathematics. Broader priority and research significance of that classification remain open, as do the source-history details specified in Appendix B. The appropriate present claim is a self-contained technical report, not evidence of a newly established research contribution merely from clearer exposition.

Appendix A. Positive observations and full signatures

The distinction between positive observation topology and complete signatures is important. The observation-topology manuscript [3, v0.9, §3] defines the topology generated by positive truth regions of background observations Φ . On our finite carrier, write $o_\Phi(x) = (\varphi(x))_{\varphi \in \Phi}$, with binary coordinates ordered by $0 < 1$. The signature representation identifies opens with inverse images of upward-closed subsets of realized signatures. This construction records oriented positive information, not unrestricted Boolean access to all coordinate values at the hidden state.

For completeness, finiteness makes the signature description direct. For a realized signature s , select, for each realized signature not above s , a coordinate witnessing that failure. There are finitely many such signatures. Intersecting the corresponding positive coordinate regions selects exactly those realized signatures above s . Unions of these principal regions give all upward-closed sets. Conversely, each positive region is upward closed, and this property survives finite intersections and arbitrary unions. This argument is confined to the finite carrier used here.

Complete-signature equivalence instead identifies x, y when all coordinates agree. Its cells belong to the Boolean algebra generated by the observations, using negative as well as positive information. For a two-state carrier $\{p, r\}$ with one coordinate 01, the positive topology is $\{\emptyset, \{r\}, \{p, r\}\}$, whereas the complete-signature partition is discrete. The Boolean algebra contains $\{p\}$, although that singleton is not positive open. Accordingly, distinguishability by complete signatures is not identical to availability of an oriented query.

Background signatures may be assumed injective as in [1], but their actual values at the hidden state are not supplied to the agent. The runtime library Q can be smaller and need not separate states. Known truth tables permit offline computation; they are not answers. Negating an acquired bit is postprocessing, not permission to perform an unsupplied complemented query. These observations explain the motivation for the stipulated order; the classification proofs do not require reconstructing it from background observations.

Appendix B. Provenance and result overlap

The matching local snapshots of [1]–[3] and the RF foundation note were recovered and hash-checked in the independent audit. They are preserved unchanged in the accompanying materials. Reference [1] supplies the fixed-state contract. Reference [2], P1-L1, contains Lemma 1; P1-T1 contains Theorem 2; P1-T2 contains Corollary 3; P1-T3 and its dual-closure section contain the predictive refinement, compatibility and generated-congruence results. P1-C1 contains Example 8 and its timing controls, and P1-C2 contains Example 9. Appendix A draws on [3, v0.9, §3]. Proofs needed here are reproduced rather than delegated to inaccessible premises.

The output-label quotient and formal moving-support proposition are additions relative to that inspected P1 version, but use standard quotient and planning principles. The reset is a synchronization illustration. Example 12 and the component-count observation were added during this revision after proof rechecking. The previous audit and the revision’s specialist checks are AI assessments, not external human peer review.

Confirmed public identifiers and authorship for [1] and [2], permanent locators for the unpublished materials, and their dissemination/submission histories remain unresolved unless supplied in the accompanying provenance record. A local hash identifies bytes; it does not authenticate priority or establish public availability. No duplicate publication is inferred from overlap with private notes. No public deposit is made by this local revision.

References

- [1] *Formal adaptive model*. Research note, 29 September 2026. Author and permanent public locator unconfirmed. The preserved RFM snapshot has SHA-256 268672674c7a1ebcfe55cab53fc1b62704eac04284d7072e2fcd852b04452b17.
- [2] *Predictive refinement under a fixed repair-safe query library*. Research note, 4 October 2026. Author and permanent public locator unconfirmed. The preserved P1 snapshot has SHA-256 7ce880f1dfef53ad16f6763d9aee555de9a9098d5fdf057c127d243aa9efee2f.
- [3] Brian Theory. *Observation Topologies for Finite Propositional Semantics: Distinguishability, Distance, and Inference*. Manuscript v0.9, September 2026. Permanent public locator unconfirmed; preserved TeX snapshot identifies the version.
- [4] John Hopcroft. *An $n \log n$ Algorithm for Minimizing States in a Finite Automaton*. Stanford Computer Science Technical Report STAN-CS-71-190, January 1971. Original scan. Relevant

printed pp1–2/PDF pp5–6.

[5] Cosma Rohilla Shalizi and James P. Crutchfield. *Computational Mechanics: Pattern and Prediction, Structure and Simplicity*. arXiv:cond-mat/9907176v2, 19 June 2000; Santa Fe Institute Working Paper 99-07-044. Versioned record, PDF. This is the cited preprint edition, not journal pagination.

[6] Natalia Kushik and Nina Yevtushenko. “Adaptive Homing is in P.” *Electronic Proceedings in Theoretical Computer Science* 180 (2015), 73–78. DOI:10.4204/EPTCS.180.5, paper. Relevant §2, printed75/PDF3.

[7] Daniel Golovin, Andreas Krause and Debajyoti Ray. “Near-Optimal Bayesian Active Learning with Noisy Observations.” *Advances in Neural Information Processing Systems* 23 (2010). Official record. Citations to §§2–3, pp2,4 refer to the expanded arXiv:1010.3091v2, 16 December 2013.

[8] Shervin Javdani, Yuxin Chen, Amin Karbasi, Andreas Krause, Drew Bagnell and Siddhartha Srinivasa. “Near Optimal Bayesian Active Learning for Decision Making.” *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, PMLR33 (2014), 430–438. Official record, paper. Relevant §2, Eq.(2), printed432/PDF3.

[9] Blai Bonet and Héctor Geffner. “Planning with Incomplete Information as Heuristic Search in Belief Space.” *Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling (AIPS2000)*, 52–61, 2000. Proceedings paper, author copy. Cited locations use the proceedings pagination: §2, printed53–54/PDF2–3.

[10] Stanley Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Graduate Texts in Mathematics78, Springer, 1981. Author-hosted Millennium edition. Cited numbering and pages refer to this edition, II§§5–6, printed38–41,50–51/PDF54–57,66–67.

[11] D. S. Ananichev, V. V. Gusev and M. V. Volkov. *Slowly synchronizing automata and digraphs*. arXiv:1005.0129v1, 2 May 2010. Versioned PDF. Relevant reset definition: §1, PDF1.

[12] Jonathan A. Barmak and Elias G. Minian. “Simple homotopy types and finite spaces.” *Advances in Mathematics* 218(1) (2008), 87–104. DOI:10.1016/j.aim.2007.11.019, author manuscript. The cited §2 and printed3–7/PDF5–9 locators refer to the author manuscript.