Execution and Observation: Two Geometries of Behaviour

Finite experiments, operational continuation, and reachability–observability pairings

Research companion to *From Behaviour Categories to Generalized Spaces*

21 October 2026 – exploratory mathematical note

Abstract

Spivak associates to a polynomial interface a free monad of finite trees, a cofree comonad of potentially infinite behaviours, and a finite-test topology on its terminal behaviour space. We make explicit how this topology arises from Libkind–Spivak's pattern–matter module action and a matching map with nontrivial residual syntax. Fairbanks–Carlson–Spivak (FCS) extract a different topology from the same comonad, with opens closed under execution. For Moore machines we identify the exact bridge: the reachable–observable minimal realization of a pointed machine is the *least FCS-open neighbourhood* of its initial behaviour, and the restrictions of finite-test evaluations to this neighbourhood are the columns of its Hankel pairing. Execution maps preserve finite-test continuity, but, for noncommuting inputs, need not be continuous in the FCS topology. We compare the two inverse-limit presentations suggested by binary and 2-adic expansions, recover Boolean polarities and classical automata minimization, and examine writer and finite-nondeterministic residues. The latter yields a standard failure of surjectivity at stage ω of the final chain; its image-finite behaviour space is a dense noncompact subspace of the finite-test Stone completion. General residual experimental interactions remain a research problem.

1 Why a second geometry of behaviour?

An operational theory tells us which experiments or commands can be performed; a runner realizes these operations and supplies their observable consequences. The associated Sweedler comonad

$$C_{T,R} = \{T, R\}$$

organizes implementations of the source theory T with residual effects in R. The preceding companion note calculates the topology, underlying category, and halos extracted from such comonads by Fairbanks–Carlson–Spivak (FCS) [3]; Garner's behaviour categories furnish the related categorical semantics of comodels [5]. A characteristic feature of its topology is *forward invariance*: if a behaviour belongs to an open set, so do its operational continuations.

Spivak's seminar notes [1] suggest a different geometry. For a polynomial p, his free monad $\mathfrak{m}_p$ contains finite experimental trees, while the cofree comonad $\mathfrak{c}_p$ describes potentially infinite behaviours. Finite trees generate a topology on $\mathfrak{c}_p(1)$, with finite observations specifying neighbourhoods. Spivak associates two toposes to p and uses the phrases *singular experience* and *counterfactual plans* for their respective interpretations. These phrases motivate, but do not formally identify, the resulting logics.

Thus even when the underlying *set* of behaviours is the same, we may have

$$\boxed{\underbrace{\tau_{\text{test}}}_{\text{agreement under finite tests}} \quad \text{and} \quad \underbrace{\tau_{\text{exec}}}_{\text{closure under execution}} .} \tag{1}$$

One should not expect either topology to determine the other. In the Moore case, however, the smallest operationally open neighbourhood of the initial behaviour is *exactly* the minimal realization, while the observational tests separate its points (Theorem 4.1). This joins the geometries without identifying their distinct topologies.

We first establish the relevant facts for deterministic Moore machines. We then consider the classical linear and Boolean dualities, the direction of the limiting processes, and two residual theories. We distinguish the mathematical consequences of the stated data from proposals for a general theory of experiments on residual runners.

2 Spivak's syntax of finite experiments

A polynomial endofunctor is expressed as

$$p(X) = \coprod_{i \in p(1)} X^{p[i]}.$$

Its coalgebras are dynamical systems with a shape i of available responses and a continuation for each position in $p[i]$. Spivak considers two related constructions in the monoidal category of polynomials [1]:

$$\begin{aligned} \mathfrak{m}_p &\simeq y + p \triangleleft \mathfrak{m}_p \quad \text{(well-founded, finite-depth } p\text{-trees)}, \\ \mathfrak{c}_p &\simeq y \times (p \triangleleft \mathfrak{c}_p) \quad \text{(cofree comonad of } p\text{-behaviours)}. \end{aligned} \tag{2}$$

Here y is the substitution unit and $\triangleleft$ denotes polynomial substitution. The displayed isomorphisms are interpreted within the relevant free/cofree constructions, not as equations for an arbitrary polynomial p.

The underlying set $\mathfrak{c}_p(1)$ is the terminal p-coalgebra. A finite well-founded tree $u \in \mathfrak{m}_p(1)$ determines a predicate on the terminal behaviours: the predicate asserts that the behaviour agrees with the finite shape and labels specified by u. The resulting basic opens generate Spivak's finite-test topology τ_{test}.

Spivak also describes a more structured *Gendlin interaction*

$$\mathfrak{m}_p \longrightarrow [\mathfrak{c}_p, \mathfrak{m}_{p+1}], \tag{3}$$

where $\mathfrak{m}_{p+1}$ is the *free monad on the polynomial* $p+1$, with an extra nullary failure constructor. It is not the endofunctor $\mathfrak{m}_p + 1$. We can now relate this interaction to the explicit "pattern runs on matter" construction of Libkind and Spivak [2].

2.1 The module action behind finite-tree tests

For arbitrary polynomials p, q, Libkind–Spivak construct a natural transformation

$$\boxed{\Xi_{p,q} : \mathfrak{m}_p \otimes \mathfrak{c}_q \longrightarrow \mathfrak{m}_{p \otimes q}.} \tag{4}$$

Their Theorem 3.4 establishes the module coherence laws. Here $\otimes$ is the *Dirichlet* product of polynomials, not their substitution product $\triangleleft$. A map $\varphi : p \otimes q \to r$ induces the composite

$$\mathfrak{m}_p \otimes \mathfrak{c}_q \xrightarrow{\Xi_{p,q}} \mathfrak{m}_{p \otimes q} \xrightarrow{\mathfrak{m}_\varphi} \mathfrak{m}_r, \tag{5}$$

which runs a finite p-pattern on a potentially infinite q-behaviour.

For matching, there is a particularly simple polynomial map. Write $p = \sum_{i \in I} y^{E_i}$. Its Dirichlet square has positions (i, j) and directions $E_i \times E_j$. Define

$$\chi_p : p \otimes p \longrightarrow p + 1 \tag{6}$$

by sending a matching pair (i, i) to $i \in I$ and a mismatched pair (i, j), $i \neq j$, to the new nullary position $\perp$. In the matching case the backwards map on directions is the diagonal $E_i \to E_i \times E_i$; in the mismatched case it is the unique map $\varnothing \to E_i \times E_j$.

Proposition 2.1 (Matching as a pattern–matter interaction)**.** *The composite*

$$\mathfrak{m}_p \otimes \mathfrak{c}_p \xrightarrow{\Xi_{p,p}} \mathfrak{m}_{p \otimes p} \xrightarrow{\mathfrak{m}_{\chi_p}} \mathfrak{m}_{p+1} \tag{7}$$

is the finite-tree matching interaction described in Spivak's seminar notes: it retains agreeing positions and produces $\perp$ on any mismatched branch. Its transpose is the Gendlin map (3).

Proof. At a node, Ξ combines the position chosen by the finite pattern with the position supplied by the current behaviour and recursively continues along paired directions. The map χ_p retains the common position and follows the diagonal direction precisely when the two positions agree; otherwise it terminates that branch with $\perp$. Induction over the well-founded pattern tree gives the matching recursion. Transposition under Dirichlet closure yields the displayed internal-hom map, with the monad-map coherence inherited from the module action and the free-monad extension. □

Interpret terminated pattern leaves as true, the extra nullary constructor $\perp$ as false, and a p-node as conjunction of its branch results. The induced predicate on complete behaviours is

$$\mathrm{match} : \mathfrak{m}_p(1) \times \mathfrak{c}_p(1) \longrightarrow 2,$$

and the basic open selected by $u \in \mathfrak{m}_p(1)$ is $[u] = \{b : \mathrm{match}(u, b) = 1\}$. For $p = 2y$, this reduces to the familiar prefix relation

$$\mathfrak{m}_{2y}(1) \cong 2^*, \qquad \mathfrak{c}_{2y}(1) \cong 2^{\mathbb{N}}, \qquad \mathrm{match}(u, b) = 1 \iff u \text{ is a prefix of } b.$$

Thus the finite-test topology is induced by a concrete monad–comonad interaction, not merely by comparing finite and infinite trees as sets. This is an interpretation of the constructions in [1, 2], not a claim that every Sweedler runner has a canonical collection of finite tests.

2.2 Sweedler duality and the distinction between syntaxes

There is a second, different, way in which $\mathfrak{m}_p$ meets a cofree comonad. In the polynomial duoidal setting, freeness, Dirichlet closure, and cofreeness give the natural isomorphism

$$\boxed{\{\mathfrak{m}_p, R\} \;\cong\; \mathfrak{c}_{[p,R]}} \tag{8}$$

for a polynomial residual monad R, under the relevant Sweedler existence hypotheses. This is the free-monoid case of the general Sweedler calculation (McDermott–Rivas–Uustalu, Proposition 2(i)) [4]. In particular,

$$\{\mathfrak{m}_p, y\} \cong \mathfrak{c}_{[p,y]}, \qquad \{T_{\mathrm{Moore}}(A, O), y\} \cong \mathfrak{c}_p \quad (p = Oy^A). \tag{9}$$

The second isomorphism is established by identifying Moore comodels with p-coalgebras over **Set**. For $p = Oy^A$, the Dirichlet dual of the interface is $[p, y] \cong A^O y^O$, whereas the Moore

operational theory is the quotient of the free monad on $Ay + y^O$ by the reader equations. It is not generally $\mathfrak{m}_p$.

The trivial-residue comparison can also be obtained directly from (4) using Dirichlet evaluation $\mathrm{ev} : p \otimes [p, y] \to y$:

$$\mathfrak{m}_p \otimes \mathfrak{c}_{[p,y]} \xrightarrow{\Xi_{p,[p,y]}} \mathfrak{m}_{p\otimes[p,y]} \xrightarrow{\mathfrak{m}_{\mathrm{ev}}} \mathfrak{m}_y \longrightarrow y.$$

The last arrow forgets the number of unary steps. For a general polynomial residual monad R, replace evaluation into y by evaluation into R and the last arrow by the counit $\mathfrak{m}_R \to R$. This is the universal-answerer interaction discussed in [2, Section 4.1, equation (6)].

The matching construction, by contrast, has *nontrivial residue*:

$$\mathfrak{c}_p \longrightarrow \{\mathfrak{m}_p, \mathfrak{m}_{p+1}\} \cong \mathfrak{c}_{[p,\mathfrak{m}_{p+1}]}. \tag{10}$$

Its first arrow is induced by (7) and the Sweedler measuring adjunction. It is not asserted to be an isomorphism. By contrast, (9) with residue y uses the *universal answerer* $[p, y]$, not p itself. Thus Sweedler duality and finite-tree matching are two specializations of a common pattern–matter action, but they answer different questions.

Remark 2.2 (Two syntaxes, one category of implementations). For the Moore interface $p = Oy^A$, a p-coalgebra is a Moore machine. Its cofree comonad agrees with $\{T_{\mathrm{Moore}}(A, O), \mathrm{id}\}$ in the setting of our earlier note. But the free *polynomial* monad $\mathfrak{m}_p$ is not the operational Moore theory $T_{\mathrm{Moore}}(A, O)$: the latter has an O-ary getter satisfying the reader equations and freely generated unary updates. The agreement $\mathfrak{c}_p \cong \{T_{\mathrm{Moore}}(A, O), y\}$ comes from the equivalence of p-coalgebras and Moore comodels over **Set**, not from equality of operational syntaxes. In particular, the Sweedler dual of $\mathfrak{m}_p$ is instead $\mathfrak{c}_{[p,y]}$, as in (9). A translation between the two kinds of experiments must therefore be constructed, not silently assumed.

3 Moore behaviours, execution, and testing

Let A be an input alphabet and O an output set. All the Moore execution and observation identities below hold for arbitrary A, O. We give O the discrete topology; finiteness of O alone is enough for O^{A^*} to be Stone. Additional finite-stage assertions will say when finiteness of A is also needed. Write A^* for the free monoid of words, with uv meaning *first execute u, then v*.

A Moore machine is $(S, o, (\delta_a)_{a\in A})$ with

$$o : S \to O, \qquad \delta_a : S \to S.$$

Its unique map to the terminal Moore coalgebra is

$$\mathrm{obs}(s)(w) = o\big(\delta_w(s)\big), \qquad B := O^{A^*}.$$

For $u \in A^*$, define the residual/continuation operator and the experiment e_v by

$$\sigma_u : B \to B, \quad (\sigma_u b)(v) = b(uv), \qquad e_v : B \to O, \quad e_v(b) = b(v). \tag{11}$$

The family (σ_u) is an action in the appropriate word-composition convention: $\sigma_v\sigma_u = \sigma_{uv}$. The e_v are all the evaluations of a behaviour.

Theorem 3.1 (Execution–experiment compatibility). *For all $u, v \in A^*$,*

$$\boxed{e_v \circ \sigma_u = e_{uv}.} \tag{12}$$

Equip O with the discrete topology and B with the initial topology generated by all e_v. Then each σ_u is continuous.

Proof. For $b \in B$, $e_v(\sigma_u b) = (\sigma_u b)(v) = b(uv) = e_{uv}(b)$. A subbasic open of the initial topology is $e_v^{-1}(V)$ for $V \subseteq O$, and

$$\sigma_u^{-1}(e_v^{-1}(V)) = e_{uv}^{-1}(V),$$

which is again subbasic open. □

3.1 The observational topology

For a finite set $F \subseteq A^*$ and assignment $q : F \to O$, define the cylinder

$$N(F, q) = \{b \in B \mid b(w) = q(w) \text{ for all } w \in F\}.$$

The nonempty such cylinders are clopen and form a basis for τ_{test}. When A and O are finite, each finite stage of observation is finite, so

$$B \cong \varprojlim_n B_n, \qquad B_n = O^{A^{<n}}, \qquad A^{<n} = \{w : |w| < n\}. \tag{13}$$

The bonding maps restrict labels from $A^{<n+1}$ to $A^{<n}$. If A and O are finite, the B_n are finite discrete spaces. More generally, when O is finite (even if A is infinite), each B_n is a product of finite discrete spaces and is Stone, and the inverse limit has the product topology. Thus $B = O^{A^*}$ is Stone whenever O is finite, by Tychonoff's theorem. This is Spivak's topology for the Moore interface $p = Oy^A$: a finite full p-tree can test the output at each node of a finite prefix-closed region of A^*, with untested branches terminated by leaves.

3.2 The operational FCS topology

For the cofree Moore comonad

$$C(X) \cong O^{A^*} \times X^{A^*},$$

the FCS topology on $B = C1$ has opens exactly the subsets closed under every residual σ_u (by the subcoalgebra definition [3]). Equivalently, put

$$b \preceq c \iff (\exists u \in A^*)\ c = \sigma_u b. \tag{14}$$

Then the opens of τ_{exec} are the upward-closed sets for this reachability preorder. We fix the specialization convention $b \preceq c$ when *every open containing b also contains c*; this is the opposite of another common convention.

Proposition 3.2 (The two topologies are incomparable)**.** *For $A = \{a\}$ and $O = \{0, 1\}$, the topologies τ_{test} and τ_{exec} on $B = 2^{\mathbb{N}}$ are incomparable. The latter is not T_0.*

Proof. The stream 0^ω is fixed by the shift, so $\{0^\omega\}$ is τ_{exec}-open. It is not τ_{test}-open, since no finite prefix singles out 0^ω. Conversely, the cylinder $\{b : b(0) = 0\}$ is τ_{test}-open, but it is not τ_{exec}-open: a stream beginning 01 belongs to it while its shift does not. Finally, $0101\ldots$ and $1010\ldots$ are distinct, mutually reachable periodic streams, so every τ_{exec}-open containing either contains the other. Hence τ_{exec} is not T_0. □

Remark 3.3 (Two kinds of neighbourhood). A τ_{test}-neighbourhood expresses finite information about a complete behaviour; an τ_{exec}-neighbourhood expresses forward closure under operational continuation. The latter is directional, not a metric notion. Theorem 3.1 asserts continuity of execution only for the *test* topology. Execution need not be continuous for the topology built from execution itself, as the next proposition shows.

Proposition 3.4 (Execution need not be FCS-continuous)**.** *For $A = \{a, c\}$ and $O = 2$, the continuation $\sigma_a : B \to B$ is τ_{test}-continuous but not τ_{exec}-continuous. For a one-letter alphabet, by contrast, all continuations are τ_{exec}-continuous.*

Proof. Define $b(w) = 1$ exactly when w begins with the two-letter word ca. We have $b \preceq \sigma_c b$, whereas $\sigma_a b = 0^{A^*}$ is the constant-zero behaviour and $\sigma_a(\sigma_c b) = \sigma_{ca} b = 1^{A^*}$ is the constant-one behaviour. The latter is not reachable from 0^{A^*}. Thus σ_a fails to preserve the specialization preorder $\preceq$, and hence is not continuous for the Alexandrov topology. Test continuity follows from Theorem 3.1. If A has one generator, all residuals commute; hence $\sigma_u(\sigma_w b) = \sigma_w(\sigma_u b)$, making every σ_u monotone for reachability and thus Alexandrov-continuous. □

Remark 3.5 (Binary versus 2-adic approximation). Let $B_n = 2^n$. Deleting the rightmost digit defines prefix bonding maps $B_{n+1} \to B_n$; deleting the leftmost defines suffix bonding maps. Both give *inverse* systems, isomorphic by reversal of each finite word. The prefix limit consists of right-infinite binary streams, while the suffix limit consists of left-infinite expansions $\mathbb{Z}_2 \cong \varprojlim_n \mathbb{Z}/2^n\mathbb{Z}$. Thus the two Cantor topologies preserve different ends of written strings. For example, $1101011_2 - 101011_2 = 2^6$, so these integers are 2^{-6} apart in the 2-adic metric because their six lowest digits agree. Prefix agreement, by contrast, underlies finite testing of Moore output streams.

The shift σ acts on the already constructed inverse limit; it is not itself one of the prefix bonding maps. Operational reachability from b is the union $\uparrow b = \bigcup_n R_n(b)$ of finite-horizon reachable sets $R_n(b) = \{\sigma_u b : |u| \le n\}$ (for finite A). This is a directed colimit of inclusions, but not the colimit of $B \xrightarrow{\sigma} B \xrightarrow{\sigma} \cdots$, which would forget the direction of reachability. The distinction is between *finite approximation of a behaviour* and *generation of its continuations*, not a real inverse limit versus a 2-adic direct limit.

4 The Hankel pairing and minimal realization

Fix a pointed Moore machine (S, s_0, o, δ). Define its reachability and observation maps

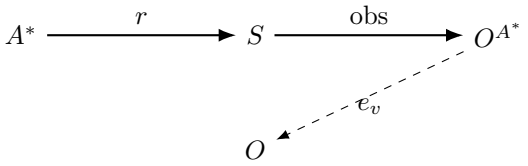
$$r : A^* \to S, \quad r(u) = \delta_u(s_0), \qquad \text{obs} : S \to O^{A^*}, \quad \text{obs}(s)(v) = o(\delta_v s). \tag{15}$$

Their evaluation pairing is the Hankel function

$$\boxed{H : A^* \times A^* \longrightarrow O, \qquad H(u, v) = o\big(\delta_{uv}(s_0)\big).} \tag{16}$$

One variable is a *history* leading to a state; the other is a *future test* performed from it. The fundamental factorization is

$$H(u, v) = e_v\big(\text{obs}(r(u))\big). \tag{17}$$



Theorem 4.1 (Minimal realization is the minimal FCS-open neighbourhood)**.** *For a pointed Moore machine (S, s_0, o, δ), let $b_0 = \text{obs}(s_0) \in B$, and let $S_{\text{reach}} = \{\delta_u(s_0) : u \in A^*\}$. On this reachable submachine set $s \sim_{\text{obs}} t$ exactly when $\text{obs}(s) = \text{obs}(t)$. Then there is a canonical isomorphism of* pointed Moore machines

$$\boxed{S_{\min} = S_{\text{reach}}/{\sim_{\text{obs}}} \;\cong\; \{\sigma_u b_0 : u \in A^*\} = \uparrow b_0 \subseteq B.} \tag{18}$$

The set on the right is the smallest τ_{exec}-open neighbourhood of b_0, equivalently the subcoalgebra of the final Moore coalgebra generated by b_0. Moreover the evaluations $e_v|_{\uparrow b_0}$ separate its states, and their values on reachable points recover the Hankel columns:

$$H(u, v) = e_v(\sigma_u b_0). \tag{19}$$

Every reachable observable pointed Moore machine with initial behaviour b_0 is isomorphic to this one.

Proof. The equivalence $\sim_{\mathrm{obs}}$ is a Moore congruence: equivalent states have the same output by the empty-word test, and $\mathrm{obs}(\delta_a s)(v) = \mathrm{obs}(s)(av)$ shows stability under all updates. Hence the quotient of the reachable submachine exists. The final-behaviour map satisfies $\mathrm{obs}(\delta_u(s_0)) = \sigma_u b_0$. Consequently its restriction to S_{reach} has image $\uparrow b_0$ and kernel $\sim_{\mathrm{obs}}$, giving the displayed isomorphism by image factorization in **Set**. This image is a subcoalgebra, since $\sigma_a \sigma_u b_0 = \sigma_{ua} b_0$. It is the least forward-invariant subset containing b_0, exactly the least τ_{exec}-open neighbourhood. Distinct functions $b, c : A^* \to O$ differ at some v, and thus the restricted evaluations separate points. Formula (19) follows from Theorem 3.1. A second reachable observable pointed machine with the same initial behaviour maps bijectively to this image, proving uniqueness. □

This is the familiar minimal-realization construction as the reachable subcoalgebra of the final behaviour object, expressed in the FCS topology. The image-factorization viewpoint belongs to classical categorical realization theory [7, 8, 6]; the *topological restatement* follows because FCS opens in this polynomial case are precisely subcoalgebras.

Corollary 4.2 (Finite orbit and Myhill–Nerode)**.** *A Moore behaviour $b \in O^{A^*}$ admits a finite-state realization if and only if its smallest τ_{exec}-open neighbourhood $\uparrow b$ is finite. If A is finite, $O = 2$, and b is the characteristic function of a language $L \subseteq A^*$, then L is regular if and only if $|\uparrow b| < \infty$; its points are the distinct left derivatives $u^{-1}L$.*

Proof. Any realization has at least as many states as its distinct reachable behaviours. Conversely $\uparrow b$, with restricted final Moore structure and initial point b, is a realization of b. In the Boolean case $\sigma_u b$ is the characteristic function of $u^{-1}L$, giving the Myhill–Nerode criterion. □

Corollary 4.3 (The restricted experimental pairing)**.** *The restrictions $e_v|_{\uparrow b_0} : \uparrow b_0 \to O$ constitute a separating family of experimental predicates (or observations) on the minimal machine. Composing them with the orbit map $u \mapsto \sigma_u b_0$ gives exactly the columns $H(-, v)$ of the Hankel pairing. Hence the two geometries meet on the same minimal realization, without becoming the same topology.*

5 Boolean polarity, automata, and linear duality

5.1 A precise polarity

Suppose $O = 2 = \{0, 1\}$. Let $U, V \subseteq A^*$ be sets of histories and future tests, respectively. Define

$$U^{\perp} = \{v : \forall u \in U,\ H(u, v) = 1\}, \qquad V^{\perp} = \{u : \forall v \in V,\ H(u, v) = 1\}. \tag{20}$$

The notation is overloaded on the two sides, but their domains make it unambiguous.

Proposition 5.1 (Galois duality)**.** *The operators in* (20) *are antitone and satisfy*

$$U \subseteq V^{\perp} \iff V \subseteq U^{\perp}.$$

Consequently $U \mapsto U^{\perp\perp}$ *and* $V \mapsto V^{\perp\perp}$ *are closure operators, and the lattices of polar-closed history sets and polar-closed test sets are order-dual.*

Proof. Both inclusions assert precisely that $H(u, v) = 1$ for every $(u, v) \in U \times V$. The remaining claims are the standard consequences of an antitone Galois connection: extensivity and idempotence of double polar, and the mutually inverse antitone bijections on fixed points. □

This is the standard *formal-concept-analysis polarity* of the Boolean Hankel relation [11]. It is an actual duality, although not an equivalence of the full powersets. The polarity depends on which outcomes count as accepting; for arbitrary O one can replace the value 1 with a chosen predicate on O or an enriched outcome object.

5.2 Languages and Brzozowski derivatives

If the initial behaviour represents a Boolean language $L \subseteq A^*$, then $H(u, v) = 1$ iff $uv \in L$. Its u-row is the left derivative

$$u^{-1}L = \{v : uv \in L\}.$$

Therefore Theorem 4.1 identifies the minimal deterministic automaton with the reachable derivatives. Right contexts supply the columns, and reversal switches the roles of left and right derivatives. The derivative construction is classical [10]; its combination with reversal is the basis of Brzozowski minimization. Bonchi, Bonsangue, Hansen, Panangaden, Rutten and Silva explicitly formulate reachability algebraically and observability coalgebraically, and derive reversal through a duality [9]. Thus the algebra–coalgebra duality is an established precedent, not an open possibility first suggested by this note. What remains open is whether an analogous mechanism arises canonically for general residual Sweedler homs.

5.3 Comparison with Kalman duality

For a finite-dimensional discrete-time linear system

$$x_{t+1} = Ax_t + Bu_t, \qquad y_t = Cx_t,$$

the reachability and observability matrices (with suitable truncation) are

$$\mathcal{R} = [B \; AB \; A^2B \; \cdots], \qquad \mathcal{O} = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \end{bmatrix}, \qquad \mathcal{H} = \mathcal{O}\mathcal{R},$$

where the (i, j) block of $\mathcal{H}$ is $CA^{i+j}B$. Under transpose dualization,

$$(A, B, C) \longmapsto (A^{\mathsf{T}}, C^{\mathsf{T}}, B^{\mathsf{T}}),$$

reachability and observability exchange roles. This stronger duality relies on finite-dimensional linear duals. The general Moore factorization in **Set** does not have an analogous transpose unless one supplies additional dualizing structure. Classical categorical work of Arbib–Manes, and the modern account of Bonchi et al., places this asymmetry in a wider realization framework [7, 9]. The Boolean polarity is one possible replacement, but not automatically the same as Kalman's dual system.

6 Modalities, finite evidence, and safety

The operational topology induces an interior operator on subsets of B:

$$\Box_{\mathrm{exec}} P = \{b : \forall u \in A^*,\ \sigma_u b \in P\} = \bigcap_{u \in A^*} \sigma_u^{-1}(P). \tag{21}$$

Because reachability is reflexive and transitive, this is an S4 interior operator. The observational topology induces a different operator

$$\Box_{\mathrm{test}} P = \mathrm{int}_{\tau_{\mathrm{test}}}(P).$$

The latter says that a finite collection of observations suffices to certify membership in P. The two modal operators are not generally equal.

Example 6.1 (Finite checks and an unbounded safety assertion)**.** Let $B = 2^{\mathbb{N}}$ and $P = \{b : b(0) = 1\}$, a clopen one-step test. Define

$$\Box_{\leq N} P = \bigcap_{n=0}^{N} \sigma^{-n}(P) = \{b : b(0) = \cdots = b(N) = 1\}.$$

Each $\Box_{\leq N} P$ is a clopen finite-test condition. But

$$\Box_{\mathrm{exec}} P = \bigcap_{N \geq 0} \Box_{\leq N} P = \{1^\omega\},$$

which is τ_{test}-closed and is not τ_{test}-open. A finite observation can witness a violation at a particular finite stage; no fixed finite horizon verifies that the property persists for all time.

Thus the same system can have a finite-test topology which expresses experimental evidence and an execution modality which expresses global invariance. The logical distinction is useful for safety. But the FCS Alexandrov interior only sees $\to^*$, not the separate one-step modalities $\langle a \rangle$ *and output atoms* of modal logics for Moore systems. Indeed the non-T_0 example in Proposition 3.2 already shows that the topology cannot characterize bisimilarity, although suitable one-step modal formulas do characterize bisimilarity for image-finite systems [12].

7 A first extension: writer residues

Take the standard Moore theory with freely generated updates and writer residue $R = M \times -$, where M is a monoid. As derived in the preceding companion note,

$$C_M(X) = \{T_{\mathrm{Moore}}(A, O), M \times -\}(X) \cong O^{A^*} \times M^{A^* \times A} \times X^{A^*}. \tag{22}$$

The Moore getter must remain pure; the effect of each update edge is an element of M. Hence the point set is

$$B_M = O^{A^*} \times M^{A^* \times A}.$$

If O and M are given the discrete topology, define a finite-test topology on B_M by pointwise evaluations of $b(w)$ and $m(w, a)$. If O and M are finite, B_M is Stone for *any* input alphabet A, by Tychonoff; finiteness of A is only needed to make bounded-depth sets of coordinates finite. If O or M is infinite and discrete, the resulting product topology need not be compact. Execution still acts by residuals:

$$\sigma_u(b, m) = (b_u, m_u), \qquad b_u(v) = b(uv), \quad m_u(v, a) = m(uv, a).$$

Both kinds of evaluation obey the extended-test identity, for example

$$\mathrm{effect}_{v,a}(\sigma_u(b,m)) = \mathrm{effect}_{uv,a}(b,m).$$

Hence the writer action is continuous for its finite-test topology. The FCS topology is, as before, the Alexandrov topology of residual reachability.

Remark 7.1 (What counts as an observation?). Making the M-labels available to the test topology is a *choice*. If writer effects are internal rather than externally observable, only the O-evaluations should be admitted as tests, and behaviours with different effect trees may be observationally indistinguishable. Thus the residual monad alone need not specify the observational logic. Any general construction $\mathrm{Test}_{T,R}$ must state what portion of the residue is observable and under which operations.

8 Finite nondeterminism and the limit obstruction

Now let $R = \mathcal{P}_{\mathrm{f}}$, the finite-powerset monad (including the empty set). The full Moore–nondeterministic Sweedler hom is the cofree comonad on

$$F(S) = O \times (\mathcal{P}_{\mathrm{f}} S)^A,$$

and its point set is the final coalgebra

$$B_{\mathrm{nd}} = \nu Z.\, O \times (\mathcal{P}_{\mathrm{f}} Z)^A. \tag{23}$$

Unlike O^{A^*}, these points record branching structure up to observation-preserving bisimulation; sets of output traces alone need not determine the behaviour. The companion note shows that some of its halos are nonrepresentable, for instance

$$H_b(X) \cong X \times \mathcal{P}_{\mathrm{f}}^+(X)$$

when b has one transition to a terminal successor. Consequently joint observations need not be recoverable from marginal observations.

A plausible finite-test semantics assigns each point its finite-depth unfoldings. But one must not immediately identify the final coalgebra with the inverse limit of *all* finite-depth approximants.

Proposition 8.1 (An inverse limit with extra behaviours)**.** *For $A = \{a\}$, $O = \{0,1\}$, and $F(Z) = 2 \times \mathcal{P}_{\mathrm{f}} Z$, the canonical map from the final F-coalgebra to the limit of the finite final chain*

$$B_{\mathrm{nd}} \longrightarrow \varprojlim_n F^n 1$$

is not surjective. The target admits compatible finite-depth descriptions that have no finitely branching realization.

Proof. For $k \geq 0$, let t_k be the deterministic infinite chain whose successive outputs are $0^k 1 0^\omega$. Each t_k can be viewed as an F-coalgebra by taking singleton successor sets at each step. Let $\pi_n(t_k)$ denote the depth-n projection into $F^n 1$. Form a compatible family of root descriptions $q_n \in F^n 1$: its root has output 0 and its finite nonempty set of children (for $n \geq 2$) is

$$K_{n-1} = \{\pi_{n-1}(t_k) : k \geq 0\} \subseteq F^{n-1} 1.$$

Each K_{n-1} is finite, and the maps of the final chain take K_n onto K_{n-1}; set q_1 to have root output 0 and one formal successor. Thus (q_n) is a point of $\varprojlim_n F^n 1$.

Suppose that it came from a behaviour $b \in B_{\mathrm{nd}}$ with a finite set D of immediate successors. Then the set of depth-$(n-1)$ approximations of elements of D would equal K_{n-1} for every n. Hence $|K_{n-1}| \leq |D|$ for every n. But the patterns $0^k 10^\omega$ become mutually distinguishable in sufficiently deep finite prefixes, so the cardinalities $|K_{n-1}|$ grow without bound, a contradiction. □

Remark 8.2 (What finite modal tests can still do). This is the familiar stage-ω obstruction in the final-sequence analysis of finitary set functors [13], illustrated here for Moore output labels. It does not imply failure of finite-depth separation of *actual* image-finite behaviours. In the standard image-finite setting, distinct bisimulation classes can be distinguished by a finite-depth modal test: if equivalence through all finite depths held, image-finiteness lets us match successors and obtain a bisimulation. Thus the finite-depth observation map can be injective while its target contains additional coherent but unrealizable approximants. This is the point at which Spivak's finite-polynomial Stone-space result should not be transferred mechanically to the finite-powerset residual case.

Proposition 8.3 (Image-finite behaviours form a dense noncompact subspace)**.** *Assume A, O are finite and nonempty, and A has at least one letter while O has at least two elements. Endow $L = \varprojlim_n F^n 1$, for $F(Z) = O \times (\mathcal{P}_{\mathrm{f}} Z)^A$, with its limit of finite discrete topologies, and B_{nd} with the finite-depth test topology. The canonical map $\iota : B_{\mathrm{nd}} \to L$ is an embedding with dense proper image. In particular B_{nd} is not compact in that topology.*

Proof. Finite-depth separation for image-finite coalgebras makes ι injective (Remark above); since the test topology is initial for the finite projections, it is a topological embedding. Every element of $F^n 1$ can be realized by a finite-depth labelled tree with terminal states supplying the missing continuations; choosing arbitrary terminal outputs realizes a point of the final coalgebra having the prescribed n-th projection. Thus ι meets every nonempty finite-stage cylinder and is dense in L. Proposition 8.1 gives a compatible point absent from the image. Since L is compact Hausdorff, a compact image would be closed; a dense proper image cannot be compact. □

A natural related construction replaces the set-based finite-powerset functor by the Vietoris functor on Stone spaces, whose successor sets are compact instead of necessarily finite [14]. This is a candidate for a compact observational completion, not an assertion that its final coalgebra is identical to the image-finite B_{nd}.

9 Toward a Sweedler interaction between execution and experiments

The explicit Moore formulas suggest a larger, but genuinely open, problem. Fix accessible monads T and R, and let

$$C_{T,R} = \{T, R\}, \qquad B_{T,R} = C_{T,R}(1).$$

We know how $C_{T,R}$ classifies residual runners and how the FCS construction extracts operational topology and generalized halos. What is *not* supplied by this data in general is a preferred language of experiments that can distinguish the points of $B_{T,R}$.

An experimental semantics would require a choice of tests and outcomes, say

$$\mathrm{Test}_{T,R} \times B_{T,R} \longrightarrow \mathrm{Outcome}_R, \tag{24}$$

with enough composition or substitution structure to express *execute, then test*. The Boolean-valued case gives an observational topology from the preimages of designated accepting outcomes. If the operational continuations on $B_{T,R}$ transport tests by substitution, the pullback of a finite-test open along execution ought again to be a finite-test open, as in Theorem 3.1.

Proposition 2.1 now provides an explicit polynomial benchmark: the matching interaction is obtained from the Libkind–Spivak module action (4), a map $p \otimes p \to p + 1$, and free-monad extension. The Sweedler measuring adjunction (McDermott–Rivas–Uustalu, Section 3, and its specialization to accessible monads in Theorem 1 [4]) gives the universal bijection

$$\mathrm{Comon}(D, \{T, R\}) \cong \mathrm{Mon}(T, [D, R]),$$

where $[D, R]$ is the appropriate residual internal hom in the duoidal setting, not simply the ordinary cartesian exponential. In the polynomial instance this classifies the map (10). The nontrivial residue $\mathfrak{m}_{p+1}$ supplies the capacity to record failed matches inside a finite computation; it is not the identity-residue Sweedler dual in (9). However, the construction does *not* provide a canonical $\mathrm{Test}_{T,R}$ for arbitrary T, R; even the two syntaxes in Remark 2.2 have to be compared carefully.

A reasonable next step is to seek a *residual Hankel pairing*, with histories, tests, and their outcomes all typed by the source and residue:

$$\langle -, - \rangle_{T,R} : \mathrm{Hist}_{T,R} \times \mathrm{Test}_{T,R} \longrightarrow \mathrm{Outcome}_R \,. \tag{25}$$

The desired theorem would show how this pairing factors through a space of runner states, how it separates reachable states, and when it induces a polarity, enrichment, or another authentic duality. The following qualifications are essential.

(i) Reachability depends on the available initial-state data and implementation; it is not determined merely by the existence of $C_{T,R}$.

(ii) Finite experimental syntax is not generally identical to the source monad T or Spivak's $\mathfrak{m}_p$.

(iii) A residual effect may be externally visible, hidden, or only partially testable; the logic depends on that choice.

(iv) Nondeterministic and probabilistic outcomes can have joint information that is absent from separate marginal tests. The nonrepresentable halos in the companion note provide a concrete obstruction to naive productwise treatment.

(v) Kalman's transpose duality requires a dualizing structure; general Set-based minimal realization alone yields an image factorization, not an equivalence between reachability and observability constructions.

Theorem 9.1 (Established Moore benchmark)**.** *For a deterministic Moore theory with finite output set O and arbitrary input alphabet A, $B = O^{A^*}$ carries both the finite-test Stone topology and the FCS Alexandrov topology. Every continuation is continuous for the former, but not necessarily for the latter. For a pointed machine, the minimal realization is the least execution-open neighbourhood of its initial behaviour, and finite tests restricted to that neighbourhood yield the Hankel columns. With Boolean outputs this pairing induces the polarity of Proposition 5.1.*

Proof. Combine Theorem 3.1, Proposition 3.4, Theorem 4.1, and Proposition 5.1, together with the product-topology description of the Moore behaviour space. □

10 Editorial placement and next calculations

This note is intended as a second companion interlude following Part IX on behaviours. The first, *From Behaviour Categories to Generalized Spaces*, investigates the category, topology, and halos extracted from Sweedler comonads. The present note studies the complementary topology generated by finite observations and the pairing between operational histories and tests. The mathematical center is Theorem 4.1: the minimal realization is the least execution-open neighbourhood of the initial behaviour, while test evaluations recover its Hankel columns. The identity $e_v\sigma_u = e_{uv}$ explains the compatibility, and Proposition 3.4 marks its asymmetric limit.

We propose three concrete next calculations rather than a general unification claim: (1) construct a syntactic comparison between Moore getter/update experiments and Spivak's finite p-trees, now taking (7) as a benchmark for the latter; (2) develop a modal-test semantics for nondeterministic Moore runners and compare their noncompact image-finite behaviour space with Vietoris/Stone coalgebras; and (3) determine how a chosen observable part of the residual monad governs finite tests and their joint correlations. The further connections to ionads, dependent quantified modalities, probability, and smooth dynamics belong to later parts of the programme.

Status and provenance. Spivak's free/cofree constructions and Gendlin map, the Libkind–Spivak module action, the MRU Sweedler–runner correspondence, FCS's generalized-space constructions, classical automata minimization, and finite-dimensional Kalman duality are inherited ideas [1, 2, 4, 3]. The identification of the Gendlin matching algorithm with the specialization (7) is spelled out here from those constructions. The Moore computations and topological reformulation are proved here; minimal automata, the Hankel factorization, FCA polarities, the algebra–coalgebra duality of reversal, and the final-chain obstruction are classical results, now cited at their appropriate points. The general residual experimental syntax and its putative duality are proposals only; no priority claims are made.

References

[1] D. I. Spivak, *The free monad, cofree comonad, and topological space associated to a polynomial*, Topos Institute Berkeley Seminar outline, 14 June 2022. https://dspivak.net/talks/pdfs/20220614-topos-scotttopology.pdf.

[2] S. Libkind and D. I. Spivak, *Pattern Runs on Matter: The Free Monad Monad as a Module over the Cofree Comonad Comonad*, EPTCS **429** (2025), 1–28, Theorem 3.4 and Section 4. doi:10.4204/EPTCS.429.1.

[3] A. D. Fairbanks, K. Carlson, and D. I. Spivak, *Comonads as Spaces*, arXiv:2607.15091 (2026). https://arxiv.org/abs/2607.15091.

[4] D. McDermott, E. Rivas, and T. Uustalu, *Sweedler Theory of Monads*, FoSSaCS 2022, LNCS 13242, 428–448. doi:10.1007/978-3-030-99253-8_22.

[5] R. Garner, *The costructure–cosemantics adjunction for comodels for computational effects*, Math. Struct. Comput. Sci. **32**(4) (2022), 374–419. doi:10.1017/S0960129521000219.

[6] J. J. M. M. Rutten, *Universal coalgebra: a theory of systems*, Theoret. Comput. Sci. **249**(1) (2000), 3–80. doi:10.1016/S0304-3975(00)00056-6.

[7] M. A. Arbib and E. G. Manes, *Machines in a category: an expository introduction*, SIAM Rev. **16**(2) (1974), 163–192. doi:10.1137/1016026.

[8] J. A. Goguen, *Realization is universal*, Math. Systems Theory **6**(4) (1972), 359–374. https://research.ibm.com/publications/realization-is-universal.

[9] F. Bonchi, M. M. Bonsangue, H. H. Hansen, P. Panangaden, J. J. M. Rutten, and A. Silva, *Algebra–coalgebra duality in Brzozowski's minimization algorithm*, ACM Trans. Comput. Log. **15**(1), article 3 (2014), 1–29. doi:10.1145/2490818.

[10] J. A. Brzozowski, *Derivatives of regular expressions*, J. ACM **11**(4) (1964), 481–494. doi:10.1145/321239.321249.

[11] B. Ganter and R. Wille, *Formal Concept Analysis: Mathematical Foundations*, Springer (1999). doi:10.1007/978-3-642-59830-2.

[12] M. Hennessy and R. Milner, *Algebraic laws for nondeterminism and concurrency*, J. ACM **32**(1) (1985), 137–161.

[13] J. Worrell, *On the final sequence of a finitary set functor*, Theoret. Comput. Sci. **338**(1–3) (2005), 184–199. doi:10.1016/j.tcs.2004.12.009.

[14] C. Kupke, A. Kurz, and Y. Venema, *Stone coalgebras*, Theoret. Comput. Sci. **327**(1–2) (2004), 109–134. doi:10.1016/j.tcs.2004.07.023.