

Separation and Gluing of Explanations on Sites of Dynamical Systems

Paul Wang

LIP6, Sorbonne Université, Paris, France

paul.wang@lip6.fr

We construct a Grothendieck site whose objects are Mealy machines over definable sets in an o-minimal structure and whose coverings are jointly surjective families of definable open immersions. On this site, we define presheaves of *explanations*—systems equipped with an interpretable interface, parameterised by a “judge.” We prove that the behavioral presheaf (quotienting by observable output equivalence) is *separated*: a global explanation is determined by its local restrictions. We show that gluing fails in general—locally consistent explanations need not assemble globally—and give, for stateless explanatory systems of the restricted-interface presheaf, a necessary and sufficient topological condition for the sheaf property in terms of robust disconnection of fibers of the judge.

1 Introduction

Interpretability in machine learning seeks to explain *why* a model produces a given output. A striking feature of current methods is their inherent **locality**: explanations are given for specific data points, specific circuits, or specific feature directions—never for the model as a whole. The fundamental question is: *when multiple local explanations are available, do they fit together into a coherent global picture?*

This is precisely the question addressed by sheaf theory. A sheaf assigns data to each open set and specifies how data on overlaps must agree; the sheaf condition captures the principle that compatible local data glue uniquely to global data. When local data fail to glue, the obstruction can in principle be measured by sheaf cohomology (§7).

Sheaf-theoretic approaches to dynamical systems have been studied before: Schultz, Spivak, and Vasilakopoulou [22] model systems as sheaves on time intervals; Goguen [9] uses sheaves on observations; Joyal, Nielsen, and Winskel [14] use presheaves on paths. In each case, a Grothendieck topology is placed on a “base” category, and systems are modelled *as* sheaves. In a different vein, Abramsky and Brandenburger [1] use sheaves to characterise non-locality and contextuality in quantum foundations: local measurement data fail to extend to a global section, an obstruction measured by sheaf cohomology. Our gluing failure (Counterexample 5.2) is structurally analogous—compatible local explanations fail to extend globally—but the mathematical content is entirely different: their presheaves live on measurement scenarios, ours on categories of dynamical systems. On the interpretability side, Geiger et al. [8] develop *causal abstraction* as a framework for mechanistic interpretability; their notion of faithfulness relates a single abstraction to a model, whereas our framework studies when *multiple* local abstractions (sections) compose. Bassan, Amir, and Katz [3] study local-vs-global interpretability from a computational complexity perspective.

We take a different approach: we define a Grothendieck topology on a category whose *objects are systems* and whose morphisms are open immersions. The sheaf condition on such a site formalises a local-to-global principle with respect to *system representations*. To the best of our knowledge, no such site structure has previously been studied.

Contributions. Working over definable sets in an o-minimal structure (a setting providing tame geometry and a canonical class of open immersions):

1. We prove that categories of homogeneous Mealy machines over an adhesive base inherit adhesivity (Theorem 2.8), and construct a Grothendieck site $(\text{Sys}_{\text{oi}}, \text{Cov})$ of definable open immersions (§2).
2. We define a hierarchy of presheaves of explanations, parameterised by a “judge” translating the system’s interface to interpretable concepts (§3).
3. We prove that the behavioral presheaf $\mathcal{F}_j^{\text{beh}}$ is *separated* (§4): a global explanation is determined by its local restrictions.
4. We show that gluing *fails* in general (§5), and give a necessary and sufficient topological condition for the sheaf property of the presheaf $\mathcal{F}_{j,1}^{\text{ri}}$ of restricted-interface stateless explanations (§5.3), in terms of robust disconnection of fibers of the judge (Theorem 5.5).
5. We introduce ε -approximate sections (§6) and show that the dimension of the output space controls the depth of higher-order reconciliation obstructions via Helly’s theorem (Theorem 6.3).

Notation. We write $\mathcal{S} = (S_b, S_a, I, O, \alpha)$ for a heterogeneous Mealy machine with before-states S_b , after-states S_a , inputs I , outputs O , and dynamics $\alpha : S_b \times I \rightarrow S_a \times O$. All systems have non-empty I and O .

2 Systems, Sites, and the O-Minimal Setting

We set up the categorical framework: Mealy machines over definable sets in an o-minimal structure, definable open immersions as morphisms, and the resulting Grothendieck site.

2.1 Heterogeneous Mealy machines

Definition 2.1. Let \mathcal{C} be a category with finite products. A **heterogeneous Mealy machine** is a quintuple $\mathcal{S} = (S_b, S_a, I, O, \alpha)$ with $\alpha : S_b \times I \rightarrow S_a \times O$, where S_b (before-states) and S_a (after-states) may differ; this arises naturally when restricting to a subset of states whose after-states may leave the subset (cf. [23] for a categorical treatment of time-varying state spaces). When $S_b = S_a =: S$, we recover a *homogeneous* Mealy machine (S, I, O, α) ; in **Set**, these are coalgebras for the endofunctor $H(X) = (X \times O)^I$.

A **morphism** $\mathcal{S} \rightarrow \mathcal{S}'$ is a 4-tuple (f_b, f_a, f_I, f_O) making the dynamics square commute:

$$\begin{array}{ccc} S_b \times I & \xrightarrow{\alpha} & S_a \times O \\ f_b \times f_I \downarrow & & \downarrow f_a \times f_O \\ S'_b \times I' & \xrightarrow{\alpha'} & S'_a \times O' \end{array}$$

We write Sys for the resulting category, and $\text{Sys}_{\text{ho}}(I, O)$ for the *non-full* subcategory of fixed- (I, O) homogeneous systems where morphisms between state objects are the same for before-state and after-state.

2.2 The o-minimal setting

ACT approaches to dynamical systems typically work over **Set**, topological spaces, or categories of smooth manifolds. We work instead over $\text{Def}(\mathbb{M})$, the category of *definable sets* in an o-minimal structure \mathbb{M} , or its wide subcategory $\text{Def}_c(\mathbb{M})$ of continuous definable maps (Remark 2.4). The motivation comes from a classical idea of Grothendieck, who proposed in his *Esquisse d'un Programme* [10] that the category of topological spaces is too wild for geometry, and called for a theory of *topologie modérée* (tame topology):

La notion de “espace topologique” [...] est manifestement inadéquate [...]. La structure topologique [...] laisse coexister des phénomènes “modérés” [...] avec des phénomènes “sauvages” [...]. C’est la prédominance de ces derniers qui rend inadéquate la notion actuelle de “variété topologique” [...]. On voit mal d’ailleurs quelle idée raisonnable on pourrait se faire d’un “espace modéré” de dimension quelconque, tant qu’on ne dispose pas d’un “principe de modération” qui permette de se mettre à l’abri. —Grothendieck, Esquisse d’un Programme [10], §5

O-minimal structures, introduced by van den Dries [6] following work of Pillay and Steinhorn, provide precisely such a “principe de modération”: a single axiom on one-dimensional definable sets forces tameness in all dimensions.

Definition 2.2 ([6, Ch. 1]). An **o-minimal structure** \mathbb{M} expanding $(\mathbb{R}, <, +, \cdot)$ is given by a sequence $(\mathcal{D}_n)_{n \geq 1}$ of Boolean algebras of subsets of \mathbb{R}^n , closed under products and projections, such that \mathcal{D}_n contains all the semialgebraic sets and \mathcal{D}_1 consists exactly of finite unions of intervals and points. Elements of \mathcal{D}_n are **definable sets**; a function between them is **definable** if its graph is.

The condition on \mathcal{D}_1 is the “principle of moderation”: it forbids Cantor sets, space-filling curves, and other pathologies. The key consequences for our purposes are:

1. *Cell decomposition.* Every definable set $X \subseteq \mathbb{R}^n$ admits a finite partition into *cells* (definable sets of a simple form, that are canonically homeomorphic to open balls), and every definable function on X can be made piecewise continuous by refining the partition [6, Ch. 3]. In particular, definable open subsets are finite unions of connected definable opens.
2. *Adhesivity.* The category $\text{Def}(\mathbb{M})$ is adhesive [15]: it has pullbacks, pushouts along monomorphisms, and these satisfy the Van Kampen condition. The proof uses the faithful embedding $\text{Def}(\mathbb{M}) \hookrightarrow \mathbf{Set}$: pullbacks and pushouts along monomorphisms are computed on underlying sets and remain definable (images of definable maps are definable), so the Van Kampen condition reduces to the known adhesivity of **Set**. Adhesivity ensures that pushouts along open immersions preserve monomorphisms—the condition will be used for cogerm gluing (Theorem 5.1).
3. *Finite cohomology.* Edmundo and Peatfield [7] develop Čech cohomology for definable sheaves over o-minimal structures, with finite admissible covers and finite simplicial nerves. Our site Sys_{oi} is not directly a definable site in their sense (it is a category of systems, not a single definable space), but coverings of a fixed system have finite definable nerves, making their techniques a natural starting point for the cohomological obstructions discussed in §7.

Each definable set $X \subseteq \mathbb{R}^n$ inherits a topology from \mathbb{R}^n . A **definable open immersion** $\iota : U \hookrightarrow X$ is a definable injection whose image is a definable open subset.

Example 2.3. Standard examples of o-minimal structures: \mathbb{R}_{alg} (semialgebraic sets—polynomials only), \mathbb{R}_{an} (globally subanalytic sets—convergent power series), $\mathbb{R}_{\text{an,exp}}$ (adding the exponential function). Neural network activation functions (ReLU, sigmoid, softmax) are definable in $\mathbb{R}_{\text{an,exp}}$, so they naturally fit in this framework. Indeed, o-minimality has found direct applications in deep learning theory: Ji and Telgarsky [12] use tools from o-minimal analysis to prove directional convergence of gradient flow on deep homogeneous networks, and Bolte and Pauwels [5] exploit Whitney stratification of definable sets for convergence in values of stochastic gradient methods.

Remark 2.4. Working over **Set** would allow all injections as morphisms, with no topological control: any partition of S_b gives a covering, and the geometry of the state space plays no role. In $\text{Def}(\mathbf{M})$, the o-minimal topology provides a canonical class of “local” inclusions, finite coverings with finitely many connected components, and the tameness needed for the sheaf characterisation of Theorem 5.5. The cost is that $\text{Def}(\mathbf{M})$ is not cartesian closed.

Morphisms in $\text{Def}(\mathbf{M})$ need not be continuous; we write $\text{Def}_c(\mathbf{M})$ for the wide subcategory with continuous definable maps. The site construction works over both $\text{Def}(\mathbf{M})$ and $\text{Def}_c(\mathbf{M})$, since open immersions are continuous by nature (inclusions of open subsets with the subspace topology). Adhesivity, separation, and cogerm gluing work only in $\text{Def}(\mathbf{M})$ for now. The sheaf characterisation (Theorem 5.5) requires $\text{Def}_c(\mathbf{M})$: its proof uses compactness and normality of the input space, which depend on maps being continuous.

2.3 The site Sys_{oi}

We now combine the system structure with the o-minimal geometry.

Definition 2.5 (Definable open immersion of systems). A morphism $\iota : \mathcal{S}' \hookrightarrow \mathcal{S}$ in Sys (over $\text{Def}(\mathbf{M})$) is called a **definable open immersion** if each component $\iota_b, \iota_a, \iota_I, \iota_O$ is a definable open immersion. Write Sys_{oi} (“open immersions”) for the wide subcategory with these as morphisms.

Definition 2.6 (Coverings). A finite family $\{m_\alpha : \mathcal{S}_\alpha \hookrightarrow \mathcal{S}\}$ of definable open immersions is a **covering** if the induced maps are jointly surjective: $\bigsqcup_\alpha (S_{\alpha,b} \times I_\alpha) \twoheadrightarrow S_b \times I$ and $\bigsqcup_\alpha (S_{\alpha,a} \times O_\alpha) \twoheadrightarrow S_a \times O$.

Joint surjectivity on $S_b \times I$ is strictly stronger than separate-factor surjectivity ($\bigsqcup S_{\alpha,b} \twoheadrightarrow S_b$ and $\bigsqcup I_\alpha \twoheadrightarrow I$ separately): the product condition ensures that for every $(s, i) \in S_b \times I$, some single patch covers both s and i .

Proposition 2.7. Sys_{oi} has pullbacks (computed componentwise as intersections of definable opens), and the coverings above define a Grothendieck pretopology [17, Ch. III], making $(\text{Sys}_{\text{oi}}, \text{Cov})$ a site.

Proof sketch. (T1) Isomorphisms cover trivially.

(T2) Given a covering $\{m_\alpha\}$ of \mathcal{S} and $n : \mathcal{S}' \hookrightarrow \mathcal{S}$, the pullback has components $(S_{\alpha,b} \cap S'_b) \times (I_\alpha \cap I')$; joint surjectivity restricts from $S_b \times I$ to $S'_b \times I'$.

(T3) Compositions of definable open immersions are definable open immersions, and joint surjectivity is transitive. \square

Theorem 2.8. For all I, O in $\text{Def}(\mathbf{M})$, the category $\text{Sys}_{\text{ho}}(I, O)$ is adhesive.

This is related to a more general result on coalgebras (Appendix B):

Theorem 2.9. Let \mathcal{C} be an adhesive category [15] with finite products, and let $H : \mathcal{C} \rightarrow \mathcal{C}$ be a pullback-preserving endofunctor. Then the category of coalgebras $\text{Coalg}(H)$ is adhesive.

For $\mathcal{C} = \mathbf{Set}$, this was established by Padberg [19]; the general case appears to be new. Adhesivity ensures colimit constructions in gluing proofs produce well-defined homogeneous systems; combined with the strict initial object (the empty definable set), it also gives extensivity [15, Lemma 4.1].

3 Judges and Presheaves of Explanations

An *explanation* of a system translates its dynamics into human-interpretable terms. Here, we model human-interpretable as being relative to a *judge*: a choice of interpretable interface, along with a map from the system’s interface to the interpretable one.

3.1 Judges and sections

Definition 3.1 (Judge). A **judge** for $\mathcal{S} = (S_b, S_a, I, O, \alpha)$ is a pair $j = (j_I, j_O)$ of definable maps $j_I : I \rightarrow I'$, $j_O : O \rightarrow O'$, where (I', O') is the **interpretable interface**.

Construction 3.2 (Presheaf of explanations). Fix a judge $j : (I, O) \rightarrow (I', O')$. For a definable open immersion $m : \mathcal{S}_U \hookrightarrow \mathcal{S}$, define $\tilde{\mathcal{F}}_j(\mathcal{S}_U)$ as the set of pairs (\mathcal{S}', ψ) where $\mathcal{S}' = (S'_b, S'_a, I', O', \alpha')$ and $\psi : \mathcal{S}_U \rightarrow \mathcal{S}'$ is a morphism with interface components $\psi_I = j_I \circ m_I$ and $\psi_O = j_O \circ m_O$. The dynamics square

$$\begin{array}{ccc} S_{U,b} \times I_U & \xrightarrow{\alpha_U} & S_{U,a} \times O_U \\ \psi_b \times (j_I \circ m_I) \downarrow & & \downarrow \psi_a \times (j_O \circ m_O) \\ S'_b \times I' & \xrightarrow{\alpha'} & S'_a \times O' \end{array}$$

says that \mathcal{S}' represents the dynamics of \mathcal{S}_U as seen through the judge. Restriction is by precomposition: $(\mathcal{S}', \psi) \mapsto (\mathcal{S}', \psi \circ n)$ for $n : \mathcal{S}_{U'} \hookrightarrow \mathcal{S}_U$.

The presheaf $\tilde{\mathcal{F}}_j$ on the slice site $(\text{Sys}/\mathcal{S})_{\text{oi}}^{\text{op}}$ is a sheaf (separation and gluing reduce to joint surjectivity, since sections are literal pairs (\mathcal{S}', ψ)), but too fine: it distinguishes explanations that differ only in inessential internal structure. The question is which quotient to take.

3.2 Equivalence relations and the presheaf hierarchy

Convention 3.3. Throughout the remainder of this paper, all **explanatory systems** \mathcal{S}' (targets of sections) are homogeneous ($S'_b = S'_a =: S'$), so that dynamics can be iterated and behavioral equivalence is well-defined. The systems being explained remain heterogeneous, for they can be given by patches of before-states and after-states that don’t coincide.

Definition 3.4 (Cogerm equivalence). Two sections (\mathcal{S}'_1, ψ_1) and (\mathcal{S}'_2, ψ_2) over \mathcal{S}_U are **cogerm-equivalent** if there exists a homogeneous system \mathcal{S}'' with interface (I', O') , monomorphisms $i_k : \mathcal{S}'' \hookrightarrow \mathcal{S}'_k$ ($k = 1, 2$) in $\text{Sys}_{\text{ho}}(I', O')$ (i.e. with identity interface components and $i_{k,b} = i_{k,a}$), and a $\varphi : \mathcal{S}_U \rightarrow \mathcal{S}''$ (necessarily unique) with $i_k \circ \varphi = \psi_k$, for $k = 1, 2$. Write $(\mathcal{S}'_1, \psi_1) \sim (\mathcal{S}'_2, \psi_2)$.

$$\begin{array}{ccccc} \mathcal{S}'_1 & \xleftarrow{i_1} & \mathcal{S}'' & \xleftarrow{i_2} & \mathcal{S}'_2 \\ & \swarrow \psi_1 & \uparrow \varphi & \searrow \psi_2 & \\ & & \mathcal{S}_U & & \end{array}$$

The name reflects an analogy with germs: two explanations agree if they share a common “essential core” on the target side. The structure is parallel to the span-based characterisation of bisimulation in Joyal–Nielsen–Winskel [14]: both use spans as equivalence witnesses, but with monomorphisms in place of open maps and on the target side rather than the path side.

We denote this quotient $\tilde{\mathcal{F}}_j^{\text{cogerm}} := \tilde{\mathcal{F}}_j / \sim$.

Definition 3.5 (Behavioral equivalence). Two explanations (\mathcal{S}'_1, ψ_1) and (\mathcal{S}'_2, ψ_2) over \mathcal{S}_U are **behaviorally equivalent**, written \approx , if for every $x \in S_{U,b}$, every n , and every $(i'_1, \dots, i'_n) \in (I')^n$, the output sequences from $\psi_{1,b}(x)$ in \mathcal{S}'_1 and $\psi_{2,b}(x)$ in \mathcal{S}'_2 coincide (well-defined since \mathcal{S}'_k is homogeneous by Convention 3.3).

Cogerm equivalence implies behavioral equivalence, but not conversely (e.g., a 3-cycle and a 6-cycle with the same constant output are behaviorally equivalent but not cogerm-equivalent). For homogeneous Mealy machines over **Set**, behavioral equivalence is related to the standard coalgebraic notion of equality of images in the final coalgebra for $H(X) = (X \times O')^I$ [20]. From the interpretability perspective, behavioral equivalence is a natural relation: an explanation serves to predict behavior through the judge, and two explanations making identical predictions are interchangeable.

Definition 3.6 (Restricted-interface presheaf). For $m : \mathcal{S}_U \hookrightarrow \mathcal{S}$, define $I'_U := \text{im}(j_I \circ m_I) \subseteq I'$. The **restricted-interface presheaf** $\mathcal{F}_j^{\text{ri}}(\mathcal{S}_U)$ consists of explanations (\mathcal{S}', ψ) where \mathcal{S}' has input interface I'_U (rather than all of I'), modulo restricted-behavioral equivalence (testing input sequences in $(I'_U)^n$ only).

The natural transformation $\eta : \mathcal{F}_j^{\text{beh}} \rightarrow \mathcal{F}_j^{\text{ri}}$ given by restricting the input interface is surjective on sections (any section of $\mathcal{F}_j^{\text{ri}}$ over \mathcal{S}_U extends to one of $\mathcal{F}_j^{\text{beh}}$ by choosing arbitrary dynamics on $I' \setminus I'_U$, which does not affect the restricted-behavioral class) and an isomorphism when $I'_U = I'$. In practice, the distinction maps onto two modes of interpretability: sections of $\mathcal{F}_j^{\text{beh}}$ would correspond to *mechanistic explanations* (“this feature detects X across all inputs”), while sections of $\mathcal{F}_j^{\text{ri}}$ correspond to *local explanations* (“this circuit implements Y on distribution \mathcal{D} ”).

We write $\mathcal{F}_j^{\text{beh}} := \tilde{\mathcal{F}}_j / \approx$ for the **behavioral presheaf**. The hierarchy is:

$$\tilde{\mathcal{F}}_j \twoheadrightarrow \mathcal{F}_j^{\text{cogerm}} \twoheadrightarrow \mathcal{F}_j^{\text{beh}} \xrightarrow{\eta} \mathcal{F}_j^{\text{ri}}.$$

$\tilde{\mathcal{F}}_j$ is a sheaf but too fine; the three quotients capture progressively coarser notions of explanatory content. In general, none of the quotients is a sheaf (§4–5).

4 Separation

The separation axiom asks: if two global sections restrict to the same class on every member of a covering, are they globally equivalent?

Theorem 4.1 (Behavioral presheaf is separated). $\mathcal{F}_j^{\text{beh}}$ satisfies the separation axiom on $(\text{Sys}_{\text{oi}}, \text{Cov})$.

Proof. Behavioral equivalence is pointwise: $(\mathcal{S}'_1, \psi_1) \approx (\mathcal{S}'_2, \psi_2)$ iff $\text{Beh}_{\mathcal{S}'_1}(\psi_{1,b}(x)) = \text{Beh}_{\mathcal{S}'_2}(\psi_{2,b}(x))$ for every $x \in S_{U,b}$, where Beh denotes the map from states to infinite output trees. If this holds on every \mathcal{S}_α in a jointly surjective covering, then for every $x \in S_{U,b}$ there exists α with $x \in S_{\alpha,b}$, so the behaviours agree at x . \square

The cogerm presheaf $\mathcal{F}_j^{\text{cogerm}}$ does *not* satisfy separation in general: local cogerm witnesses may contain “extra states” outside the image of the section, and these extra states can disagree across patches. The passage to the behavioral quotient eliminates this problem by testing only observable output sequences.

Definition 4.2 (j -full covering). A covering $\{m_\alpha : \mathcal{S}_\alpha \hookrightarrow \mathcal{S}\}$ is **j -full** if $I'_{U_\alpha} = I'_U$ for every α , i.e. each patch’s inputs hit every fiber of j_I . Coverings where $m_{I,\alpha} = \text{id}_I$, which we call **data-local**, are automatically j -full.

Proposition 4.3 (Separation for $\mathcal{F}_j^{\text{ri}}$). $\mathcal{F}_j^{\text{ri}}$ satisfies separation for any j -full covering.

Proof. When $I'_{U_\alpha} = I'_U$ for all α , local and global restricted-behavioral equivalences test the same input sequences $(I'_U)^n$. Joint surjectivity on S_b (by projection from $S_b \times I$) completes the argument. \square

Separation can fail when the covering splits inputs across patches:

Counterexample 4.4. Consider $\mathcal{S} = (\{s_1, s_2\}, \{s_1, s_2\}, \{a, b\}, \{0, 1\}, \alpha)$ with $\alpha(s_1, a) = (s_1, 0)$, $\alpha(s_1, b) = (s_2, 0)$, $\alpha(s_2, a) = (s_1, 1)$, $\alpha(s_2, b) = (s_2, 1)$, and $j = \text{id}$. The input-splitting covering $\mathcal{S}_1 = \mathcal{S}|_{\{a\}}$, $\mathcal{S}_2 = \mathcal{S}|_{\{b\}}$ has $I'_{U_1} = \{a\} \neq \{b\} = I'_{U_2}$. An alternative explanation \mathcal{S}'_2 with different internal routing agrees with the identity explanation on each patch's pure input sequences, but disagrees on the mixed sequence (a, b) (see Figure 1).




Figure 1: The 2-state system of Counterexample 4.4 (separation failure for $\mathcal{F}_j^{\text{ri}}$). Edge labels: input/output. The input-splitting covering $\mathcal{S}|_{\{a\}}$, $\mathcal{S}|_{\{b\}}$ restricts each patch to a single input; an alternative explanation agrees on pure-input sequences but disagrees on the mixed sequence (a, b) .

5 The Gluing Landscape

While the cogerm presheaf $\mathcal{F}_j^{\text{cogerm}}$ always satisfies gluing (Theorem 5.1), the behavioral and restricted-interface presheaves do not. We map out the landscape of positive and negative results.

5.1 Gluing for $\mathcal{F}_j^{\text{cogerm}}$

Theorem 5.1. $\mathcal{F}_j^{\text{cogerm}}$ satisfies the gluing axiom: given a covering $\{m_\alpha : \mathcal{S}_\alpha \hookrightarrow \mathcal{S}\}$ and a compatible family of sections $[\mathcal{S}'_\alpha, \psi_\alpha] \in \mathcal{F}_j^{\text{cogerm}}(\mathcal{S}_\alpha)$, there exists $[\mathcal{S}', \psi] \in \mathcal{F}_j^{\text{cogerm}}(\mathcal{S})$ restricting to each $[\mathcal{S}'_\alpha, \psi_\alpha]$.

Proof sketch. Choose cogerm witnesses $\mathcal{S}''_{\alpha\beta}$ for compatibility on overlaps. The global target \mathcal{S}' is the colimit of $\dots \mathcal{S}'_\alpha \xleftarrow{i_\alpha} \mathcal{S}''_{\alpha\beta} \xrightarrow{i_\beta} \mathcal{S}'_\beta \dots$ in Sys; this exists as iterated pushouts along monomorphisms in the adhesive category $\text{Def}(\mathbf{M})$. Adhesivity ensures the canonical maps $\mathcal{S}'_\alpha \hookrightarrow \mathcal{S}'$ remain monic; the compatible morphisms ψ_α then glue to $\psi : \mathcal{S} \rightarrow \mathcal{S}'$ by joint surjectivity. \square

5.2 Gluing fails for $\mathcal{F}_j^{\text{beh}}$

Counterexample 5.2 (Gluing failure for $\mathcal{F}_j^{\text{beh}}$). Let $\mathcal{S} = (\{s_0, s_1, s_2, s_3\}, \{s_0, s_1, s_2, s_3\}, \{a, b\}, \{0, 1\}, \alpha)$ (see Figure 2) with dynamics:

$$\begin{aligned} \alpha(s_0, a) &= (s_0, 0), & \alpha(s_0, b) &= (s_2, 0), \\ \alpha(s_1, a) &= (s_1, 1), & \alpha(s_1, b) &= (s_1, 1), \\ \alpha(s_2, a) &= (s_1, 1), & \alpha(s_2, b) &= (s_1, 1), \\ \alpha(s_3, a) &= (s_3, 1), & \alpha(s_3, b) &= (s_2, 1), \end{aligned}$$

with judge $j_I : \{a, b\} \rightarrow \{\bullet\}$, $j_O = \text{id}$. The covering $\mathcal{S}_1 = \mathcal{S}|_{\{s_0, s_1, s_2\}}$, $\mathcal{S}_2 = \mathcal{S}|_{\{s_1, s_2, s_3\}}$ has non-trivial overlap $\{s_1, s_2\}$ and is data-local.

Over \mathcal{S}_1 : the minimal realisation $M_1 = (\{p_0, p_1\}, \{\bullet\}, \{0, 1\})$ maps $s_0 \mapsto p_0$ (constant output 0) and $s_1, s_2 \mapsto p_1$ (constant output 1). Over \mathcal{S}_2 : the realisation $M_2 = (\{q\}, \{\bullet\}, \{0, 1\})$ maps $s_1, s_2, s_3 \mapsto q$ (constant output 1). Compatibility holds on $\{s_1, s_2\}$: both produce constant output 1.

A global section $(M, \bar{\psi})$ would require: $\bar{\psi}(s_0)$ producing output 0, and $\bar{\psi}(s_3)$ producing output 1. The dynamics square at (s_0, b) forces $\bar{\psi}_a(s_2)$ to continue with output 0; at (s_3, b) , it forces $\bar{\psi}_a(s_2)$ to continue with output 1. Contradiction: a single after-state cannot produce both behaviours.




Figure 2: The 4-state system of Counterexample 5.2 (gluing failure for $\mathcal{F}_j^{\text{beh}}$). Edge labels: input/output; the judge collapses inputs ($j_I : \{a, b\} \rightarrow \{\bullet\}$). The dashed and dotted boxes show the covering patches; the overlap is $\{s_1, s_2\}$. The obstruction: $s_0 \xrightarrow{b} s_2$ and $s_3 \xrightarrow{b} s_2$ force incompatible behavioral requirements on the after-state s_2 .

The obstruction is that the dynamics sends states from distinct behavioral regions (s_0 and s_3 , in different patches) to a common after-state s_2 . Compatibility constrains the overlap $\{s_1, s_2\}$ but not the after-state reached from *outside* the overlap. The non-injectivity of j_I is essential: collapsing inputs forces a single dynamics-square constraint per initial state.

5.3 Sheaf characterisation for stateless explanations

Write $\mathcal{F}_{j,1}^{\text{ri}} \subseteq \mathcal{F}_j^{\text{ri}}$ for the **stateless subpresheaf**: sections (\mathcal{S}', ψ) with $|S'| = 1$. (Restriction preserves $|S'| = 1$, so this is indeed a subpresheaf; the notation generalises to $\mathcal{F}_{j,N}^{\text{ri}}$ for $|S'| \leq N$.) A stateless section reduces to a function $h : I_U \rightarrow O'$ with $h = g \circ j_I$ for a uniquely determined g , and gluing becomes a local-to-global question about $\ker(j_I)$ -invariance of continuous definable functions. We give a precise characterisation in the o-minimal setting.

Definition 5.3 (Robustly disconnected fiber). A fiber $j_I^{-1}(t_0)$ is **robustly disconnected** if there exist an open $N \ni t_0$ in I' and disjoint non-empty opens V, W in I with $j_I^{-1}(N) \subseteq V \cup W$, $V \cap j_I^{-1}(t_0) \neq \emptyset$, $W \cap j_I^{-1}(t_0) \neq \emptyset$.

Connected fibers are never robustly disconnected. The converse is strict: a fiber can be disconnected but not *robustly* so, when the disconnection “heals” in nearby fibers.

Example 5.4 (Punctured square). Let $I = [0, 1]^2 \setminus \{(0.5, 0.5)\}$, $j_I(x, y) = x$. The fiber at $x = 0.5$ is disconnected, but nearby fibers $\{x\} \times [0, 1]$ ($x \neq 0.5$) are connected and cannot be separated by $V \cup W$. No robust disconnection exists, so $\mathcal{F}_{j,1}^{\text{ri}}$ is a sheaf despite the disconnected fiber.

Theorem 5.5 (Sheaf characterisation for stateless explanations). In $\text{Def}_c(\mathbf{M})$ (the wide subcategory with continuous definable maps), suppose I is compact and $\text{im}(j_O)$ is definably connected with $|\text{im}(j_O)| \geq 2$. Then

$$\mathcal{F}_{j,1}^{\text{ri}} \text{ is a sheaf} \iff \text{no fiber of } j_I \text{ is robustly disconnected.}$$

When $\text{im}(j_O)$ is totally disconnected, the sheaf condition holds trivially.

Proof sketch. (\Rightarrow) Given a robust disconnection $j_I^{-1}(N) \subseteq V \cup W$, use the connected image of j_O and a definable bump function to construct a continuous $f : I \rightarrow \mathcal{O}$ that is $\ker(j_I)$ -invariant on V and on W separately, but with $j_O(f(a)) \neq j_O(f(b))$ for $a \in V \cap j_I^{-1}(t_0)$, $b \in W \cap j_I^{-1}(t_0)$. The covering $U_1 := V \cup (I \setminus j_I^{-1}(N))$, $U_2 := W \cup (I \setminus j_I^{-1}(N))$ provides compatible local sections that do not glue.

(\Leftarrow) If h is locally $\ker(j_I)$ -invariant but $h(x) \neq h(y)$ with $j_I(x) = j_I(y) = t_0$, separate the level sets of $h|_{j_I^{-1}(t_0)}$ by disjoint opens (normality of compact I), then use compactness to find $N \ni t_0$ with $j_I^{-1}(N)$ contained in the union—a robust disconnection. \square

Corollary 5.6. Connected fibers of $j_I \Rightarrow$ no robust disconnection $\Rightarrow \mathcal{F}_{j,1}^{\text{ri}}$ is a sheaf. The first implication is strict (Example 5.4).

5.4 Summary of the gluing landscape

Presheaf	Separated?	Gluing?
$\tilde{\mathcal{F}}_j$ (unquotiented)	Yes	Yes (sheaf)
$\mathcal{F}_j^{\text{cogerm}}$ (cogerm)	No	Yes
$\mathcal{F}_j^{\text{beh}}$ (behavioral)	Yes	No (Cex. 5.2)
$\mathcal{F}_j^{\text{ri}}$ (restricted-interface)	Yes for j -full	No in general
$\mathcal{F}_{j,1}^{\text{ri}}$ (stateless)	Yes for j -full	Thm. 5.5: iff no robust disconn.

The pattern: finer equivalence relations admit gluing but break separation; coarser ones yield separation but break gluing. The “right” quotient depends on the interpretability task.

6 Approximate Sections and Helly Obstructions

In practice, explanations are rarely exact: a probe achieves a given accuracy, a LIME fit has a fidelity bound, a concept bottleneck leaks information. When the dynamics square is relaxed from exact equality to ε -approximation, a new phenomenon emerges: pairwise consistency of local explanations no longer implies global consistency, and the *dimension of the output space* controls the depth of the obstruction.

6.1 ε -Sections

Let (\mathcal{O}', ρ) be a metric space. For a system \mathcal{S}_U with singleton state space, i.e. $|S_{U,b}| = |S_{U,a}| = 1$, write $f : I_U \rightarrow \mathcal{O}$ for the output function $f(i) := \pi_{\mathcal{O}}(\alpha_U(*, i))$, where $\pi_{\mathcal{O}}$ is the projection onto \mathcal{O} . A stateless explanatory section ($|S'| = 1$) over \mathcal{S}_U reduces to a function $g : I' \rightarrow \mathcal{O}'$ satisfying the dynamics square $g(j_I(i)) = j_{\mathcal{O}}(f(i))$ for all $i \in I_U$.

Definition 6.1 (ε -section). An ε -section over \mathcal{S}_U is a function $g : I' \rightarrow \mathcal{O}'$ satisfying

$$\rho(g(j_I(i)), j_{\mathcal{O}}(f(i))) \leq \varepsilon \quad \text{for all } i \in I_U.$$

Write $\mathcal{F}_j^\varepsilon(\mathcal{S}_U)$ for the set of ε -sections. Note that g is defined on the full interface I' , not the restricted $I'_U := \text{im}(j_I|_{I_U})$; this is the ε -relaxation of the stateless behavioral presheaf (stateless sections of $\mathcal{F}_j^{\text{beh}}$ with $|S'| = 1$), not of $\mathcal{F}_{j,1}^{\text{ri}}$. At $\varepsilon = 0$, this recovers the exact stateless sections of $\mathcal{F}_j^{\text{beh}}$.

For each interpretable input $i' \in I'_U$, define the **target set**

$$T(i', U) := \{j_{\mathcal{O}}(f(i)) : i \in I_U, j_I(i) = i'\} \subseteq \mathcal{O}'$$

and the **feasibility set** $A_\varepsilon(i', U) := \bigcap_{t \in T(i', U)} \bar{B}(t, \varepsilon)$, the intersection of closed ε -balls centred at the target values. An ε -section exists iff $A_\varepsilon(i', U) \cap \mathcal{O}' \neq \emptyset$ for every $i' \in I'_U$.

For a covering $\{U_1, \dots, U_n\}$, the target sets compose by union: $T(i', \bigcup_\alpha U_\alpha) = \bigcup_\alpha T(i', U_\alpha)$, so the global feasibility set decomposes as

$$A_\varepsilon(i', \bigcup_\alpha U_\alpha) = \bigcap_\alpha A_\varepsilon(i', U_\alpha). \quad (1)$$

Global section existence at i' thus requires the intersection of the per-patch feasibility sets to be non-empty.

Remark 6.2. $\mathcal{F}_j^\varepsilon$ is a sheaf: restriction is the identity on functions (a global $g : I' \rightarrow \mathcal{O}'$ is either ε -close or not on each patch), so separation is trivial, and gluing follows because the ε -constraint is pointwise in $i \in I_U$. This sheaf property relies on the stateless restriction; a presheaf of multi-state ε -sections would have a more complex structure.

6.2 Obstruction depth

At $\varepsilon = 0$, pairwise section existence implies global existence: if every pair of patches admits a common exact section, then the target sets are all singletons (the same point), so the global section exists. At $\varepsilon > 0$, this fails: the feasibility sets $A_\varepsilon(i', U_\alpha)$ are balls whose pairwise intersections may be non-empty even when the full intersection is empty. This is the setup of Helly's theorem.

Theorem 6.3 (Helly obstruction depth). *Suppose $\mathcal{O}' \subseteq \mathbb{R}^d$ is closed convex and ρ is the Euclidean metric. Let $\{U_1, \dots, U_n\}$ be a covering with $n \geq d + 1$, and fix $i' \in I'$. If $A_\varepsilon(i', U_{\alpha_1} \cup \dots \cup U_{\alpha_{d+1}}) \neq \emptyset$ for every $(d + 1)$ -element sub-covering, then $A_\varepsilon(i', U_1 \cup \dots \cup U_n) \neq \emptyset$.*

The bound $d + 1$ is sharp (Example 6.4).

Proof. Each $A_\varepsilon(i', U_\alpha) \cap \mathcal{O}'$ is convex (intersection of closed balls with a convex set in \mathbb{R}^d). By (1), $A_\varepsilon(i', \bigcup_{\alpha \in S} U_\alpha) \cap \mathcal{O}' = \bigcap_{\alpha \in S} A_\varepsilon(i', U_\alpha) \cap \mathcal{O}'$ for any subset S . The hypothesis says every $(d + 1)$ -element sub-family of $\{A_\varepsilon(i', U_\alpha) \cap \mathcal{O}'\}_\alpha$ has non-empty intersection. By Helly's theorem [11] in \mathbb{R}^d , the whole family has non-empty intersection. \square

Example 6.4 (Equilateral triangle). Let $O' = \mathbb{R}^2$, $I = \{a, b, c\}$, $j_I \equiv \star$ (all inputs collapse), $j_O = \text{id}$, and $f(a) = (0, 0)$, $f(b) = (2, 0)$, $f(c) = (1, \sqrt{3})$ (equilateral triangle, side 2). Three singleton patches $U_\alpha = \{\alpha\}$.

The pairwise Chebyshev radius (circumradius of a segment) is 1; the triple Chebyshev radius (circumradius of the triangle) is $2/\sqrt{3} \approx 1.155$. For $\varepsilon \in (1, 2/\sqrt{3})$:

- ε -sections exist on every pair of patches (pairwise Chebyshev radius = $1 \leq \varepsilon$),
- but no ε -section exists on the triple ($2/\sqrt{3} > \varepsilon$).

This is a genuine 3-fold obstruction in \mathbb{R}^2 , not reducible to any pairwise failure. The construction generalises: a regular $(d + 1)$ -simplex in \mathbb{R}^d gives a $(d + 1)$ -fold obstruction.

6.3 Interpretability reading

The Helly obstruction has an interpretability reading: *even if every pair of local explanations can be reconciled into a joint ε -approximate explanation, the full set may be globally irreconcilable, at least at ε level—and the dimension of the output space determines how many explanations must be simultaneously checked.*

For **hard** k -class classification ($O' = \{1, \dots, k\}$, discrete metric), the Helly number is 2: pairwise reconcilability always suffices. For **soft** classification ($O' = \Delta^{k-1}$, the probability simplex in \mathbb{R}^{k-1}), the Helly number is k : all k local explanations must be checked simultaneously.

For unrestricted multi-state explanations with ψ_b injective ($|S'| \geq |S_{U,b}|$, $\psi_b = \text{id}_{S_{U,b}}$), the trivial partition (each state in its own class) decouples the output constraints across states. The obstruction depth remains at most $\dim(O') + 1$ per state, with states contributing independently. The stateless case ($|S'| = 1$) is thus the **extremal case**: all states share one explanatory state, creating maximal constraint coupling. Understanding this in depth remains future work.

7 Discussion

Concrete instantiation. Consider a neural network as a dynamical system: the state space S_b is the space of hidden activations at a given layer, the dynamics α maps an activation–input pair to the next-layer activation and the network’s output, and the judge $j = (j_I, j_O)$ projects the raw input/output spaces onto interpretable concepts. A *section* over a region of activation space is then an explanation: a simpler system \mathcal{S}' with interpretable interface (I', O') that represents the network’s dynamics as seen through the judge. *Note that there are several ways of representing neural networks or Large Language Models in this framework, some of them stateless.*

Several interpretability methods illustrate this pattern. *Representation Engineering* [25] identifies “reading vectors” $v \in \mathbb{R}^d$ such that the projection $v^T \cdot h$ of a hidden state $h \in S_b$ tracks a high-level concept (honesty, sentiment, factuality). In our framework, j_O projects the network’s output onto the concept score, and a stateless section ($|S'| = 1$) is a function $g : I' \rightarrow O'$ predicting the concept score from interpretable inputs alone—in effect, a *probe*. *Automata extraction* [24] recovers a finite-state machine \mathcal{S}' simulating a transformer on a formal language: a *multi-state section* where the explanatory system has $|S'| > 1$ internal states and the judge projects tokens to a finite alphabet. *Feature attribution* methods (saliency maps [2], SHAP [16]) assign importance scores to input features; Bilodeau et al. [4] show that such methods can provably fail—in our language, the presheaf of feature-attribution explanations may have empty sections on certain regions. The **judge is the central design choice**: it formalises

what “interpretable” means, and different judges yield different presheaves with different local-to-global properties.

What the results say. *Separation* (Theorem 4.1) says that if two behavioral explanations agree on every data slice, they are globally identical. The *failure of gluing* (Counterexample 5.2) says that locally consistent explanations need not assemble into a global one. This is an analogous symptom to *polysemanticity*: when a neuron responds to multiple unrelated concepts, a “one concept per neuron” explanation works on each concept’s data region but fails to glue globally. The mechanism differs—polysemanticity arises from linear-algebraic superposition in activation space, while the gluing obstruction is dynamical—but the local-success-global-failure pattern is analogous.

The sheaf characterisation (Theorem 5.5) makes this topological: in the continuous definable setting, $\mathcal{F}_{j,l}^{\text{ri}}$ is a sheaf iff no fiber of j_l is robustly disconnected. Concretely, if the interpretable input projection groups raw inputs into clusters, gluing holds when these clusters are topologically connected—it fails when a cluster splits into components that cannot be healed by nearby clusters.

The ε -relaxation (§6) addresses the practical regime where explanations are approximate. A probe whose worst-case error is at most ε (in the metric on O') is an ε -section. (This is a uniform bound, stronger than the average-case accuracy typical in practice; the Helly bound applies to the worst-case regime.) Theorem 6.3 then predicts: for regression probes with d -dimensional output ($d \geq 2$), checking pairwise consistency of local probes is insufficient—one must check $(d + 1)$ -wise consistency. For hard classification, pairwise checks suffice (target values must all agree exactly); for soft k -class classifiers ($O' = \Delta^{k-1} \subset \mathbb{R}^{k-1}$, convex), Theorem 6.3 gives Helly number k . This gives a concrete prediction about when local interpretability audits miss global inconsistencies.

The role of the categorical framework. The results above are proved by direct topological and combinatorial arguments: separation is pointwise, the gluing obstruction is an explicit counterexample, the sheaf characterisation uses compactness and normality, and the Helly bound uses convex geometry. The categorical framework—the site of open immersions, the presheaf hierarchy, the covering conditions—provides the *questions*: Do local probes glue? Is separation satisfied? We expect future work—multi-state characterisations, graph locality, cohomological obstructions—to make the categorical tools directly load-bearing; the framework’s current value lies in organising concepts and identifying questions.

Abstract framework. The results factor through abstract axioms on an admissible pair $(\mathcal{C}, \mathcal{M})$ (Appendix A), clarifying what each result depends on: separation needs only joint epimorphicity of coverings; cogerm gluing needs full adhesivity (available in $\text{Def}(\mathcal{M})$ but not $\text{Def}_c(\mathcal{M})$; see Remark A.5); the sheaf characterisation additionally needs a topological enrichment. This factorisation identifies the interpretability-relevant results (separation, sheaf characterisation) as independent of adhesivity, and opens the framework to other base categories.

Future directions. When gluing fails, we can try to use Čech cohomology to measure the obstruction: the compatible family of Counterexample 5.2 suggests a non-trivial 1-cocycle, but the behavioral presheaf is a presheaf of sets, not of abelian groups, requiring care with non-abelian Čech cohomology or a linearisation step. Beyond data locality, two other dimensions of interpretability locality—*graph locality* (decomposing the computational graph, in the sense of [21, 18]) and *feature locality* (decomposing the activation space via coarsenings of the judge)—require different site structures whose interactions are an important direction for future work.

References

- [1] Samson Abramsky & Adam Brandenburger (2011): *The Sheaf-Theoretic Structure of Non-Locality and Contextuality*. *New Journal of Physics* 13(11), p. 113036, doi:10.1088/1367-2630/13/11/113036.
- [2] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt & Been Kim (2018): *Sanity Checks for Saliency Maps*. In: *Advances in Neural Information Processing Systems*, 31, Curran Associates, Inc. Available at https://proceedings.neurips.cc/paper_files/paper/2018/hash/294a8ed24b1ad22ec2e7efea049b8737-Abstract.html.
- [3] Shahaf Bassan, Guy Amir & Guy Katz (2024): *Local vs. Global Interpretability: A Computational Complexity Perspective*. In: *Proceedings of the 41st International Conference on Machine Learning (ICML)*.
- [4] Blair Bilodeau, Natasha Jaques, Pang Wei Koh & Been Kim (2024): *Impossibility theorems for feature attribution*. *Proceedings of the National Academy of Sciences* 121(2), p. e2304406120, doi:10.1073/pnas.2304406120.
- [5] Jérôme Bolte & Edouard Pauwels (2021): *Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning*. *Mathematical Programming* 188, pp. 19–51, doi:10.1007/s10107-020-01501-5.
- [6] Lou van den Dries (1998): *Tame Topology and O-minimal Structures*. *London Mathematical Society Lecture Note Series* 248, Cambridge University Press, doi:10.1017/CBO9780511525919.
- [7] Mário J. Edmundo & Nicholas J. Peatfield (2008): *O-minimal Čech cohomology*. *Quarterly Journal of Mathematics* 59(2), pp. 213–220, doi:10.1093/qmath/ham033.
- [8] Atticus Geiger, Duligur Ibeling, Amir Zur, Thomas Icard & Joseph Y. Halpern (2025): *Causal Abstraction: A Theoretical Foundation for Mechanistic Interpretability*. *Journal of Machine Learning Research*. To appear.
- [9] Joseph A. Goguen (1992): *Sheaf Semantics for Concurrent Interacting Objects*. *Mathematical Structures in Computer Science* 2(2), pp. 159–191, doi:10.1017/S0960129500001420.
- [10] Alexander Grothendieck (1997): *Esquisse d'un Programme*. In Leila Schneps & Pierre Lochak, editors: *Geometric Galois Actions, London Mathematical Society Lecture Note Series* 242, Cambridge University Press, pp. 5–48, doi:10.1017/CBO9780511758874.003. Written in 1984.
- [11] Eduard Helly (1923): *Über Mengen konvexer Körper mit gemeinschaftlichen Punkten*. *Jahresbericht der Deutschen Mathematiker-Vereinigung* 32, pp. 175–176.
- [12] Ziwei Ji & Matus Telgarsky (2020): *Directional convergence and alignment in deep learning*. In: *Advances in Neural Information Processing Systems (NeurIPS)*, 33, pp. 17176–17186.
- [13] Peter Johnstone, John Power, Toru Tsujishita, Hiroshi Watanabe & James Worrell (2001): *On the Structure of Categories of Coalgebras*. *Theoretical Computer Science* 260(1–2), pp. 87–117, doi:10.1016/S0304-3975(00)00124-9.
- [14] André Joyal, Mogens Nielsen & Glynn Winskel (1996): *Bisimulation from Open Maps*. *Information and Computation* 127(2), pp. 164–185, doi:10.1006/inco.1996.0057.
- [15] Stephen Lack & Paweł Sobociński (2005): *Adhesive and quasiadhesive categories*. *RAIRO - Theoretical Informatics and Applications* 39(3), pp. 511–545, doi:10.1051/ita:2005028.
- [16] Scott M. Lundberg & Su-In Lee (2017): *A Unified Approach to Interpreting Model Predictions*. In: *Advances in Neural Information Processing Systems*, 30, Curran Associates, Inc., pp. 4765–4774.
- [17] Saunders Mac Lane & Ieke Moerdijk (1994): *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Universitext, Springer-Verlag, doi:10.1007/978-1-4612-0927-0.
- [18] Nelson Niu & David I. Spivak (2025): *Polynomial Functors: A Mathematical Theory of Interaction*. *LMS Lecture Note Series* 498, Cambridge University Press.
- [19] Julia Padberg (2017): *Hierarchical Graph Transformation Revisited*. In: *Graph Transformation (ICGT 2017)*, LNCS 10373, Springer, pp. 20–35, doi:10.1007/978-3-319-61470-0_2. Extended version: arXiv:1702.04650.

- [20] J. J. M. M. Rutten (2000): *Universal coalgebra: a theory of systems*. *Theoretical Computer Science* 249(1), p. 3–80, doi:10.1016/S0304-3975(00)00056-6.
- [21] Patrick Schultz & David I. Spivak (2019): *Temporal Type Theory: A Topos-Theoretic Approach to Systems and Behavior*. doi:10.1007/978-3-030-00704-1.
- [22] Patrick Schultz, David I. Spivak & Christina Vasilakopoulou (2020): *Dynamical Systems and Sheaves*. *Applied Categorical Structures* 28, pp. 1–57, doi:10.1007/s10485-019-09565-x.
- [23] David Sprunger & Shin-ya Katsumata (2019): *Differentiable Causal Computations via Delayed Trace*. In: *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pp. 1–12, doi:10.1109/LICS.2019.8785670.
- [24] Yihao Zhang, Zeming Wei & Meng Sun (2024): *Automata Extraction from Transformers* (arXiv:2406.05564). doi:10.48550/arXiv.2406.05564. Available at <http://arxiv.org/abs/2406.05564>.
- [25] Andy Zou, Long Phan, Sarah Chen, James Campbell, Phillip Guo, Richard Ren, Alexander Pan, Xuwang Yin, Mantas Mazeika, Ann-Kathrin Dombrowski, Shashwat Goel, Nathaniel Li, Michael J. Byun, Zifan Wang, Alex Mallen, Steven Basart, Sanmi Koyejo, Dawn Song, Matt Fredrikson, J. Zico Kolter & Dan Hendrycks (2023): *Representation Engineering: A Top-Down Approach to AI Transparency*. doi:10.48550/ARXIV.2310.01405. Available at <https://arxiv.org/abs/2310.01405>.

A Abstract framework

We give precise statements and proofs for the axiomatic framework described in §7. Each result uses only the axioms it needs: *separation* (Theorem 4.1) requires only that the covering family is jointly epimorphic and the pointwise nature of behavioral equivalence; *cogerm gluing* (Theorem 5.1) uses full adhesivity of \mathcal{C} ; the *sheaf characterisation* (Theorem 5.5) additionally requires a topological enrichment $T : \mathcal{C} \rightarrow \mathbf{Top}$ sending \mathcal{M} -morphisms to open embeddings, with compactness and normality of $T(I)$. The concrete category $\mathbf{Def}(\mathbf{M})$ satisfies all axioms except (T), giving every result except the sheaf characterisation. $\mathbf{Def}_c(\mathbf{M})$ additionally satisfies (T) and (A'), enabling the sheaf characterisation, but full adhesivity fails (Remark A.5), so cogerm gluing is not established in that setting. $\mathbf{Def}_c(\mathbf{M})$ is *not* regular ($f(x, y) = (x, xy)$ collapses the y -axis to a point with no definable coequaliser), but no result in this paper depends on regularity of $\mathbf{Def}_c(\mathbf{M})$.

Definition A.1 (Admissible pair). An *admissible pair* $(\mathcal{C}, \mathcal{M})$ consists of:

- (C1) A finitely complete, extensive (finite coproducts are disjoint and universal), regular (every morphism factors as regular epi + mono, and regular epis are stable under pullback) category \mathcal{C} ;
- (C2) A class $\mathcal{M} \subseteq \mathbf{Mono}(\mathcal{C})$ stable under pullback and composition, containing all isomorphisms.

An \mathcal{M} -covering family of an object B is a finite family $\{m_i : A_i \hookrightarrow B\}_{i=1}^n$ of \mathcal{M} -monomorphisms that is jointly epimorphic (i.e. any two morphisms $g, h : B \rightarrow C$ agreeing on every A_i must be equal; in concrete categories such as $\mathbf{Def}(\mathbf{M})$, this is equivalent to being jointly surjective).

Definition A.2 (Modular add-ons). Given an admissible pair $(\mathcal{C}, \mathcal{M})$ we define the following conditions:

- (C3') *Descent*: every \mathcal{M} -covering family is effective epimorphic (i.e. B is the colimit of the Čech nerve of the family).
- (A') *Open-cover VK*: for any finite family of \mathcal{M} -subobjects $\{m_i : A_i \hookrightarrow X\}$ of a common object X , the pushout $\bigcup_i A_i$ along the pairwise overlaps $A_i \times_X A_j$ exists in \mathcal{C} and satisfies the Van Kampen condition (pulling back along any morphism yields a pushout).
- (T) *Topological enrichment*: a faithful functor $T : \mathcal{C} \rightarrow \mathbf{Top}$ sending \mathcal{M} -morphisms to open embeddings.

Proposition A.3 (Subcanonical pretopology). *If $(\mathcal{C}, \mathcal{M})$ satisfies (C1)–(C2) and (C3'), then \mathcal{M} -covering families form a subcanonical Grothendieck pretopology on \mathcal{C} .*

Proof. We verify the pretopology axioms.

(T1) *Isomorphisms cover.* $\{\text{id}_B\}$ is an \mathcal{M} -family (isos $\in \mathcal{M}$ by (C2)) and is trivially jointly epimorphic.

(T2) *Stability under pullback.* Let $\{m_i : A_i \hookrightarrow B\}$ be a covering and $f : C \rightarrow B$ a morphism. The pullback projections $p_i : C \times_B A_i \rightarrow C$ are \mathcal{M} -monomorphisms (pullback of $m_i \in \mathcal{M}$, stable by (C2)). It remains to show $\{p_i\}$ is jointly epimorphic. By (C3'), $\{m_i\}$ is effective epimorphic; by extensivity, this is equivalent to the canonical map $c : \bigsqcup_i A_i \rightarrow B$ being an effective epimorphism (the Čech nerve of the family identifies with the kernel pair of c). In particular c is a regular epimorphism, hence stable under pullback by regularity: f^*c is a regular epi. By extensivity again, $f^*c \cong \bigsqcup_i (C \times_B A_i) \rightarrow C$ is the canonical map of $\{p_i\}$, and a regular epi in a regular category is effective, so $\{p_i\}$ is jointly epimorphic.

(T3) *Transitivity.* \mathcal{M} is closed under composition (C2); compositions of jointly epimorphic families are jointly epimorphic (if $g \circ m_i \circ n_{ij} = h \circ m_i \circ n_{ij}$ for all i, j , then $g \circ m_i = h \circ m_i$ by joint epimorphicity of $\{n_{ij}\}$, then $g = h$ by joint epimorphicity of $\{m_i\}$).

Subcanonicity. Effective epimorphic families satisfy descent by definition: compatible data on the covering glues uniquely. \square

Proposition A.4 (Open-cover VK for $\text{Def}_c(\mathbf{M})$). *Let \mathbf{M} be an o -minimal expansion of $(\mathbb{R}, <, +, \cdot)$, $\text{Def}_c(\mathbf{M})$ the category of definable sets with continuous definable maps, and $\mathcal{M} = \{\text{definable open immersions}\}$. Then $(\text{Def}_c(\mathbf{M}), \mathcal{M})$ satisfies (A').*

Proof. Let U_1, U_2 be definable open subsets of a definable set X , with $V = U_1 \cap U_2$. We show $P = U_1 \cup U_2$ (with the subspace topology from X) is the pushout of $V \hookrightarrow U_1$ and $V \hookrightarrow U_2$ in $\text{Def}_c(\mathbf{M})$.

Universal property. Given continuous definable $g_1 : U_1 \rightarrow Z$ and $g_2 : U_2 \rightarrow Z$ with $g_1|_V = g_2|_V$, define $g : P \rightarrow Z$ by $g(x) = g_i(x)$ for $x \in U_i$. This is well-defined on V , continuous by the pasting lemma (U_1 and U_2 are open in P), and definable: $\text{graph}(g) = \text{graph}(g_1) \cup \text{graph}(g_2)$ is a finite union of definable sets.

Van Kampen. The forgetful functor $U : \text{Def}_c(\mathbf{M}) \rightarrow \mathbf{Set}$ preserves finite limits. The set-theoretic pushout $U(P) = U(U_1) \cup U(U_2)$ coincides with the pushout in \mathbf{Set} , which satisfies VK by adhesivity of \mathbf{Set} . For any $f : W \rightarrow P$, the pullback gives definable open subsets $W_i = f^{-1}(U_i) \subseteq W$ covering W , with $W_1 \cap W_2 = f^{-1}(V)$. The VK equivalence at the set level (top face is a pushout \Leftrightarrow front/right faces are pullbacks) lifts to $\text{Def}_c(\mathbf{M})$: the unique set-theoretic mediating map is continuous (pasting lemma on the open cover $\{W_1, W_2\}$) and definable (piecewise definability on a definable open cover).

The proof handles the binary case $n = 2$; the general case of n subobjects follows by induction, taking $U_1 \cup \dots \cup U_{n-1}$ and U_n as the two pieces at each step (the intermediate union is a definable open subset of X , hence an \mathcal{M} -subobject). \square

Remark A.5. Full \mathcal{M} -adhesivity of $\text{Def}_c(\mathbf{M})$ fails: pushouts along an \mathcal{M} -mono $m : V \hookrightarrow U_1$ with an arbitrary continuous definable $f : V \rightarrow U_2$ need not satisfy VK. A counterexample: $V = (-1, 1)$, $U_1 = \mathbb{R}$, $U_2 = \mathbb{R}^2$, $m : V \hookrightarrow U_1$ the inclusion, $f : V \rightarrow U_2$ defined by $f(t) = (t^2 - 1, t(t^2 - 1))$ (nodal cubic parametrisation; f is injective on V but not an open immersion). As $t \rightarrow \pm 1$, $f(t) \rightarrow (0, 0)$; in the pushout $P = U_1 \sqcup_V U_2$, any continuous map out of P must send $\text{can}_1(\pm 1)$ and $\text{can}_2(0, 0)$ to the same point, so the Hausdorff quotient identifies them. The pullback $U_1 \times_P U_2$ is then $[-1, 1]$, not $(-1, 1) = V$: the pushout square is not a pullback. Property (A') avoids this by requiring *both* legs to be \mathcal{M} -monomorphisms.

B Adhesivity of $\text{Sys}_{\text{ho}}(I, O)$

Proof of adhesivity (Theorem 2.8). Let \mathcal{C} be adhesive with finite products, and let $U : \text{Sys}_{\text{ho}}(I, O) \rightarrow \mathcal{C}$ be the forgetful functor $(S, \alpha) \mapsto S$.

Pullbacks. U creates pullbacks: given $(S_1, \alpha_1) \xrightarrow{f} (S_3, \alpha_3) \xleftarrow{g} (S_2, \alpha_2)$, the pullback state space $P = S_1 \times_{S_3} S_2$ in \mathcal{C} acquires a unique dynamics $\alpha_P : P \times I \rightarrow P \times O$ making the projections into morphisms, because $- \times I$ and $- \times O$ preserve pullbacks.

Pushouts along monomorphisms. Let $m : (C, \gamma) \hookrightarrow (A, \alpha_A)$ be a mono (i.e. $m : C \hookrightarrow A$ is mono in \mathcal{C}) and $f : (C, \gamma) \rightarrow (B, \alpha_B)$ a morphism. The pushout $P = A \sqcup_C B$ exists in \mathcal{C} by adhesivity. The dynamics $\alpha_P : P \times I \rightarrow P \times O$ is defined using the canonical identification

$$P \times I \cong (A \times I) \sqcup_{C \times I} (B \times I),$$

which follows from the Van Kampen condition applied to the cube whose bottom face is the pushout $A \sqcup_C B$ and whose vertical maps are the product projections $- \times I \rightarrow -$ (all four vertical faces are pullbacks: the pullback of a projection $\pi_Y : Y \times Z \rightarrow Y$ along any $g : X \rightarrow Y$ is $X \times Z$). On the $A \times I$ component, α_P is $(\text{can}_A \times \text{id}_O) \circ \alpha_A$; on the $B \times I$ component, α_P is $(\text{can}_B \times \text{id}_O) \circ \alpha_B$. Compatibility on $C \times I$ follows from the morphism conditions on m and f ; maps out of the pushout $(A \times I) \sqcup_{C \times I} (B \times I)$ are determined by their restrictions, so the dynamics is unique.

Van Kampen condition. Given a commutative cube in $\text{Sys}_{\text{ho}}(I, O)$ whose bottom is a pushout along a mono and whose back and left faces are pullbacks, apply U : by adhesivity of \mathcal{C} , the top of the underlying cube is a pushout iff the front and right faces are pullbacks. Since U creates both pullbacks (above) and pushouts along monomorphisms—the distributivity identification determines unique dynamics on any such pushout—the equivalence lifts to $\text{Sys}_{\text{ho}}(I, O)$: the “if” direction holds because U creates pullbacks, and the “only if” direction holds because a pushout in $\text{Sys}_{\text{ho}}(I, O)$ is a pushout in \mathcal{C} equipped with the unique compatible dynamics. \square

No exponentiability hypothesis is used: the entire argument works with the uncurried dynamics $\alpha : S \times I \rightarrow S \times O$, never forming $(-)^I$. However, assuming exponentiability, the above is a special case of the general Theorem 2.9.

Proof of Theorem 2.9. Let $U : \text{Coalg}(H) \rightarrow \mathcal{C}$ be the forgetful functor.

Pullbacks. U creates pullbacks: given coalgebra morphisms $f : (A, \alpha) \rightarrow (C, \gamma)$ and $g : (B, \beta) \rightarrow (C, \gamma)$, the pullback $P = A \times_C B$ in \mathcal{C} acquires a unique coalgebra structure $\pi : P \rightarrow H(P)$ via $H(P) \cong H(A) \times_{H(C)} H(B)$ (using that H preserves pullbacks).

Monomorphisms. Since U is faithful, it reflects monomorphisms. Conversely, if $f : (A, \alpha) \rightarrow (B, \beta)$ is mono in $\text{Coalg}(H)$, its kernel pair (R, π_1, π_2) in \mathcal{C} lifts to a coalgebra kernel pair by pullback creation, so $\pi_1 = \pi_2$ gives $U(f)$ mono in \mathcal{C} . Thus monomorphisms in $\text{Coalg}(H)$ correspond to monomorphisms in \mathcal{C} .

Pushouts along monomorphisms. Let $m : (C, \gamma) \hookrightarrow (A, \alpha_A)$ be a mono and $f : (C, \gamma) \rightarrow (B, \alpha_B)$ a morphism. The pushout $P = A \sqcup_C B$ exists in \mathcal{C} by adhesivity. The colimit cocone α_A, α_B induces a coalgebra structure on P via the universal property: the cocone $D(j) \xrightarrow{\alpha_j} H(D(j)) \xrightarrow{H(t_j)} H(P)$ is compatible by the coalgebra morphism conditions on m and f , so the universal property yields $\alpha_P : P \rightarrow H(P)$. Uniqueness: a map out of the pushout is determined by its restrictions to A and B , and the coalgebra morphism conditions determine these.

Van Kampen condition. Given a commutative cube in $\text{Coalg}(H)$ whose bottom is a pushout along a mono and whose back and left faces are pullbacks, apply U : by adhesivity of \mathcal{C} , the top of the underlying

cube is a pushout iff the front and right faces are pullbacks. This equivalence lifts to $\text{Coalg}(H)$: the “if” direction uses pullback creation, and the “only if” direction uses that a pushout in $\text{Coalg}(H)$ maps to a pushout in \mathcal{C} equipped with the unique compatible coalgebra structure. \square

For $\mathcal{C} = \mathbf{Set}$, this result was established by Padberg [19] in the framework of \mathcal{M} -adhesive categories; the general statement for an arbitrary adhesive base appears to be new. Whether the pullback-preservation properties we used can be weakened to weak pullbacks—which would cover many more functors of coalgebraic interest [13]—remains open.

C Proof of Theorem 5.1

Proof. The proof uses adhesivity of $\text{Def}(\mathbf{M})$ (§2.2) and the fact that all systems in the construction share the interface (I', O') .

Step 1: Choose representatives and overlap witnesses. Choose representatives $(\mathcal{S}'_\alpha, \Psi_\alpha)$ from each equivalence class. Compatibility on overlaps means: for each pair (α, β) , the restrictions $(\mathcal{S}'_\alpha, \Psi_\alpha \circ n_\alpha)$ and $(\mathcal{S}'_\beta, \Psi_\beta \circ n_\beta)$ are cogerm-equivalent in $\mathcal{F}_j^{\text{cogerm}}(\mathcal{S}_\alpha \cap \mathcal{S}_\beta)$, where $n_\alpha : \mathcal{S}_\alpha \cap \mathcal{S}_\beta \hookrightarrow \mathcal{S}_\alpha$. So there exists a system $\mathcal{S}''_{\alpha\beta}$ with interface (I', O') and monomorphisms $i_\alpha : \mathcal{S}''_{\alpha\beta} \hookrightarrow \mathcal{S}'_\alpha$, $i_\beta : \mathcal{S}''_{\alpha\beta} \hookrightarrow \mathcal{S}'_\beta$ with identity interface components, and a morphism $\varphi_{\alpha\beta} : \mathcal{S}_\alpha \cap \mathcal{S}_\beta \rightarrow \mathcal{S}''_{\alpha\beta}$ with $i_\alpha \circ \varphi_{\alpha\beta} = \Psi_\alpha \circ n_\alpha$ and $i_\beta \circ \varphi_{\alpha\beta} = \Psi_\beta \circ n_\beta$.

Step 2: Construct the global target by colimit. Form the system \mathcal{S}' as the colimit of the diagram

$$\dots \mathcal{S}'_\alpha \xleftarrow{i_\alpha} \mathcal{S}''_{\alpha\beta} \xrightarrow{i_\beta} \mathcal{S}'_\beta \dots$$

in Sys . Concretely, $S'_b = \text{colim}(S'_{\alpha,b} \leftarrow S''_{\alpha\beta,b} \rightarrow S'_{\beta,b})$, and similarly for S'_a ; since all \mathcal{S}'_α and $\mathcal{S}''_{\alpha\beta}$ are homogeneous and all morphisms lie in $\text{Sys}_{\text{ho}}(I', O')$, the two colimit diagrams coincide, so $S'_b = S'_a$ and \mathcal{S}' is homogeneous. The interface of \mathcal{S}' is (I', O') since all pieces share this interface and all identifications are identity on interface. The colimit exists in $\text{Def}(\mathbf{M})$: it is the quotient of $\bigsqcup_\alpha S'_{\alpha,b}$ by the definable equivalence relation generated by $s \sim s'$ whenever $s = (i_\alpha)_b(t)$ and $s' = (i_\beta)_b(t)$ for some $t \in S''_{\alpha\beta,b}$. Dynamics compatibility propagates through the generators of this equivalence relation: if s and s' are identified via $t \in S''_{\alpha\beta,b}$, their images under the dynamics are identified via the same t (since $\mathcal{S}''_{\alpha\beta}$ is a subsystem), so the relation is compatible with the dynamics by induction on chain length.

Since the covering is finite, this colimit is constructed as iterated binary pushouts along monomorphisms. By adhesivity, each binary pushout along a mono preserves monos, so inductively the canonical maps $\text{can}_\alpha : \mathcal{S}'_\alpha \hookrightarrow \mathcal{S}'$ are monomorphisms with identity interface components.

The dynamics $\alpha' : S'_b \times I' \rightarrow S'_a \times O'$ is well-defined on the quotient: if $s = (i_\alpha)_b(t)$ is identified with $s' = (i_\beta)_b(t)$ for $t \in S''_{\alpha\beta,b}$, then $\alpha'_\alpha(s, i')$ and $\alpha'_\beta(s', i')$ are identified in S'_a by the same equivalence relation, since $\mathcal{S}''_{\alpha\beta}$ is a subsystem of both \mathcal{S}'_α and \mathcal{S}'_β .

Step 3: Glue the morphisms. The compositions $\text{can}_\alpha \circ \Psi_\alpha : \mathcal{S}_\alpha \rightarrow \mathcal{S}'$ form a compatible family: for $x \in \mathcal{S}_\alpha \cap \mathcal{S}_\beta$,

$$(\text{can}_\alpha \circ \Psi_\alpha)(n_\alpha(x)) = \text{can}_\alpha(i_\alpha(\varphi_{\alpha\beta}(x))) = \text{can}_\beta(i_\beta(\varphi_{\alpha\beta}(x))) = (\text{can}_\beta \circ \Psi_\beta)(n_\beta(x)),$$

where the middle equality holds by the defining property of the colimit. Since the covering is jointly surjective on $S_b \times I$ and $S_a \times O$, the compatible family defines a unique morphism $\Psi : \mathcal{S} \rightarrow \mathcal{S}'$: compatibility ensures well-definedness on overlaps, and the dynamics square holds globally because it holds on each patch.

Step 4: Verify interface and restriction. On each \mathcal{S}_α , $\psi_I|_{I_\alpha} = (\text{can}_\alpha)_I \circ (j_I \circ m_{I,\alpha}) = j_I \circ m_{I,\alpha}$ since can_α has identity interface. Joint surjectivity gives $\psi_I = j_I$; similarly $\psi_O = j_O$. Hence $(\mathcal{S}', \psi) \in \mathcal{F}_j^{\text{cogerm}}(\mathcal{S})$.

The restriction to \mathcal{S}_α satisfies $\psi \circ m_\alpha = \text{can}_\alpha \circ \psi_\alpha$. Since $\text{can}_\alpha : \mathcal{S}'_\alpha \hookrightarrow \mathcal{S}'$ is a monomorphism with identity interface, taking $\mathcal{S}'' = \mathcal{S}'_\alpha$, $i_1 = \text{can}_\alpha$, $i_2 = \text{id}$, $\varphi = \psi_\alpha$ witnesses $(\mathcal{S}', \text{can}_\alpha \circ \psi_\alpha) \sim (\mathcal{S}'_\alpha, \psi_\alpha)$ in $\mathcal{F}_j^{\text{cogerm}}(\mathcal{S}_\alpha)$. \square

D Proof of Theorem 5.5

Proof. (\Rightarrow): robustly disconnected \Rightarrow not a sheaf. Suppose $j_I^{-1}(t_0)$ is robustly disconnected, with $j_I^{-1}(N) \subseteq V \cup W$. Pick a closed definable $\bar{N}_0 \subset N$ with $t_0 \in \text{int}(\bar{N}_0)$, and set $C := I \setminus j_I^{-1}(\bar{N}_0)$. Since $\text{im}(j_O)$ is definably connected with ≥ 2 points, some definably connected component C_0 of O has $j_O|_{C_0}$ non-constant. Pick $o_0, o_1 \in C_0$ with $j_O(o_0) \neq j_O(o_1)$, connected by a definable arc $\tilde{\gamma} : [0, 1] \rightarrow C_0$ with $\tilde{\gamma}(0) = o_0$, $\tilde{\gamma}(1) = o_1$, and let $\beta : \bar{N}_0 \rightarrow [0, 1]$ be a definable bump with $\beta|_{\partial\bar{N}_0} = 0$, $\beta(t_0) = 1$ (existence: normality of compact definable sets, cf. [6, Ch. 6]). Define $f : I \rightarrow O$ by $f(x) = o_0$ on $V \cap j_I^{-1}(\bar{N}_0)$ and $I \setminus j_I^{-1}(\bar{N}_0)$, and $f(x) = \tilde{\gamma}(\beta(j_I(x)))$ on $W \cap j_I^{-1}(\bar{N}_0)$. Then f is continuous (the V and W parts are clopen in $j_I^{-1}(\bar{N}_0)$, and $f \rightarrow o_0$ at the boundary). On U_1 , f is constant o_0 , so $j_O \circ f|_{U_1}$ factors through j_I . On U_2 , the $W \cap j_I^{-1}(\bar{N}_0)$ part depends on $j_I(x)$ only, the C part is constant, and no point in $W \cap j_I^{-1}(\bar{N}_0)$ shares a j_I -value with any point in C , so factorisation holds. The covering $U_1 := V \cup C$, $U_2 := W \cup C$ gives two sections $j_O \circ f|_{U_i}$ that agree on $U_1 \cap U_2 = C$ but satisfy $j_O(f(a)) = j_O(o_0) \neq j_O(o_1) = j_O(f(b))$ for $a \in V \cap j_I^{-1}(t_0)$, $b \in W \cap j_I^{-1}(t_0)$.

(\Leftarrow): no robust disconnection \Rightarrow sheaf. Suppose $h : I \rightarrow O'$ is locally $\ker(j_I)$ -invariant on a covering $\{I_\alpha\}$ but not globally: there exist x, y with $j_I(x) = j_I(y) = t_0$ and $h(x) \neq h(y)$. Since h is constant on $j_I^{-1}(t_0) \cap I_\alpha$ for each α , the function $h|_{j_I^{-1}(t_0)}$ is locally constant with clopen level sets. Set $A := \{z \in j_I^{-1}(t_0) : h(z) = h(x)\}$ and $B := j_I^{-1}(t_0) \setminus A$ (both non-empty, disjoint, closed in I). By normality of I (compact Hausdorff), there exist disjoint open $V \supseteq A$, $W \supseteq B$. Since $S := I \setminus (V \cup W)$ is closed with $S \cap j_I^{-1}(t_0) = \emptyset$, compactness gives $t_0 \notin j_I(S)$, so some open $N \ni t_0$ satisfies $j_I^{-1}(N) \subseteq V \cup W$ —a robust disconnection. Since every locally $\ker(j_I)$ -invariant function that is not globally invariant yields a robust disconnection, the contrapositive gives: no robust disconnection implies every locally invariant function is globally invariant, i.e. $\mathcal{F}_{j,1}^{\text{ri}}$ is a sheaf. \square

E Index of Notation

Categories and spaces

\mathcal{C}	Base category (typically adhesive with finite products)
$\text{Def}(\mathbf{M})$	Category of definable sets in the o-minimal structure \mathbf{M}
$\text{Def}_c(\mathbf{M})$	Wide subcategory of $\text{Def}(\mathbf{M})$: continuous definable maps
$\text{Coalg}(H)$	Category of coalgebras for the endofunctor H
Sys	Category of heterogeneous Mealy machines (varying interfaces)
$\text{Sys}_{\text{ho}}(I, O)$	Non-full subcategory: homogeneous systems, $f_b = f_a$
Sys_{oi}	Wide subcategory of Sys : definable open immersions
Set	Category of sets

Systems

(S_b, S_a, I, O, α)	Heterogeneous Mealy machine (system)
$\alpha : S_b \times I \rightarrow S_a \times O$	Dynamics map
$\iota : \mathcal{S}' \hookrightarrow \mathcal{S}$	Definable open immersion of systems

Judge and interface

$j = (j_I, j_O)$	Judge: morphisms to an interpretable interface (I', O')
I'_U	Restricted interface $\text{im}(j_I \circ m_I)$ for subsystem \mathcal{S}_U

Presheaves and equivalences

$\mathcal{F}_j^{\text{cogerm}}$	Cogerm-quotiented presheaf of explanations
$\mathcal{F}_j^{\text{beh}}$	Behavioral presheaf (quotient by \approx ; separated)
$\mathcal{F}_j^{\text{ri}}$	Restricted-interface presheaf
$\mathcal{F}_{j,1}^{\text{ri}}$	Stateless subpresheaf of $\mathcal{F}_j^{\text{ri}}$ (sections with $ S' = 1$)
$\tilde{\mathcal{F}}_j$	Unquotiented presheaf of explanations (sheaf)
$\mathcal{F}_j^\varepsilon$	Presheaf of ε -approximate sections
\sim	Cogerm equivalence on sections
\approx	Behavioral equivalence on sections