

An intrinsic interval-deletion characterization of digraph support families

Abstract

Let E be a finite set. A directed graph on E , with loops permitted, *supports* the subsets of E in which every vertex has an in-neighbour inside the subset. The supported subsets form a union-closed family containing the empty set, and Steel (*Acta Biotheoretica*, 2023) asked for a set-theoretic characterization of the families that arise in this way. We give a recursive characterization that refers only to the given family $\mathcal{F}$. Starting from the full power set, one repeatedly selects an inclusion-maximal set $S \notin \mathcal{F}$, chooses a root $r \in S$ that lies in no member of $\mathcal{F}$ contained in S and has not been used before, and deletes the interval $\{X : r \in X \subseteq S\}$. The family $\mathcal{F}$ is the support family of a digraph on E if and only if some legal sequence of deletions ends exactly at $\mathcal{F}$. The residual set may be chosen by any fixed rule; only the root requires branching. Every sequence has at most $|E|$ steps, and a successful sequence reconstructs a realizing digraph on the original ground set by the predecessor rows $P(r_i) = E \setminus S_i$. The proof passes through complementation to single-head definite Horn theories, where the key step is a normalization lemma replacing one rule of a hypothetical single-head completion while preserving its models and all previously accepted rules. We derive the exact root set and length of every successful sequence, a finite recognizer whose rejections are proofs of non-realizability, membership in NP for explicitly listed families, closure of the class under coordinate projection (elimination of hidden vertices), induced deletion and disjoint products, a Hall-type necessary condition, and a three-element family showing that the class is not closed under intersection on a shared ground set, so that no collection of closure axioms can characterize it. For catalytic reaction systems we obtain that a family is the fixed family of an elementary system if and only if it passes the interval-deletion test, and that same-ground RAF families are exactly intersections of an antimatroid with a family passing that test. The characterization, its consequences and an executable trace checker are verified in Lean 4.

Keywords. union-closed families; directed graphs; interior operators; definite Horn theories; single-head formulas; autocatalytic sets; RAF theory; formal verification.

MSC 2020. 05C20, 06A15, 05D05, 68R05, 68T27, 92E20, 68V20.

1 Introduction

1.1 The question

Let E be a finite set and let D be a directed graph on the vertex set E ; loops are permitted. Following Steel [13, Section 2.1], the *support family* of D is

$$\mathcal{C}(D) = \{S \subseteq E : \forall r \in S \exists u \in S \text{ with } u \rightarrow r\}. \quad (1)$$

Thus $S \in \mathcal{C}(D)$ exactly when every vertex of the induced subgraph $D|_S$ has in-degree at least one. The empty set belongs to $\mathcal{C}(D)$, and $\mathcal{C}(D)$ is closed under unions [13, Proposition 2]: a vertex of $S \cup T$ lies in S or in T , where it already has an in-neighbour. Consequently $\mathcal{C}(D)$ is the family of fixed sets of an interior operator on the Boolean lattice 2^E , namely the operator sending X to the largest supported subset of X .

Support families arise in the theory of autocatalytic networks. A reflexively autocatalytic and food-generated set (RAF) is a set of reactions that can be built up from a food set and in which every reaction is catalysed by a food molecule or by a product of the set [8, 9, 14]. When every reactant is food, food generation is automatic and the RAFs are exactly the nonempty supported sets of the digraph on the reactions whose arcs record which reaction’s product catalyses which reaction, with a loop recording catalysis by food [13, Section 3.2]. Steel showed that many generic facts about RAFs follow from the interior-operator property alone, and in his concluding section asked two questions: to characterize in set-theoretic terms the union-closed families of the form $\mathcal{C}(D)$, and, “more difficult”, to characterize the families realizable as the RAFs of an arbitrary catalytic reaction system [13, Section 4]. A companion paper [2] answers the second question for realizations on the given ground set: the fixed family of such a system is exactly an intersection $\mathcal{A} \cap \mathcal{C}(D)$ of an antimatroid $\mathcal{A}$ (recording food generation) with a digraph support family (recording catalysis). That result leaves the first question open, and it is the subject of the present paper. We stress that the question concerns a digraph on the *original* ground set E , with loops allowed; no auxiliary vertices are permitted.

1.2 The answer

The characterization is recursive. For a family $\mathcal{F} \subseteq 2^E$ and a set $S \subseteq E$ write

$$I_{\mathcal{F}}(S) = \bigcup \{T \in \mathcal{F} : T \subseteq S\} \quad (2)$$

for the union of the members of $\mathcal{F}$ below S ; when $\mathcal{F}$ is union-closed and contains the empty set, $I_{\mathcal{F}}(S)$ is the largest member of $\mathcal{F}$ contained in S , that is, the interior of S . For $r \in S$ the *rooted interval* beneath S is $[\{r\}, S] = \{X \subseteq E : r \in X \subseteq S\}$.

Starting from $\mathcal{U}_0 = 2^E$ with no roots used, a *legal step* selects an inclusion-maximal set S of the residual family $\mathcal{U}_i \setminus \mathcal{F}$, then a root $r \in S \setminus I_{\mathcal{F}}(S)$ that has not been used before, and passes to $\mathcal{U}_{i+1} = \mathcal{U}_i \setminus [\{r\}, S]$. The sequence *succeeds* if it reaches $\mathcal{U}_m = \mathcal{F}$. Our main result (Theorem 2.2) states that $\mathcal{F} = \mathcal{C}(D)$ for some digraph D on E if and only if some legal sequence succeeds; that for *every* fixed rule selecting the maximal residual set there is a successful sequence following that rule, so the only branching needed is over roots; that every legal sequence has at most $|E|$ steps; and that a successful sequence $(S_i, r_i)_{i < m}$ reconstructs a realizing digraph on exactly E through the predecessor rows

$$P(r_i) = E \setminus S_i, \quad P(v) = \{v\} \text{ for every unused vertex } v.$$

Each step therefore commits to one row of the graph, and the residual set fixes the entire row.

The proof complements every set. The complements of the supported sets are the models of the definite Horn rules $P(r) \rightarrow r$, one rule per head r , so realizability becomes the question whether the complementary family $\mathcal{K} = \{E \setminus S : S \in \mathcal{F}\}$ is the model class of a *single-head* definite Horn theory. Soundness is a direct computation. Completeness rests on Lemma 4.1: if a partial theory G extends to some single-head theory L with $\text{Mod}(L) = \mathcal{K}$, and B is a minimal model of G outside $\mathcal{K}$, then some rule $A \rightarrow r$ of L violated by B has $\text{cl}_G(A) = B$, and replacing $A \rightarrow r$ by $B \rightarrow r$ yields a single-head theory with the same models that still contains G . The invariant is closure equality, not equality of the original body with the residual, and it is precisely this invariant that survives an arbitrary choice of the residual set.

1.3 Contributions

Beyond the characterization itself we establish the following.

- *Exact trace information* (Section 6). Every successful sequence uses exactly the roots r with $\{r\} \notin \mathcal{F}$, each once; its length is therefore determined by $\mathcal{F}$, and $\mathcal{F}$ determines the loops of every realizing digraph although not its other arcs.

- *A complete finite recognizer* (Section 7). Under a fixed residual-selection rule the search tree over roots has at most $b! \sum_{k \leq b} 1/k!$ nodes, where b is the number of excluded singletons, and exhausting it is a proof of non-realizability. For a family given as an explicit list the recognition problem is in NP, with the digraph as certificate (Proposition 7.2).
- *Why the criterion is dynamic* (Section 8). Two realizable families on three elements have a non-realizable intersection, so no collection of closure axioms characterizes the class (Corollary 8.2). A Hall-type condition on maximal non-members is necessary (Proposition 8.5) but not sufficient (Example 8.7).
- *Composition* (Section 9). The class is closed under coordinate projection, by an explicit vertex-elimination formula (Theorem 9.1), so hidden vertices cannot rescue a non-realizable family; it is also closed under induced deletion and disjoint products.
- *Autocatalytic networks* (Section 10). A family is the fixed family of an elementary catalytic reaction system if and only if it passes the interval-deletion test, and same-ground RAF families are exactly intersections of an antimatroid with a family passing the test. The non-realizable three-element family above is the fixed family of a three-reaction system, which separates elementary from general systems.
- *Formal verification* (Section 11). The characterization, the trace checker and the consequences are formalized in Lean 4 [3] over Mathlib [15]; the few statements proved only by hand are listed explicitly.

Three worked examples (Section 3) exhibit a successful trace, a root choice that fails on a realizable family, a family with two realizations, and a family rejected by every branch.

1.4 Relation to single-head Horn theory

The complementary formulation belongs to established theory. Definite Horn theories, their least models and their closure systems are classical [4, 16]; a single-head theory is one in which no two rules share a head. Liberatore [12] studies single-head formulas in connection with logical forgetting and gives a complete algorithm for reconstructing a single-head formula equivalent to a given definite Horn formula [12, Theorem 4.1]; unit clauses are handled by propagation and reinsertion in the paragraph following that theorem. Our treatment allows empty bodies directly, which is a convenience rather than a capability absent from that work. Elimination of variables from single-head formulas is polynomial [11, Section 5.7], which is the logical counterpart of the projection theorem of Section 9. The contribution of this paper is the direct interval-deletion formulation on the original family, the completion-preserving normalization proof that makes the residual choice free, the exact reconstruction on the original ground set, the consequences listed above, and the machine-checked connection to digraph support and RAF semantics. No priority claim for general single-head reconstruction is made.

1.5 Organization

Section 2 states the criterion and the main theorem. Section 3 works three examples. Section 4 proves the normalization lemma and Section 5 the theorem. Sections 6–10 develop the consequences, Section 11 describes the formalization, and Section 12 lists open problems.

2 The intrinsic criterion

Throughout, E is a finite set with $n = |E|$, and a *family* is a subset $\mathcal{F} \subseteq 2^E$. A *predecessor map* is a function $P : E \rightarrow 2^E$; it encodes the digraph with an arc $u \rightarrow r$ whenever $u \in P(r)$, and a loop at r whenever $r \in P(r)$. A set S is *P-supported* if $S \cap P(r) \neq \emptyset$ for every $r \in S$,

and $\mathcal{C}(P)$ denotes the family of P -supported sets, which is (1) for the digraph encoded by P . The family $\mathcal{F}$ is *digraph-realizable* (on E) if $\mathcal{F} = \mathcal{C}(P)$ for some predecessor map P on E .

Two conventions deserve mention. A vertex with an empty row, $P(r) = \emptyset$, belongs to no supported set. A loop at r makes the support condition at r automatic. Neither $E \in \mathcal{F}$ nor $\bigcup \mathcal{F} = E$ is assumed anywhere.

Definition 2.1 (Legal interval deletion). Let $\mathcal{F} \subseteq 2^E$. A *state* is a pair $(\mathcal{U}, H)$ with $\mathcal{U} \subseteq 2^E$ and $H \subseteq E$; the initial state is $(2^E, \emptyset)$. A *legal step* from a state $(\mathcal{U}, H)$ with $\mathcal{U} \neq \mathcal{F}$ consists of

- (i) an inclusion-maximal member S of $\mathcal{U} \setminus \mathcal{F}$, called the *residual set*, and
- (ii) a *root* $r \in S \setminus I_{\mathcal{F}}(S)$ with $r \notin H$,

and leads to the state $(\mathcal{U} \setminus [\{r\}, S], H \cup \{r\})$. A *legal sequence* is a sequence of states $(\mathcal{U}_i, H_i)_{i \leq m}$ starting at the initial state in which each state is obtained from its predecessor by a legal step; we record it as $(S_i, r_i)_{i < m}$. It *succeeds* if $\mathcal{U}_m = \mathcal{F}$. A state with $\mathcal{U} \neq \mathcal{F}$ from which no legal step exists is a *dead end*. A *residual policy* is a function π assigning to every family $\mathcal{U}$ with $\mathcal{F} \subsetneq \mathcal{U}$ an inclusion-maximal member of $\mathcal{U} \setminus \mathcal{F}$; a legal sequence *follows* π if $S_i = \pi(\mathcal{U}_i)$ for every $i < m$.

The root condition is a condition on $\mathcal{F}$ alone. Unwinding (2),

$$r \notin I_{\mathcal{F}}(S) \iff \forall T \in \mathcal{F} : T \subseteq S \Rightarrow r \notin T. \quad (3)$$

Hence a legal step never deletes a member of $\mathcal{F}$: a member $T \in \mathcal{F}$ lies in $[\{r\}, S]$ only if $T \subseteq S$ and $r \in T$, which (3) excludes. Since $S \in [\{r\}, S]$ is deleted, $\mathcal{F} \subseteq \mathcal{U}_{i+1} \subsetneq \mathcal{U}_i$ along every legal sequence. In particular a family that does not contain the empty set, or is not union-closed, admits no successful sequence, in the first case because no rooted interval contains $\emptyset$, and in the second case by Theorem 2.2, since a successful sequence produces a realizing digraph.

Theorem 2.2 (Intrinsic characterization). *For every finite set E and every family $\mathcal{F} \subseteq 2^E$ the following are equivalent.*

- (1) *$\mathcal{F}$ is digraph-realizable on E : there is a directed graph D on exactly the vertex set E , loops permitted, with $\mathcal{F} = \mathcal{C}(D)$.*
- (2) *Some legal sequence for $\mathcal{F}$ succeeds.*
- (3) *For every residual policy π there is a successful legal sequence following π .*

Every legal sequence has at most $|E|$ steps. If $(S_i, r_i)_{i < m}$ is successful, then the predecessor map

$$P(r_i) = E \setminus S_i \quad (i < m), \quad P(v) = \{v\} \quad (v \notin \{r_0, \dots, r_{m-1}\}) \quad (4)$$

satisfies $\mathcal{C}(P) = \mathcal{F}$.

Remark 2.3. (a) The theorem is stated for arbitrary families. Its restriction to union-closed families containing the empty set answers Steel's question, and the procedure itself certifies those two properties on success.

(b) A policy fixes residual sets, not roots. Statement (3) asserts that a successful branch exists whichever maximal residual set is selected at each stage; it does not assert that every eligible root leads to success. Example 3.1 shows a realizable family with a dead-end root choice. Consequently a recognizer may select residual sets deterministically and must branch only over roots (Section 7).

(c) The reconstruction (4) is exact on every subset of E , not only on maximal or minimal members of $\mathcal{F}$. It need not be the only realization: Example 3.1 has four realizations, Example 3.2 has two.

(d) Boundary cases: for $E = \emptyset$ the only realizable family is $\{\emptyset\}$; for $\mathcal{F} = 2^E$ the empty sequence succeeds and the reconstruction places a loop at every vertex; for $\mathcal{F} = \{\emptyset\}$ the edgeless graph is a realization, and a successful sequence has exactly n steps.

3 Three worked examples

Let $E = \{a, b, c\}$. In the tables below a trace is listed as its residual sets and roots, followed by the rows reconstructed by (4). All three examples were also checked by exhaustive enumeration of the 2^9 predecessor maps on three vertices.

Example 3.1 (A chain with a loop). Let D have a loop at a and arcs $a \rightarrow b$, $b \rightarrow c$, so $P(a) = \{a\}$, $P(b) = \{a\}$, $P(c) = \{b\}$. Its support family is the chain

$$\mathcal{F} = \{\emptyset, \{a\}, \{a, b\}, \{a, b, c\}\},$$

and the complementary family is $\mathcal{K} = \{\emptyset, \{c\}, \{b, c\}, \{a, b, c\}\}$, the models of the rules $a \rightarrow b$ and $b \rightarrow c$ (the loop at a is the tautology $a \rightarrow a$). The residual family $2^E \setminus \mathcal{F} = \{\{b\}, \{c\}, \{a, c\}, \{b, c\}\}$ has two maximal members, $\{a, c\}$ and $\{b, c\}$, and $I_{\mathcal{F}}(\{a, c\}) = \{a\}$, $I_{\mathcal{F}}(\{b, c\}) = \emptyset$.

Policy 1: select $\{a, c\}$ first. The only eligible root is c . Deleting $[\{c\}, \{a, c\}] = \{\{c\}, \{a, c\}\}$ leaves the residual family $\{\{b\}, \{b, c\}\}$, whose maximal member is $\{b, c\}$; the eligible roots are b and c , but c is used, so $r_1 = b$. Deleting $\{\{b\}, \{b, c\}\}$ reaches $\mathcal{F}$. The trace $(\{a, c\}, c), (\{b, c\}, b)$ reconstructs $P(c) = E \setminus \{a, c\} = \{b\}$, $P(b) = E \setminus \{b, c\} = \{a\}$, $P(a) = \{a\}$: the original digraph.

Policy 2: select $\{b, c\}$ first. Now both b and c are eligible.

root at $\{b, c\}$	next residual sets	eligible unused roots	outcome
b	$\{a, c\}$	c	success
c	$\{a, c\}, \{b\}$	none at $\{a, c\}$	dead end

The successful branch reconstructs the same rows as Policy 1. Choosing c at $\{b, c\}$ commits to the row $P(c) = \{a\}$; together with the loop at a this would support $\{a, c\} \notin \mathcal{F}$, and indeed the residual set $\{a, c\}$ then has c as its only eligible root, which is already used. Thus the same realizable family succeeds or fails depending on the root, while the residual choice is immaterial, exactly as Theorem 2.2(3) predicts.

The family has four realizations: $P(b) = \{a\}$ and $P(c) = \{b\}$ are forced, and $P(a)$ may be any set containing a , because arcs into a looped vertex change no support condition. The reconstruction returns the sparsest of them. Two singletons are excluded from $\mathcal{F}$, and both traces have length two (Proposition 6.1).

Example 3.2 (The directed three-cycle). Let $\mathcal{F} = \{\emptyset, E\}$. The six proper nonempty subsets are residual and the three pairs are maximal; every interior $I_{\mathcal{F}}(S)$ with $S \neq E$ is empty, so every element of a pair is eligible. Select $\{b, c\}$, then $\{a, c\}$, then $\{a, b\}$ whenever available.

trace	outcome	reconstructed rows
$(\{b, c\}, b), (\{a, c\}, c), (\{a, b\}, a)$	success	$P(b) = \{a\}, P(c) = \{b\}, P(a) = \{c\}$
$(\{b, c\}, c), (\{a, c\}, a), (\{a, b\}, b)$	success	$P(c) = \{a\}, P(a) = \{b\}, P(b) = \{c\}$
$(\{b, c\}, b), (\{a, c\}, a), (\{a, b\}, -)$	dead end	

The dead end is the greedy failure of the complementary formulation: accepting $a \rightarrow b$ and then $b \rightarrow a$ closes a two-cycle on $\{a, b\}$ whose two heads are then unavailable for the residual set $\{a, b\}$: the only member of $\mathcal{F}$ below it is $\emptyset$, so its eligible roots are a and b , and both are used. A short closed cycle can consume heads that the rest of the closure class requires. The two successful traces reconstruct the two orientations $a \rightarrow b \rightarrow c \rightarrow a$ and $a \rightarrow c \rightarrow b \rightarrow a$ of the three-cycle, and these are the only realizations of $\mathcal{F}$. At each head, both other singletons are minimal sets whose complementary closure contains the head, so a single-head realization can coexist with several minimal “triggers” per head.

Example 3.3 (A family rejected by every branch). Let

$$\mathcal{F}_\wedge = \{\emptyset, \{a\}, \{b\}, \{a, b\}, \{a, b, c\}\}.$$

The residual family is $\{\{c\}, \{a, c\}, \{b, c\}\}$ with maximal members $\{a, c\}$ and $\{b, c\}$, and $I_{\mathcal{F}}(\{a, c\}) = \{a\}$, $I_{\mathcal{F}}(\{b, c\}) = \{b\}$, so c is the only eligible root at either. Selecting $\{a, c\}$ and deleting $\{\{c\}, \{a, c\}\}$ leaves the residual set $\{b, c\}$, whose only eligible root c is used: a dead end. Selecting $\{b, c\}$ first fails symmetrically. Under either policy the whole search tree has three nodes and no success, so by Theorem 2.2(3) the family is not digraph-realizable. Directly: $\{a\}, \{b\} \in \mathcal{F}_\wedge$ force loops at a and b , $\{c\} \notin \mathcal{F}_\wedge$ forbids a loop at c , and supporting c inside E then needs $a \in P(c)$ or $b \in P(c)$, which would support $\{a, c\}$ or $\{b, c\}$. The family asks for a *conjunction*, c requires both a and b , whereas a predecessor row offers *alternatives*. This family recurs in Sections 8 and 10.

4 Definite Horn theories and the normalization lemma

4.1 Rules, models and least models

A *definite Horn rule* on E is a pair $A \rightarrow r$ with $A \subseteq E$ (the body) and $r \in E$ (the head). A set $X \subseteq E$ *models* the rule if $A \subseteq X$ implies $r \in X$. For a set G of rules, $\text{Mod}(G)$ denotes the family of sets modelling every rule of G , and $\text{Heads}(G)$ the set of heads occurring in G . The theory G is *single-head* if any two rules of G with the same head are identical. A rule with $r \in A$ is tautological and may be omitted; empty bodies are allowed and act as unit clauses $\emptyset \rightarrow r$, forcing r into every model.

For $A \subseteq E$ let

$$\text{cl}_G(A) = \bigcap \{X \subseteq E : A \subseteq X, X \in \text{Mod}(G)\}. \quad (5)$$

The indexing family contains E . An intersection of models is a model, because a body contained in the intersection is contained in each model, whose head then lies in each of them. Hence $\text{cl}_G(A)$ is a model, contains A , and is contained in every model containing A ; it is the least model above A , computable by forward chaining in linear time [4]. In particular $A \mapsto \text{cl}_G(A)$ is a closure operator whose closed sets are exactly $\text{Mod}(G)$.

4.2 Residual normalization

Lemma 4.1 (Residual normalization and replacement). *Let $G \subseteq L$ be sets of rules with L single-head, and let B be inclusion-minimal in $\text{Mod}(G) \setminus \text{Mod}(L)$. Then there is a rule $A \rightarrow r \in L$ with*

$$A \subseteq B, \quad r \notin B, \quad r \notin \text{Heads}(G), \quad \text{cl}_G(A) = B. \quad (6)$$

Moreover $L' = (L \setminus \{A \rightarrow r\}) \cup \{B \rightarrow r\}$ is single-head, contains $G \cup \{B \rightarrow r\}$, and $\text{Mod}(L') = \text{Mod}(L)$.

Proof. Since $B \notin \text{Mod}(L)$, some rule $A \rightarrow r$ of L is violated by B : $A \subseteq B$ and $r \notin B$. If r were a head of G , then since L is single-head and $G \subseteq L$, the rule of G with head r would be this same rule, contradicting $B \in \text{Mod}(G)$. So r is fresh.

Put $T = \text{cl}_G(A)$. Because B is a G -model containing A , $T \subseteq B$, hence $r \notin T$, while $A \subseteq T$. So T violates $A \rightarrow r$ and lies in $\text{Mod}(G) \setminus \text{Mod}(L)$. Minimality of B gives $T = B$, which is (6).

The removed rule is not in G because its head is fresh, so $L' \supseteq G$; and $B \rightarrow r$ replaces the unique rule of L with head r , so L' is single-head and contains $G \cup \{B \rightarrow r\}$.

If $X \in \text{Mod}(L)$ and $B \subseteq X$, then $A \subseteq X$, so $r \in X$; thus every L -model satisfies $B \rightarrow r$ and is an L' -model. Conversely let $X \in \text{Mod}(L')$. All rules of L other than $A \rightarrow r$ belong to L' . If $A \subseteq X$, then X models $G \subseteq L'$, so $B = \text{cl}_G(A) \subseteq X$, and the rule $B \rightarrow r$ gives $r \in X$. Thus X also satisfies $A \rightarrow r$, and $X \in \text{Mod}(L)$. $\square$

The equality in (6) concerns the closure of the original body under the *accepted* rules G , not the body itself; the accepted rules are exactly what makes the replacement model-preserving. The following example shows that the distinction is real.

Example 4.2 (The body genuinely changes). On $E = \{a, b, c\}$ let $G = \{a \rightarrow b\}$, $L = \{a \rightarrow b, a \rightarrow c\}$ and $B = \{a, b\}$. Then B models G but not L , and it is minimal with this property since $\emptyset$, $\{a\}$ and $\{b\}$ model L or fail G . The violated rule $a \rightarrow c$ has body $\{a\} \subsetneq B$, and $\text{cl}_G(\{a\}) = \{a, b\} = B$. Replacing it by $\{a, b\} \rightarrow c$ gives $L' = \{a \rightarrow b, \{a, b\} \rightarrow c\}$ with the same models $\{\emptyset, \{b\}, \{c\}, \{b, c\}, \{a, b, c\}\}$, because the accepted rule $a \rightarrow b$ is retained. Replacing $a \rightarrow c$ by $\{a, b\} \rightarrow c$ in the *absence* of $a \rightarrow b$ would admit the new model $\{a\}$.

5 Proof of Theorem 2.2

5.1 Soundness and exact reconstruction

Let $(S_i, r_i)_{i < m}$ be a successful legal sequence and define P by (4); the roots are distinct, so P is well defined. For a used root r_i , a set X violates the support condition at r_i exactly when $r_i \in X$ and $X \cap (E \setminus S_i) = \emptyset$, that is, when $r_i \in X \subseteq S_i$. So the row at r_i excludes precisely the interval $[\{r_i\}, S_i]$ deleted at step i , and the loop at an unused vertex excludes nothing. Intersecting the row conditions,

$$\mathcal{C}(P) = 2^E \setminus \bigcup_{i < m} [\{r_i\}, S_i] = \mathcal{U}_m = \mathcal{F}.$$

This proves (2) $\Rightarrow$ (1) together with the reconstruction statement. Since a legal step uses a fresh root, every legal sequence has at most n steps.

5.2 Complementary semantics

Let $\mathcal{K} = \{E \setminus S : S \in \mathcal{F}\}$. For a predecessor map P and $X = E \setminus S$,

$$S \in \mathcal{C}(P) \iff \forall r \in S : S \cap P(r) \neq \emptyset \iff \forall r \in E : P(r) \subseteq X \Rightarrow r \in X. \quad (7)$$

Hence $\mathcal{C}(P)$ is the complement family of $\text{Mod}(G_P)$, where $G_P = \{P(r) \rightarrow r : r \in E\}$ is single-head. A loop gives a tautology; an empty row gives the unit clause $\emptyset \rightarrow r$. Conversely a single-head theory L defines a predecessor map by taking the body of the rule with head r , or $\{r\}$ if no rule has head r , and (7) shows that its support family is the complement family of $\text{Mod}(L)$. Thus

$$\mathcal{F} \text{ is digraph-realizable} \iff \mathcal{K} = \text{Mod}(L) \text{ for some single-head } L. \quad (8)$$

Along a legal sequence put $B_j = E \setminus S_j$ and $G_i = \{B_j \rightarrow r_j : j < i\}$. Deleting $[\{r_j\}, S_j]$ from $\mathcal{U}_j$ removes exactly the sets X with $r_j \in X \subseteq S_j$, whose complements are the sets Y with $B_j \subseteq Y$ and $r_j \notin Y$, that is, the violators of $B_j \rightarrow r_j$. By induction

$$\{E \setminus X : X \in \mathcal{U}_i\} = \text{Mod}(G_i) \quad (i \leq m). \quad (9)$$

Complementation reverses inclusion, so a maximal member S_i of $\mathcal{U}_i \setminus \mathcal{F}$ corresponds to a minimal member B_i of $\text{Mod}(G_i) \setminus \mathcal{K}$. Likewise, by (3), $r \in S \setminus I_{\mathcal{F}}(S)$ holds iff every member of $\mathcal{K}$ containing $E \setminus S$ contains r , and $r \in S$ iff $r \notin E \setminus S$. So a legal step is, in complementary terms, the choice of a minimal residual model B and a fresh head $r \notin B$ forced by B in $\mathcal{K}$.

5.3 Completeness for every residual policy

Assume (1), so by (8) there is a single-head L_0 with $\text{Mod}(L_0) = \mathcal{K}$; tautologies may be dropped. Fix a residual policy π . We construct a legal sequence following π while maintaining a single-head theory L_i with

$$G_i \subseteq L_i, \quad \text{Mod}(L_i) = \mathcal{K}. \quad (10)$$

Initially $G_0 = \emptyset \subseteq L_0$. Suppose $(\mathcal{U}_i, H_i)$ has been reached with (10), where $H_i = \{r_j : j < i\} = \text{Heads}(G_i)$. If $\mathcal{U}_i = \mathcal{F}$ we are done. Otherwise $\mathcal{F} \subsetneq \mathcal{U}_i$, and π selects a maximal $S_i \in \mathcal{U}_i \setminus \mathcal{F}$. By (9), $B_i = E \setminus S_i$ is inclusion-minimal in $\text{Mod}(G_i) \setminus \text{Mod}(L_i)$. Lemma 4.1 yields a head $r_i \notin B_i$ with $r_i \notin \text{Heads}(G_i) = H_i$, and a single-head L_{i+1} containing $G_{i+1} = G_i \cup \{B_i \rightarrow r_i\}$ with $\text{Mod}(L_{i+1}) = \text{Mod}(L_i) = \mathcal{K}$.

It remains to check that (S_i, r_i) is a legal step. We have $r_i \in S_i$ because $r_i \notin B_i$, and $r_i \notin H_i$. If $T \in \mathcal{F}$ and $T \subseteq S_i$, then $E \setminus T \in \mathcal{K} = \text{Mod}(L_{i+1})$ contains B_i , so it contains r_i by the rule $B_i \rightarrow r_i$; hence $r_i \notin T$. By (3), $r_i \notin I_{\mathcal{F}}(S_i)$. So the step is legal, and (10) holds at $i + 1$.

Each step uses a new element of E , so the construction cannot continue for more than n steps; but it continues whenever $\mathcal{U}_i \neq \mathcal{F}$. Hence some $\mathcal{U}_m = \mathcal{F}$ with $m \leq n$, and the sequence following π succeeds. This proves (1) $\Rightarrow$ (3). For (3) $\Rightarrow$ (2) take any policy, for instance the one selecting the first maximal residual set in a fixed linear order of 2^E . $\square$

Remark 5.1. The formal proof (Section 11) runs the induction on the finite family $\mathcal{U}_i$, which strictly shrinks at each step, rather than on a step counter, and therefore never assumes that a successful sequence exists. The Lean statement of the policy theorem quantifies over an arbitrary function π subject only to the validity condition of Definition 2.1.

6 What a successful sequence knows

Proposition 6.1 (Exact roots and length). *Let $b_{\mathcal{F}} = \{r \in E : \{r\} \notin \mathcal{F}\}$. The roots of every successful legal sequence for $\mathcal{F}$ are distinct and their set is exactly $b_{\mathcal{F}}$; hence every successful sequence has length $|b_{\mathcal{F}}|$.*

Proof. If $\{r\} \in \mathcal{F}$ then r is never eligible, because $\{r\}$ is a member of $\mathcal{F}$ below any residual set containing r . Conversely, if $\{r\} \notin \mathcal{F}$ then $\{r\} \in \mathcal{U}_0$ must be deleted before success, and $\{r\} \in [\{s\}, S]$ forces $s = r$; so r occurs as a root. Freshness gives distinctness. $\square$

Corollary 6.2 (Loops are determined). *For every realization P of $\mathcal{F}$ and every $r \in E$, $r \in P(r)$ if and only if $\{r\} \in \mathcal{F}$. Other arcs need not be determined: if every vertex has a loop, arbitrary additional arcs leave the support family equal to 2^E .*

Proof. $\{r\}$ is P -supported exactly when $\{r\} \cap P(r) \neq \emptyset$. $\square$

Exact reconstruction also preserves every interior query. If $\mathcal{F}$ is union-closed and contains $\emptyset$, its interior operator is $S \mapsto I_{\mathcal{F}}(S)$, and any interior operator with fixed family $\mathcal{F}$ coincides with it (a fixed set below S lies below the interior of S , and the interior is itself fixed). Hence the interior operator of the reconstructed digraph, which sends X to its largest P -supported subset, equals the interior operator of $\mathcal{F}$. This is the statement formalized as `reconstructed_interior_eq`.

7 Search and complexity

7.1 A complete finite recognizer

Theorem 2.2(3) turns the criterion into a recognizer. Fix a residual policy π and run the following recursion from $(\mathcal{F}, 2^E, \emptyset)$.

```

RECOGNIZE( $\mathcal{F}, \mathcal{U}, H$ ):
  if  $\mathcal{U} = \mathcal{F}$ : return success with the empty suffix
   $S \leftarrow \pi(\mathcal{U})$ 
  for each  $r \in S \setminus (I_{\mathcal{F}}(S) \cup H)$ :
     $\sigma \leftarrow \text{RECOGNIZE}(\mathcal{F}, \mathcal{U} \setminus [\{r\}, S], H \cup \{r\})$ 
    if  $\sigma$  succeeds: return  $(S, r)$  followed by  $\sigma$ 
  return failure

```

Proposition 7.1. *RECOGNIZE returns a successful legal sequence if $\mathcal{F}$ is digraph-realizable, and returns failure otherwise; a returned sequence reconstructs a realization by (4). With $b = |b_{\mathcal{F}}|$, the number of recursive calls is at most*

$$\sum_{j=0}^b \frac{b!}{(b-j)!} = b! \sum_{k=0}^b \frac{1}{k!} < e \cdot b!. \quad (11)$$

Proof. Every sequence explored is legal and follows π ; success is sound by Theorem 2.2(2) $\Rightarrow$ (1), and if $\mathcal{F}$ is realizable then Theorem 2.2(3) guarantees a successful sequence following π , which the exhaustive loop over roots finds. Because π is deterministic, a call is determined by its list of roots so far, which is a sequence of distinct elements of $b_{\mathcal{F}}$ of length $j \leq b$, because every root used is eligible when it is used and hence, by the first half of the proof of Proposition 6.1, is an excluded singleton. There are $b!/(b-j)!$ such sequences at depth j . $\square$

The bound counts calls, not elementary operations, and holds for any policy; a policy that tends to select large residual sets is natural because a large S removes more sets and typically leaves fewer eligible roots. Note that a rejection by RECOGNIZE is a proof of non-realizability under a single policy; there is no need to vary the residual selection. Example 3.3 was rejected by a three-node tree.

7.2 Representation and complexity

The cost of the recursion depends on how $\mathcal{F}$ is presented. If $\mathcal{F}$ is given as a truth table over 2^E , membership, interiors, maximal residual sets and interval deletions are all obtained by scans of the table, and the total running time is $2^{O(n)}$ times the bound (11). If $\mathcal{F}$ is given as an explicit list of m subsets, membership costs $O(nm)$ and $I_{\mathcal{F}}(S)$ costs $O(nm)$, but locating a maximal member of $\mathcal{U}_i \setminus \mathcal{F}$ is not a scan of the list. If $\mathcal{F}$ is given succinctly, for instance by a Horn formula for $\mathcal{K}$ or by a membership oracle, then a trace of $\leq n$ steps is short, but verifying its maximality conditions is not obviously polynomial. The next statement identifies a certificate that avoids maximality altogether.

Proposition 7.2 (Explicit-list recognition is in NP). *Given E with $|E| = n$ and a family $\mathcal{F}$ listed explicitly as m subsets, deciding whether $\mathcal{F}$ is digraph-realizable is in NP. A certificate is a predecessor map P with $\mathcal{C}(P) = \mathcal{F}$, and it can be verified in time polynomial in $n + m$.*

Proof. A predecessor map has n rows of n bits. Checking $\mathcal{F} \subseteq \mathcal{C}(P)$ takes $O(n^2m)$ operations. For $\mathcal{C}(P) \subseteq \mathcal{F}$, recall from (7) that $\mathcal{C}(P)$ is the complement family of $\text{Mod}(G_P)$, and that $\text{Mod}(G_P)$ is the closed-set family of the closure operator cl_{G_P} , computable by forward chaining

in polynomial time [4]. Ganter’s NEXTCLOSURE algorithm enumerates the closed sets of a closure operator with polynomial delay, using $O(n)$ closure computations per closed set [7]. Enumerate the models of G_P , testing for each whether its complement lies in $\mathcal{F}$ (cost $O(nm)$), and stop at the first failure or after $m + 1$ models have been produced. Since distinct models have distinct complements, at most m models can pass the test, so the enumeration either reports a failure or terminates within $m + 1$ steps. In the latter case $\mathcal{C}(P) \subseteq \mathcal{F}$, and together with the first check $\mathcal{C}(P) = \mathcal{F}$. The total time is polynomial in $n + m$. $\square$

We do not know whether explicit-list recognition is in P. In the succinct setting Liberatore’s algorithm decides whether a definite Horn formula is equivalent to a single-head one [12, Theorem 4.1], and general definite Horn theories are learnable in polynomial time from membership and equivalence queries [1]; whether the single-head restriction admits an efficient recognizer under some natural presentation of $\mathcal{F}$ is a question we leave open (Section 12).

8 Why the criterion is dynamic

A characterization of a class of set families is often expected to be static: a list of closure conditions, exchange axioms or forbidden configurations. Union-closed families, antimatroids and convex geometries are of this kind [5, 10]. Theorem 2.2 is instead an elimination scheme, in the tradition of perfect elimination orderings for chordal graphs [6] and shelling orders for antimatroids [10]. This section explains why a purely closure-axiomatic description is impossible, records two natural static conditions that fail, and identifies the Hall-type condition that the procedure enforces dynamically.

8.1 No closure axioms suffice

Proposition 8.1 (Intersection on a shared ground). *Let P_1 and P_2 on $E = \{a, b, c\}$ both have loops at a and b , with $P_1(c) = \{a\}$ and $P_2(c) = \{b\}$. Then $\mathcal{C}(P_1) \cap \mathcal{C}(P_2) = \mathcal{F}_\wedge$, the family of Example 3.3, which is not digraph-realizable.*

Proof. $\mathcal{C}(P_1) = \mathcal{F}_\wedge \cup \{\{a, c\}\}$ and $\mathcal{C}(P_2) = \mathcal{F}_\wedge \cup \{\{b, c\}\}$; the intersection is $\mathcal{F}_\wedge$, and Example 3.3 shows that $\mathcal{F}_\wedge$ is not realizable. $\square$

Call a condition on families a *closure rule* if it has the form “for all $S_1, \dots, S_k \in \mathcal{F}$, $\Phi(S_1, \dots, S_k) \in \mathcal{F}$ ”, where $k \geq 0$ and $\Phi : (2^E)^k \rightarrow 2^E$ is an arbitrary function. Union-closure ($\Phi(S, T) = S \cup T$) and $\emptyset \in \mathcal{F}$ ($k = 0$) are closure rules.

Corollary 8.2. *The class of digraph-realizable families on a fixed ground set with at least three elements is not the class of families satisfying a set of closure rules.*

Proof. If two families satisfy a closure rule, so does their intersection: the hypothesis $S_1, \dots, S_k \in \mathcal{F}_1 \cap \mathcal{F}_2$ places $\Phi(S_1, \dots, S_k)$ in both. So the class defined by any set of closure rules is closed under intersection, while the realizable class is not, by Proposition 8.1. For a larger ground set give every additional vertex a loop in both digraphs; the intersection is then $\mathcal{F}_\wedge \boxtimes 2^{E \setminus \{a, b, c\}}$, whose realizability would imply that of $\mathcal{F}_\wedge$ by the induced-deletion identity (13) of Section 9. $\square$

The argument does not exclude characterizations with negative hypotheses, such as forbidden configurations; we do not know whether one exists.

8.2 Two static conditions that fail

Example 8.3 (Interpolation). Every support family satisfies the following: whenever $U \subsetneq W$ are members of $\mathcal{F}$, either $W \setminus U \in \mathcal{F}$ or $U \cup \{w\} \in \mathcal{F}$ for some $w \in W \setminus U$. (If no $w \in W \setminus U$ is supported by $U \cup \{w\}$, then every such w has all its in-neighbours in W outside U , so $W \setminus U$ is supported.) The family $\mathcal{F}_\wedge$ satisfies this condition: pairs inside $\{a, b\}$ are all in $\mathcal{F}_\wedge$, and for $W = E$ the choices $U = \emptyset$, U a singleton and $U = \{a, b\}$ give $E \in \mathcal{F}_\wedge$, $\{a, b\} \in \mathcal{F}_\wedge$ and $E \in \mathcal{F}_\wedge$ respectively. Yet $\mathcal{F}_\wedge$ is not realizable. The condition does not see that one row must serve all supersets at once.

Example 8.4 (Locally admissible rows). Let $\mathcal{F}_{\text{loc}} = \{\emptyset, \{a\}, E\}$. For each r one can choose a row that is compatible with every member of $\mathcal{F}_{\text{loc}}$: the loop $\{a\}$ at a , and $\{c\}$ at b , $\{b\}$ at c (each member containing b or c is E , which contains the other). But the rows jointly support $\{b, c\} \notin \mathcal{F}_{\text{loc}}$, and indeed the family is not realizable: the loop at a is forced, loops at b, c are forbidden, $\{a, b\}, \{a, c\} \notin \mathcal{F}_{\text{loc}}$ forbid $a \in P(b)$ and $a \in P(c)$, so supporting E forces $c \in P(b)$ and $b \in P(c)$, which supports $\{b, c\}$. Rows chosen separately can add an unwanted self-supporting component; the procedure avoids this by fixing rows one at a time against the current family.

8.3 A Hall-type necessary condition

Lemma 4.1 assigns to every minimal residual model a fresh forced head. These heads are distinct.

Proposition 8.5 (Residual matching). *Let $G \subseteq L$ with L single-head, and let $\mathcal{B}_G$ be the set of inclusion-minimal members of $\text{Mod}(G) \setminus \text{Mod}(L)$. There is an injective map $\rho : \mathcal{B}_G \rightarrow E$ with*

$$\rho(B) \in \text{cl}_L(B) \setminus (B \cup \text{Heads}(G)) \quad (B \in \mathcal{B}_G).$$

Consequently $|\mathcal{Q}| \leq |\bigcup_{B \in \mathcal{Q}} \text{cl}_L(B) \setminus (B \cup \text{Heads}(G))|$ for every $\mathcal{Q} \subseteq \mathcal{B}_G$.

Proof. For $B \in \mathcal{B}_G$ let $A_B \rightarrow r_B$ be the rule of L supplied by Lemma 4.1, and put $\rho(B) = r_B$. Then $r_B \notin B \cup \text{Heads}(G)$, and $r_B \in \text{cl}_L(B)$ since every L -model containing $B \supseteq A_B$ contains r_B . If $r_B = r_C$ then, L being single-head, $A_B \rightarrow r_B$ and $A_C \rightarrow r_C$ are the same rule, so $B = \text{cl}_G(A_B) = \text{cl}_G(A_C) = C$. The inequality is the cardinality of the injective image of $\mathcal{Q}$. $\square$

In the original family, with $G = \emptyset$, the minimal members of $2^E \setminus \mathcal{K}$ are the complements of the maximal members S of $2^E \setminus \mathcal{F}$, and $\text{cl}_L(E \setminus S) \setminus (E \setminus S) = S \setminus I_{\mathcal{F}}(S)$ is the set of eligible roots at S . Hence:

Corollary 8.6 (Hall condition on maximal non-members). *If $\mathcal{F}$ is digraph-realizable, then the inclusion-maximal members of $2^E \setminus \mathcal{F}$ admit distinct representatives $r(S) \in S \setminus I_{\mathcal{F}}(S)$. More generally, at any state $(\mathcal{U}_i, H_i)$ of a legal sequence that extends to a successful one, the maximal members of $\mathcal{U}_i \setminus \mathcal{F}$ admit distinct representatives in $S \setminus (I_{\mathcal{F}}(S) \cup H_i)$.*

For $\mathcal{F}_\wedge$ the maximal non-members $\{a, c\}$ and $\{b, c\}$ both have c as their only eligible root, so the Hall condition already rejects it. A violation at a later state rejects the current branch only. The condition is not sufficient.

Example 8.7 (Initial matching without a realization). Let $E = A \sqcup B$ with $A = \{a, b, c\}$, $B = \{d, e, f\}$, and let $\mathcal{K}$ consist of E together with the sixteen sets meeting each of A and B in at most one element; $\mathcal{F}$ is the complement family. The minimal non-members of $\mathcal{K}$ are the six pairs inside A or inside B , each of which is contained in no member of $\mathcal{K}$ except E . Matching each pair to the third element of its own triple satisfies the Hall condition at the initial state, with room to spare.

Suppose L is single-head with $\text{Mod}(L) = \mathcal{K}$. From the pair $\{a, b\}$ forward chaining must reach E . The first rule fired has body contained in $\{a, b\}$; a proper subset of $\{a, b\}$ is a member of $\mathcal{K}$ and cannot be the body of a non-tautological rule, since it would violate that rule itself. So L contains a non-tautological rule with body exactly $\{a, b\}$, and similarly for each of the six pairs. Six distinct bodies give six rules with, by single-headedness, six distinct heads, so every element of E is the head of a pair-bodied rule and L has no other non-tautological rules. If the head of the rule with body $\{a, b\}$ lies in B , say d , then from $\{a, b, d\}$ no further rule fires (no other pair inside a triple is contained in $\{a, b, d\}$), so $\text{cl}_L(\{a, b\}) \neq E$. Hence each pair's head is the third element of its own triple, all six rules stay inside their triples, and $\text{cl}_L(\{a, b\}) = A \neq E$. Either way $\text{Mod}(L) \neq \mathcal{K}$, so $\mathcal{F}$ is not realizable. The recognizer used for the checks of Section 11 rejects it after exploring 322 memoized states.

The failure is one of connectivity: the heads matched to the pairs must also chain the two triples together, and a matching does not see that. The normalization lemma sees it because it carries a whole completion L along.

9 Projection, deletion and products

Which operations on families preserve realizability? This matters for composing and reducing networks. For $V \subseteq E$ define the *coordinate projection* $\pi_V(\mathcal{F}) = \{S \cap V : S \in \mathcal{F}\}$: a set of visible coordinates is admitted if some assignment of the hidden coordinates $W = E \setminus V$ extends it to a member of $\mathcal{F}$.

9.1 Elimination of one vertex

Let P be a predecessor map on E and $v \in E$. For $r \neq v$ define

$$Q(r) = \begin{cases} P(r), & v \notin P(r), \\ \{r\}, & v \in P(r) \text{ and } v \in P(v), \\ (P(r) \setminus \{v\}) \cup P(v), & v \in P(r) \text{ and } v \notin P(v), \end{cases} \quad (12)$$

and $Q(v) = \emptyset$. Every row $Q(r)$ with $r \neq v$ lies in $E \setminus \{v\}$, and v belongs to no Q -supported set.

Theorem 9.1 (Single-vertex projection). $\mathcal{C}(Q) = \{S \setminus \{v\} : S \in \mathcal{C}(P)\}$.

Proof. (The right side is contained in $\mathcal{C}(Q)$.) Let $S \in \mathcal{C}(P)$ and $T = S \setminus \{v\}$; take $r \in T$ and $u \in S \cap P(r)$. If $v \in P(r)$ and $v \in P(v)$, then $r \in Q(r)$ and r supports itself in T . Otherwise, if $u \neq v$ then $u \in T \cap Q(r)$ in each remaining case of (12). If $u = v$ then $v \in P(r)$ and $v \notin P(v)$; since $v \in S$ is P -supported there is $w \in S \cap P(v)$ with $w \neq v$, so $w \in T$ and $w \in Q(r)$ by the third case.

($\mathcal{C}(Q)$ is contained in the right side.) Let $T \in \mathcal{C}(Q)$, so $v \notin T$. If $v \in P(v)$, put $S = T \cup \{v\}$: v supports itself, a vertex r with $v \in P(r)$ is supported by v , and a vertex with $v \notin P(r)$ keeps its Q -predecessor, which is a P -predecessor. If $v \notin P(v)$ and $T \cap P(v) \neq \emptyset$, again put $S = T \cup \{v\}$: v is supported by $T \cap P(v)$, vertices with $v \in P(r)$ are supported by v , and the others as before. If $v \notin P(v)$ and $T \cap P(v) = \emptyset$, put $S = T$: for $r \in T$ with $v \in P(r)$ the Q -predecessor of r in T lies in $(P(r) \setminus \{v\}) \cup P(v)$ but not in $P(v)$, hence in $P(r)$; for the other r the Q -predecessor is a P -predecessor. In every case $S \in \mathcal{C}(P)$ and $S \setminus \{v\} = T$. $\square$

Corollary 9.2. Let $\mathcal{F}$ be digraph-realizable on E and $V \subseteq E$.

- (i) $\pi_V(\mathcal{F})$ is digraph-realizable, by a digraph on exactly the vertex set V .

(ii) (*Hidden vertices cannot rescue.*) If a family $\mathcal{H} \subseteq 2^V$ is not digraph-realizable on V , then there is no digraph on any $E \supseteq V$ whose support family projects onto $\mathcal{H}$.

(iii) For $X \subseteq V$, $I_{\pi_V(\mathcal{F})}(X) = I_{\mathcal{F}}(X \cup (E \setminus V)) \cap V$; this holds for every family $\mathcal{F}$.

Proof. (i) Eliminate the vertices of $E \setminus V$ one at a time by Theorem 9.1; successive set differences compose, and the final map has all rows inside V , so it restricts to V . (ii) is the contrapositive of (i). (iii) If $r \in S \cap V \subseteq X$ with $S \in \mathcal{F}$, then $S \subseteq X \cup (E \setminus V)$ and $r \in I_{\mathcal{F}}(X \cup (E \setminus V)) \cap V$; conversely if $r \in V$ lies in $S \in \mathcal{F}$ with $S \subseteq X \cup (E \setminus V)$, then $S \cap V \subseteq X$ is a member of $\pi_V(\mathcal{F})$ containing r . $\square$

Formula (12) is the graph-level form of variable elimination in single-head Horn theories, for which polynomial forgetting algorithms are known [11, Section 5.7]; the statement here adds the treatment of loops and the exact visible-ground conclusion.

9.2 Deletion, products, and the composition boundary

For $V \subseteq E$ the *induced deletion* of $E \setminus V$ restricts rows: $P|_V(r) = P(r) \cap V$ for $r \in V$. A set already contained in V has the same predecessors inside itself before and after the restriction, so

$$\mathcal{C}(P|_V) = \{S \in \mathcal{C}(P) : S \subseteq V\}. \quad (13)$$

Projection and deletion differ. For the two-cycle $a \leftrightarrow h$, $\mathcal{C}(P) = \{\emptyset, \{a, h\}\}$; projecting onto $\{a\}$ gives $\{\emptyset, \{a\}\}$, realized by a loop at a , while deleting h leaves $\{\emptyset\}$. Projection retains support that was supplied by hidden vertices; deletion removes it. Neither is a dynamical statement about a physical intervention.

For families on disjoint ground sets E_1, E_2 , the *independent product* is $\mathcal{F}_1 \boxtimes \mathcal{F}_2 = \{S_1 \cup S_2 : S_i \in \mathcal{F}_i\}$. The disjoint union of realizing digraphs realizes it, because with no cross arcs a set is supported iff each part is supported in its own component.

Operation	Preserves realizability?	Reason
Induced deletion	yes	(13)
Coordinate projection	yes	Theorem 9.1
Independent product	yes	disjoint union of digraphs
Intersection on a shared ground	no	Proposition 8.1

Combining Corollary 9.2(ii) with Proposition 8.1: no digraph with hidden vertices implements the conjunction family $\mathcal{F}_\wedge$ by coordinate projection. The next section shows that food generation does.

10 Autocatalytic networks

10.1 Catalytic reaction systems

A *catalytic reaction system* (CRS) consists of a finite set M of molecule types, a finite set R of reactions, each with a set of reactants and a set of products in M , a food set $F \subseteq M$, and a catalysis relation between molecules and reactions [8, 14]. A set $S \subseteq R$ is *food-generated* if its reactions can be ordered so that the reactants of each lie in F or among the products of earlier reactions of S ; equivalently, every reactant of S lies in the closure of F under S . A nonempty food-generated S is a RAF if every reaction of S is catalysed by a molecule of F or by a product of a reaction of S . Unions of RAFs are RAFs, so the map sending $T \subseteq R$ to the largest RAF inside T (or $\emptyset$) is an interior operator; its fixed family, the RAFs together with $\emptyset$, is the *fixed family* of the system [13]. A CRS is *elementary* if every reactant of every

reaction is food; then every subset of R is food-generated. We use the definitions of [2], whose closure-based food generation agrees with the ordering-based one of [8].

For an arbitrary CRS define the predecessor map on R by

$$P(r) = \{u \in R : \text{some product of } u \text{ catalyses } r\} \cup \{r : \text{some food molecule catalyses } r\}, \quad (14)$$

so a food catalyst appears as a loop. On a food-generated set the RAF condition is exactly P -support [13, Section 3.2], and in an elementary system the fixed family is therefore $\mathcal{C}(P)$ [2].

10.2 Elementary systems and the RAF factorization

Theorem 10.1 (Elementary fixed families). *A family $\mathcal{F} \subseteq 2^E$ is the fixed family of an elementary CRS with reaction set E if and only if $\mathcal{F}$ admits a successful legal interval-deletion sequence.*

Proof. If $\mathcal{F}$ is the fixed family of an elementary system, then $\mathcal{F} = \mathcal{C}(P)$ for the map (14), and Theorem 2.2 supplies a successful sequence. Conversely, given a successful sequence reconstruct P by (4) and build the system with one food molecule f , one private product c_v for each $v \in E$, reactions $r_v : f \rightarrow c_v$, and catalysis c_u catalyses r_v iff $u \in P(v)$, with no other catalysis. Every subset is food-generated because all reactants are food. The product c_u is available exactly when r_u is selected, so the RAF condition on S is P -support of S , and the fixed family is $\mathcal{C}(P) = \mathcal{F}$. $\square$

Note that a loop $v \in P(v)$ is implemented by self-catalysis of r_v by its own product, and an empty row by a reaction with no catalyst, which then belongs to no RAF. The construction is structural: it asserts nothing about kinetics or thermodynamics.

The companion paper proves that an interior operator ψ on E is the maxRAF operator of some CRS with reaction set E (a *same-ground realization*) if and only if $\text{Fix}(\psi) = \mathcal{A} \cap \mathcal{C}(P)$ for an antimatroid $\mathcal{A}$ on E and a predecessor map P [2]; here an antimatroid is a union-closed family containing $\emptyset$ in which every nonempty member has an element whose removal leaves a member [10]. Theorem 2.2 removes the graph from that statement.

Corollary 10.2 (Intrinsic RAF factorization). *An interior operator ψ on E has a same-ground RAF realization if and only if $\text{Fix}(\psi) = \mathcal{A} \cap \mathcal{H}$ for some antimatroid $\mathcal{A}$ on E and some family $\mathcal{H} \subseteq 2^E$ admitting a successful legal interval-deletion sequence.*

Proof. Replace $\mathcal{C}(P)$ by an intrinsically characterized $\mathcal{H}$ using Theorem 2.2 in both directions; the converse direction of the companion theorem supplies the CRS from $\mathcal{A}$ and the reconstructed P . $\square$

The corollary is a factorization statement; it asserts neither uniqueness of the factors nor an efficient way of finding them.

10.3 A conjunction that food generation supplies

Proposition 10.3. *The family $\mathcal{F}_\wedge = \{\emptyset, \{a\}, \{b\}, \{a, b\}, \{a, b, c\}\}$ of Example 3.3 is the fixed family of a CRS with reaction set $\{a, b, c\}$, but of no elementary CRS on that reaction set.*

Proof. Take food $\{f\}$, molecules x, y, z , reactions $a : f \rightarrow x$, $b : f \rightarrow y$, $c : x + y \rightarrow z$, and let f catalyse all three. Catalysis is then automatic, and a set is food-generated iff it does not contain c without both a and b ; the nonempty such sets are exactly the nonempty members of $\mathcal{F}_\wedge$. The second claim is Theorem 10.1 with Example 3.3. Alternatively, $\mathcal{F}_\wedge$ is an antimatroid (every nonempty member loses some element and stays a member), and the companion theorem with the loop digraph realizes it. $\square$

Thus the difference between elementary and general systems is already visible on three reactions: predecessor rows encode alternatives, food generation encodes conjunctions, and by Corollary 9.2(ii) no number of hidden elementary reactions can substitute for the latter under coordinate projection.

11 Formal verification and computational checks

11.1 The formal development

The results of this paper are formalized in Lean 4 (toolchain `leanprover/lean4:v4.30.0`) over a pinned Mathlib [3, 15], in the namespace `DigraphRealizability` spread over thirteen modules. Sets are finite sets over an arbitrary finite type with decidable equality. The support predicate is the existing predicate `RAFInteriorRealizability.PredSupported` of the companion development [2], so that the digraph and RAF results share one semantics. The interval-deletion procedure is the inductive predicate `Peels`, whose constructors are exactly conditions (i)–(ii) of Definition 2.1, and `IntrinsicPeelable F` states that a legal sequence from the full power set succeeds. The exported characterization

`digraphRealizable_iff_intrinsicPeeling`

assumes only a finite type, decidable equality and the target family. It has no hypothesis about normalizers, completions or supplied traces; those objects appear only inside the proof.

Table 1 lists the correspondence between the statements of this paper and the formal declarations. Every module compiles in a strict verification with warnings treated as errors, and the exported root declarations depend only on Lean’s standard axioms (propositional extensionality, classical choice and quotient soundness). No `sorry`, native computation, or finite case enumeration is used to discharge any universal statement.

11.2 What is proved only on paper

The following statements are established by the ordinary proofs in this paper and are not separately compiled declarations: the node bound (11); Proposition 7.2; Corollary 8.2; the cardinality inequality in Proposition 8.5, which is the immediate counting consequence of the compiled injection; the explicit three-reaction system in Proposition 10.3, whose formal counterpart is the existence statement via the antimatroid construction of [2]; Examples 8.3, 8.4 and 8.7; and the worked examples of Section 3, which were checked by exhaustive enumeration.

11.3 Computational checks

Before the formalization, the criterion was compared with a brute-force oracle on every labelled digraph with at most four vertices, in both the direct and the complementary formulation.

n	labelled digraphs (2^{n^2})	union-closed families containing $\emptyset$	digraph-realizable families
0	1	1	1
1	2	2	2
2	16	7	7
3	512	61	55
4	65 536	2 480	1 359

There were no disagreements in classification, in reconstruction, or in the length of successful sequences. For $n = 3$ the six non-realizable union-closed families are, up to relabelling, $\mathcal{F}_\wedge$ (three labellings), $\mathcal{F}_{\text{loc}}$ (three labellings). Formula (12) was compared with direct projection

Statement	Formal declaration (module)
Least models: (5) is a model containing A , below every such model	<code>subset_closure, closure_le, closure_models</code> (HornClosure)
Lemma 4.1, existence of $A \rightarrow r$ with (6)	<code>normalize</code> (HornClosure)
Lemma 4.1, replacement and completion extension	<code>replacement, extend_completion</code> (ResidualNormalization)
Completeness and soundness in complementary form	<code>elimination_complete, elimination_sound, singleHead_iff_elimimates</code> (HornElimination)
Duality (7) and reconstruction (4)	<code>support_iff_models, reconstruct_models, reconstructed_support</code> (ComplementBridge)
Translation between the direct and complementary procedures	<code>Eliminates.toPeels, Peels.toEliminates</code> (IntervalDeletion)
Theorem 2.2, (1) $\Leftrightarrow$ (2)	<code>digraphRealizable_iff_intrinsicPeeling</code> (IntervalDeletion)
Eligibility equals (3) for the interior of a union-closed family	<code>eligible_iff_interior</code> (IntervalDeletion)
Theorem 2.2, (1) $\Rightarrow$ (3) for an arbitrary valid policy	<code>policy_complete</code> (ResidualPolicy)
Executable trace checker and its equivalence with the criterion	<code>checkTrace, digraphRealizable_iff_checked_trace</code> (TraceChecker)
Proposition 6.1, Corollary 6.2	<code>exact_trace_roots, exact_trace_length, realizing_loops</code> (TraceConsequences)
Interior operator of the reconstruction equals that of $\mathcal{F}$	<code>reconstructed_interior_eq</code> (RAFConsequences)
Theorem 10.1	<code>elementary_iff_peeling</code> (RAFConsequences)
Corollary 10.2	<code>raf_iff_antimatroid_peeling</code> (RAFConsequences)
Theorem 9.1; Corollary 9.2(i)–(iii)	<code>single_vertex_projection, arbitrary_projection, source_projection_on_visible, hidden_vertices_cannot_rescue, projected_interior</code> (Projection, PublicationResolution)
Equation (13); independent products	<code>induced_deletion, independent_product</code> (Composition)
Proposition 8.1; Proposition 10.3, existence part	<code>conjunction_intersection, conjunction_not_realizable, conjunction_general_realizable</code> (Composition)
Proposition 8.5, the injection ρ	<code>residual_matching</code> (Obstructions)
Aggregate export	<code>publication_resolution</code> (PublicationResolution)

Table 1: Concordance between the paper and the Lean development.

on all 263 714 pairs (digraph, vertex) with at most four vertices, and the disjoint-union construction on all 256 products of two two-vertex digraphs, without mismatch. These finite checks test conventions and implementations; the universal statements rest on the proofs and the formalization.

12 Discussion and open problems

Theorem 2.2 answers Steel’s first question with a criterion that uses only membership, inclusion, maximality and distinctness of roots, and it does so with an exact reconstruction on the original ground set. The residual set may be chosen freely; the price of that freedom is that the root

must be searched, and Example 3.1 shows the search is genuinely needed. Section 8 explains why a static description by closure axioms is impossible, and Sections 9–10 show that the class behaves well under projection, deletion and products, that the missing closure under intersection is exactly what food generation adds in a general catalytic reaction system, and that the graph factor of the same-ground RAF characterization is now intrinsic.

Several questions remain.

1. *Complexity.* Is digraph-realizability of an explicitly listed family decidable in polynomial time? Proposition 7.2 places it in NP; Proposition 7.1 gives a search exponential only in the number of excluded singletons.
2. *Static characterizations.* Corollary 8.2 rules out closure rules. Is there a characterization by forbidden configurations, or by a condition with negative hypotheses that is checkable locally?
3. *Identifiability.* The family determines all loops (Corollary 6.2) and, in Examples 3.1 and 3.2, most arcs. Which arcs are determined by $\mathcal{F}$ in general, and is there a canonical (for instance arc-minimal) realization computable from a successful trace?
4. *Obstructions.* The Hall condition of Corollary 8.6 is necessary and Example 8.7 shows it is not sufficient. Is there a complete and interpretable family of obstruction certificates, in the spirit of the derivation-connectivity argument of that example?
5. *Beyond elementary systems.* Corollary 10.2 is existential in the antimatroid factor. Given a same-ground RAF family, how can the antimatroid and the digraph factors be found, and when are they unique?

References

- [1] Dana Angluin, Michael Frazier, and Leonard Pitt. Learning conjunctions of Horn clauses. *Machine Learning*, 9(2–3):147–164, 1992. doi:10.1007/BF00992675.
- [2] Anonymous. Interior operators realizable by autocatalytic networks: an antimatroid characterization. Companion manuscript, 2026; arXiv preprint in preparation, 2026.
- [3] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi:10.1007/978-3-030-79876-5_37.
- [4] William F. Dowling and Jean H. Gallier. Linear-time algorithms for testing the satisfiability of propositional Horn formulae. *Journal of Logic Programming*, 1(3):267–284, 1984. doi:10.1016/0743-1066(84)90014-1.
- [5] Paul H. Edelman and Robert E. Jamison. The theory of convex geometries. *Geometriae Dedicata*, 19(3):247–270, 1985. doi:10.1007/BF00149365.
- [6] Delbert R. Fulkerson and Oliver A. Gross. Incidence matrices and interval graphs. *Pacific Journal of Mathematics*, 15(3):835–855, 1965. doi:10.2140/pjm.1965.15.835.
- [7] Bernhard Ganter. Two basic algorithms in concept analysis. In Léonard Kwuida and Barış Sertkaya, editors, *Formal Concept Analysis (ICFCA 2010)*, volume 5986 of *Lecture Notes in Computer Science*, pages 312–340. Springer, Berlin, 2010. doi:10.1007/978-3-642-11928-6_22. Reprint of the 1984 Technische Hochschule Darmstadt preprint FB4-831.

- [8] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004. doi:[10.1016/j.jtbi.2003.11.020](https://doi.org/10.1016/j.jtbi.2003.11.020).
- [9] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi:[10.1016/S0022-5193\(86\)80047-9](https://doi.org/10.1016/S0022-5193(86)80047-9).
- [10] Bernhard Korte, László Lovász, and Rainer Schrader. *Greedoids*, volume 4 of *Algorithms and Combinatorics*. Springer, Berlin, 1991.
- [11] Paolo Liberatore. Common equivalence and size after forgetting. arXiv:2006.11152, 2020.
- [12] Paolo Liberatore. Reconstructing a single-head formula to facilitate logical forgetting. *Journal of Logic and Computation*, 33(7):1596–1624, 2023. doi:[10.1093/logcom/exac086](https://doi.org/10.1093/logcom/exac086).
- [13] Mike Steel. Interior operators and their relationship to autocatalytic networks. *Acta Biotheoretica*, 71(4):21, 2023. doi:[10.1007/s10441-023-09472-8](https://doi.org/10.1007/s10441-023-09472-8). arXiv:2305.16714.
- [14] Mike Steel, Wim Hordijk, and Joana C. Xavier. Autocatalytic networks in biology: structural theory and algorithms. *Journal of the Royal Society Interface*, 16(151):20180808, 2019. doi:[10.1098/rsif.2018.0808](https://doi.org/10.1098/rsif.2018.0808).
- [15] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi:[10.1145/3372885.3373824](https://doi.org/10.1145/3372885.3373824).
- [16] Marcel Wild. The joy of implications, aka pure Horn formulas: Mainly a survey. *Theoretical Computer Science*, 658:264–292, 2017. doi:[10.1016/j.tcs.2016.03.018](https://doi.org/10.1016/j.tcs.2016.03.018).