The Geometry of Admissibility

Quotients, Lattices, Sheaves, and Functorial Reachability

Flyxion

Independent Researcher

June 2026

Abstract

A representation of a state space is *admissible* for planning if it does not collapse states with different reachable futures. This condition generates a rich mathematical structure that has not been developed in full. The present paper does so. We begin with a self-contained treatment of admissibility structures: the reachability equivalence relation $\sim_A$, the admissibility quotient $\mathcal{X}/\sim_A$, and the admissibility distortion $D_A(\varphi)$ measuring departure from the quotient that planning requires. From these primitive notions we build four interlocking geometric objects. First, the *category* **Adm** of admissibility structures, in which admissible representations are exactly the morphisms and the admissibility quotient $q_A \colon \mathcal{X} \to \mathcal{X}/\sim_A$ is the initial object in the subcategory of admissible projections—a universal property that is strictly stronger than calling q_A the "largest safe equivalence relation." Second, the *quotient lattice* $\mathcal{L}(\mathcal{X})$ of equivalence relations on $\mathcal{X}$ ordered by coarseness, on which admissibility distortion is a monotone height function and every compression morphism is a step toward higher distortion; the compression–admissibility tradeoff is a functor from the category of encoders to the poset of non-negative reals. Third, the *reachability sheaf* $\mathcal{F}_A$ over the Alexandrov topology τ_A generated by reachability sets: its global sections are exactly the planning-adequate representations, while the observation presheaf $\mathcal{F}_O$ over the coarser observation topology τ_O yields only local sections; the impossibility of assembling local observational data into a complete reachability description is a failure of the sheaf gluing condition, and we propose it as the geometric form of the Observational–Interventional Separation Theorem. Fourth, *temporal admissibility functors* $\mathbf{X} \colon \mathbf{T} \to \mathbf{Adm}$ representing evolving state spaces: meta-admissibility—the requirement that self-modifications preserve reachability structure—is equivalent to the functor condition, and comparisons between developmental trajectories are natural transformations. We conclude with a precise conjecture relating admissibility distortion to the first Čech cohomology of the reachability sheaf, and with an open problems section surveying the cohomological, topos-theoretic, and persistent-homological directions that the framework opens.

Contents

1. Introduction

Planning in a learned state space requires the representation to preserve a specific kind of information: it must not identify states that differ in which futures are reachable under the available policies. A representation satisfying this condition is called *admissible*. One that violates it incurs *admissibility distortion*: it conflates states that are behaviourally distinct, causing plans constructed in the representation space to fail when executed in the environment.

The admissibility condition is a natural object of study independently of its motivating context. It defines an equivalence relation on state spaces, an induced quotient, a notion of morphism between structured spaces, a partial order on representations, and a quantitative functional measuring deviation from the ideal. Each of these objects belongs to a classical mathematical tradition—category theory, lattice theory, sheaf theory—and each connection yields results that would be invisible from the planning-theoretic perspective alone.

The purpose of this paper is to develop these connections systematically. The paper is entirely self-contained: all notions from the admissibility programme are defined from scratch, and the main results are proved in full without external reference. The mathematical prerequisites are basic category theory at the level of the standard graduate texts on the subject and elementary sheaf theory; both are recalled where needed. No background in reinforcement learning or representation learning is assumed.

The paper is organised as follows. Section 2 develops the admissibility framework from primitive notions: state spaces, policy classes, reachability, the admissibility equivalence relation, the quotient, and distortion. Section 3 constructs the category **Adm** of admissibility structures and proves the universal property of the admissibility quotient. Section 4 develops the quotient-lattice structure and the compression–admissibility tradeoff as a functor. Section 5 constructs the reachability sheaf over the Alexandrov topology, defines the observation presheaf, and proves the sheaf formulation of the Observational–Interventional Separation Theorem. Section 6 proposes the cohomological interpretation of admissibility distortion and states the central open conjecture. Section 7 treats evolving state spaces via temporal admissibility functors and characterises meta-admissibility as a functor condition. Section 8 assembles the four structures into a unified geometric picture. Section 9 surveys open problems.

Notation. Throughout, $\mathcal{X}$ denotes a state space (a set), $\mathcal{M}$ a representation space, Π a policy class, $H \in \mathbb{N}$ a planning horizon, $\mathcal{P}(\mathcal{X})$ the power set of $\mathcal{X}$, and $(\mathbb{R}_{\geq 0}, \leq)$ the poset of non-negative reals. Category-theoretic notation follows standard usage: objects, morphisms, composition $\circ$, identity id.

Motivating example: small-object detection. The admissibility condition appears concretely in visual inspection systems. In industrial fault detection, small defects—cracked insulators, loose bolts—may occupy regions of an image that are statistically nearly indistinguishable from background texture after standard compression and pooling operations. Yet their reachable operational futures differ dramatically: one leads to continued safe operation, the other to equipment failure. Standard convolutional encoders, optimised for reconstruction or classification accuracy, may collapse these visually similar regions into identical or nearby representation codes. Recent architectures for small-object detection address this by explicitly augmenting the encoder with edge-sensitive features and multi-scale context, reintroducing precisely the distinctions that would otherwise be lost. From the admissibility geometry perspective, the admissibility-critical states occupy a low-measure but high-consequence region of the quotient lattice: they lie just above the admissibility threshold $\sim_A$, and any compression that passes through this boundary incurs positive admissibility distortion. The architectural interventions can be understood as admissibility-preserving design choices that prevent the encoder from crossing this boundary.

2. Foundations: Admissibility Structures

2.1. State spaces and reachability

Definition 2.1 (State Space and Policy Class)**.** A *state space* is a set $\mathcal{X}$. A *policy class* Π is a collection of functions $\pi \colon \mathcal{X} \to \Delta(\mathcal{X})$, where $\Delta(\mathcal{X})$ denotes the set of probability distributions over $\mathcal{X}$. Each $\pi \in \Pi$ is a *policy*: a map from states to distributions over successor states.

Definition 2.2 (Reachability Set)**.** Let $H \in \mathbb{N}$, Π a policy class. The *H-step reachability set* of a state $x \in \mathcal{X}$ under Π is

$$\mathcal{R}_H^{\Pi}(x) \;=\; \big\{x' \in \mathcal{X} \;\big|\; \exists\, \pi \in \Pi,\ \exists\,\text{path } x = x_0, x_1, \dots, x_H = x' \text{ with } x_{t+1} \in \operatorname{supp}(\pi(x_t))\ \forall t\big\}.$$

We sometimes write $\mathcal{R}(x)$ when H and Π are fixed.

The reachability set $\mathcal{R}_H^{\Pi}(x)$ is the set of states that can be reached from x in exactly H steps by some policy in Π. Two states are *behaviourally distinct* at horizon H if they differ in which states they can reach.

2.2. The admissibility equivalence relation

Definition 2.3 (Admissibility Equivalence)**.** The *admissibility equivalence relation* $\sim_A$ on $\mathcal{X}$ is defined by

$$x_1 \sim_A x_2 \iff \mathcal{R}_H^{\Pi}(x_1) = \mathcal{R}_H^{\Pi}(x_2).$$

The equivalence class of x under $\sim_A$ is denoted $[x]_A$.

Lemma 2.4. *$\sim_A$ is an equivalence relation on $\mathcal{X}$.*

Proof. Reflexivity: $\mathcal{R}_H^\Pi(x) = \mathcal{R}_H^\Pi(x)$. Symmetry: if $\mathcal{R}_H^\Pi(x_1) = \mathcal{R}_H^\Pi(x_2)$ then $\mathcal{R}_H^\Pi(x_2) = \mathcal{R}_H^\Pi(x_1)$. Transitivity: if $\mathcal{R}_H^\Pi(x_1) = \mathcal{R}_H^\Pi(x_2)$ and $\mathcal{R}_H^\Pi(x_2) = \mathcal{R}_H^\Pi(x_3)$ then $\mathcal{R}_H^\Pi(x_1) = \mathcal{R}_H^\Pi(x_3)$. □

2.3. Admissible representations and distortion

Definition 2.5 (Admissible Representation)**.** A function $\varphi\colon \mathcal{X} \to \mathcal{M}$ is an *admissible representation* (or *admissible encoder*) if

$$\varphi(x_1) = \varphi(x_2) \implies x_1 \sim_A x_2,$$

i.e., φ does not merge states with different reachable futures. Equivalently, $x_1 \not\sim_A x_2 \Rightarrow \varphi(x_1) \neq \varphi(x_2)$: admissibility-inequivalent states receive distinct codes.

Remark 2.6. Admissibility is an injectivity condition relative to the admissibility equivalence classes: φ must be injective on the set of equivalence classes $\mathcal{X}/{\sim_A}$. It does not require φ to be injective on $\mathcal{X}$ itself; it may merge admissibility-equivalent states.

The *morphism defect* of an arbitrary (possibly inadmissible) encoder φ is the set of pairs that φ collapses but $\sim_A$ does not:

$$\mathrm{Def}(\varphi) \;=\; \big\{(x_1, x_2) \in \mathcal{X} \times \mathcal{X} \;\big|\; \varphi(x_1) = \varphi(x_2),\ x_1 \not\sim_A x_2\big\}.$$

φ is admissible if and only if $\mathrm{Def}(\varphi) = \emptyset$.

Definition 2.7 (Admissibility Distortion)**.** Let $d\colon \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}_{\geq 0}$ be a pseudometric on subsets of $\mathcal{X}$ (for example, the symmetric difference pseudometric or the Hausdorff metric). The *admissibility distortion* of φ is

$$D_A(\varphi) \;=\; \mathbb{E}_{(x_1,x_2)\sim\mu}\big[\mathbf{1}[\varphi(x_1) = \varphi(x_2)] \cdot d(\mathcal{R}_H^\Pi(x_1),\, \mathcal{R}_H^\Pi(x_2))\big],$$

where μ is a distribution on $\mathcal{X} \times \mathcal{X}$. The expectation sums the reachability separation over collapsed pairs.

Proposition 2.8 (Zero-Distortion Characterisation)**.** $D_A(\varphi) = 0$ *if and only if* $\mathrm{Def}(\varphi) = \emptyset$ *almost surely under* μ*, i.e., if and only if* φ *is admissible* μ*-almost surely.*

Proof. Since d is a pseudometric, $d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2)) = 0$ iff $\mathcal{R}_H^\Pi(x_1) = \mathcal{R}_H^\Pi(x_2)$ iff $x_1 \sim_A x_2$. The integrand $\mathbf{1}[\varphi(x_1) = \varphi(x_2)] \cdot d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2))$ is non-negative and vanishes on a pair (x_1, x_2) iff either $\varphi(x_1) \neq \varphi(x_2)$ (the pair is not collapsed) or $x_1 \sim_A x_2$ (the collapsed pair is admissibility-equivalent). Hence the integral vanishes iff there are no collapsed, admissibility-inequivalent pairs μ-almost surely. □

2.4. The admissibility quotient

Definition 2.9 (Admissibility Quotient)**.** The *admissibility quotient* is the quotient set $\mathcal{X}/{\sim_A} = \mathcal{X}/{\sim_A}$ together with the canonical surjection $q_A\colon \mathcal{X} \to \mathcal{X}/{\sim_A}$, $q_A(x) = [x]_A$.

The admissibility quotient is the "most compressed" admissible representation of $\mathcal{X}$: it merges exactly the pairs that admissibility permits merging, and no others. This informal description will be given precise content by the universal property in Section 3.

Example 2.10 (Finite Grid)**.** Let $\mathcal{X} = \{1, \ldots, n\}^2$ be a grid world, Π the class of deterministic policies, $H = 1$. States x_1 and x_2 are admissibility-equivalent iff their sets of adjacent cells coincide. Interior states with the same neighbourhood pattern are equivalent; corner and edge states with distinct adjacency structures are inequivalent. The admissibility quotient identifies states by their local connectivity type.

Example 2.11 (Observation-Equivalent but Reachability-Distinct States)**.** Let $\mathcal{X} = \{s_0^0, s_0^1, s_{\text{safe}}, s_{\text{danger}}\}$. Both s_0^0 and s_0^1 have identical passive observation distributions—they are *observationally equivalent*—but $\mathcal{R}_H^\Pi(s_0^0) = \{s_{\text{safe}}\}$ and $\mathcal{R}_H^\Pi(s_0^1) = \{s_{\text{danger}}\}$. Hence $s_0^0 \not\sim_A s_0^1$ despite their observational indistinguishability. Any admissible encoder must assign them distinct codes. This is the canonical minimal instance of the Observational–Interventional Separation phenomenon, developed at length in Section 5.

3. The Category of Admissibility Structures

3.1. Objects and morphisms

Definition 3.1 (Admissibility Structure)**.** An *admissibility structure* is a pair $(\mathcal{X}, \sim_A)$ where $\mathcal{X}$ is a set and $\sim_A$ is the admissibility equivalence relation on $\mathcal{X}$ generated by a reachability map $x \mapsto \mathcal{R}_H^\Pi(x)$ as in Definition 2.3. We will sometimes write $(\mathcal{X}, \mathcal{R})$ to emphasise the reachability map, with $\sim_A$ understood as the kernel of $\mathcal{R}$.

Definition 3.2 (The Category **Adm**)**.** The category **Adm** has admissibility structures as objects. A *morphism* $f\colon (\mathcal{X}, \sim_A) \to (\mathcal{X}', \sim_A')$ is a function $f\colon \mathcal{X} \to \mathcal{X}'$ satisfying the *congruence condition*:

$$x_1 \sim_A x_2 \implies f(x_1) \sim_A' f(x_2).$$

Composition is ordinary function composition; the identity morphism on $(\mathcal{X}, \sim_A)$ is $\mathrm{id}_{\mathcal{X}}$.

Lemma 3.3 (**Adm** is a Category)**.** **Adm** *as defined above satisfies the axioms of a category.*

Proof. Identities. $\mathrm{id}_{\mathcal{X}}\colon (\mathcal{X}, \sim_A) \to (\mathcal{X}, \sim_A)$ satisfies $x_1 \sim_A x_2 \Rightarrow \mathrm{id}_{\mathcal{X}}(x_1) = x_1 \sim_A x_2 = \mathrm{id}_{\mathcal{X}}(x_2)$, so the congruence condition holds trivially.

Composition. If $f\colon (\mathcal{X}, \sim_A) \to (\mathcal{X}', \sim_A')$ and $g\colon (\mathcal{X}', \sim_A') \to (\mathcal{X}'', \sim_A'')$ are morphisms, then $x_1 \sim_A x_2 \Rightarrow f(x_1) \sim_A' f(x_2) \Rightarrow g(f(x_1)) \sim_A'' g(f(x_2))$, so $g \circ f$ satisfies the congruence condition.

Unit laws and associativity follow from the corresponding properties of function composition. □

3.2. Admissible representations as morphisms: a precise statement

The relationship between the admissibility condition of Definition 2.5 and morphisms in **Adm** requires care. The two conditions are logically distinct: Definition 2.5 is an injectivity condition (equal codes imply admissibility-equivalent states), while the morphism condition is a congruence condition (admissibility-equivalent states have equivalent images). We now show precisely how they are related.

Definition 3.4 (Induced Equivalence on the Target)**.** Given $\varphi\colon \mathcal{X} \to \mathcal{M}$, define the equivalence relation $\sim_A^{\mathcal{M}}$ on $\mathcal{M}$ by:

$$m_1 \sim_A^{\mathcal{M}} m_2 \iff \exists\, x_1, x_2 \in \mathcal{X}\colon \varphi(x_1) = m_1,\ \varphi(x_2) = m_2,\ x_1 \sim_A x_2.$$

This is the equivalence relation on $\mathcal{M}$ generated by identifying image points of admissibility-equivalent states.

Proposition 3.5 (Admissible Representations and Morphisms)**.** *Let $\varphi\colon \mathcal{X} \to \mathcal{M}$ and let $\sim_A^{\mathcal{M}}$ be the induced equivalence on $\mathcal{M}$ as above.*

(i) φ is a morphism $(\mathcal{X}, \sim_A) \to (\mathcal{M}, \sim_A^{\mathcal{M}})$ in **Adm** *for any choice of φ.*
(ii) φ is admissible (Definition 2.5) if and only if φ is injective on admissibility classes: the map $\bar{\varphi}\colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ given by $\bar{\varphi}([x]_A) = \varphi(x)$ is injective.
(iii) If φ is admissible, then φ is a morphism from $(\mathcal{X}, \sim_A)$ to $(\mathcal{M}, \sim_\varphi)$, where $\sim_\varphi$ is the kernel equivalence $m_1 \sim_\varphi m_2 \iff m_1 = m_2$ (the discrete equivalence on $\mathrm{im}(\varphi)$).

Proof. (i). If $x_1 \sim_A x_2$, then by definition $\varphi(x_1) \sim_A^{\mathcal{M}} \varphi(x_2)$ (taking x_1, x_2 as witnesses). So the congruence condition holds universally.

(ii). $\bar{\varphi}$ is well-defined by any function: $\bar{\varphi}([x]_A) = \varphi(x)$ requires checking that if $[x_1]_A = [x_2]_A$ (i.e., $x_1 \sim_A x_2$) then $\varphi(x_1) = \varphi(x_2)$. This is exactly the admissibility condition: $\varphi(x_1) = \varphi(x_2) \Rightarrow x_1 \sim_A x_2$ is contraposited to $x_1 \not\sim_A x_2 \Rightarrow \varphi(x_1) \neq \varphi(x_2)$, which means the map $[x]_A \mapsto \varphi(x)$ is well-defined and injective. Conversely, if $\bar{\varphi}$ is injective and well-defined, then $\varphi(x_1) = \varphi(x_2)$ implies $\bar{\varphi}([x_1]_A) = \bar{\varphi}([x_2]_A)$, hence $[x_1]_A = [x_2]_A$ by injectivity, hence $x_1 \sim_A x_2$.

(iii). If φ is admissible, then $x_1 \sim_A x_2 \Rightarrow \varphi(x_1) = \varphi(x_2)$: admissibility-equivalent states map to identical codes. Identical codes are trivially equivalent under the discrete equivalence on $\mathrm{im}(\varphi)$, so the congruence condition is satisfied. □

Remark 3.6. Part (i) of Proposition 3.5 shows that every function $\varphi\colon \mathcal{X} \to \mathcal{M}$ is a morphism from $(\mathcal{X}, \sim_A)$ to $(\mathcal{M}, \sim_A^{\mathcal{M}})$ once the target is equipped with the induced equivalence. The interesting content of admissibility is (ii): it characterises when φ factors injectively through the admissibility quotient. The morphism condition in **Adm** is a congruence condition, while admissibility is an injectivity condition; they are equivalent only when the target equivalence is the discrete one (equal images), which is the case for admissible encoders by (iii).

3.3. The universal property of the admissibility quotient

The universal property of q_A classifies maps that *factor through* the quotient. A map $\varphi \colon \mathcal{X} \to \mathcal{M}$ factors through q_A at the set level—meaning $\varphi = \bar{\varphi} \circ q_A$ for some $\bar{\varphi} \colon \mathcal{X}/{\sim_A} \to \mathcal{M}$—if and only if φ is *constant on admissibility classes*: $x_1 \sim_A x_2 \Rightarrow \varphi(x_1) = \varphi(x_2)$. We call such maps *class-constant*. Admissibility (Definition 2.5) is a different, dual condition: a map is admissible iff it is *injective* on admissibility classes, meaning the induced map $\bar{\varphi} \colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ is injective. The universal property governs factoring; admissibility governs the injectivity of the factor.

Theorem 3.7 (Universal Property of q_A)**.** *For any class-constant map $\varphi \colon \mathcal{X} \to \mathcal{M}$ (i.e., satisfying $x_1 \sim_A x_2 \Rightarrow \varphi(x_1) = \varphi(x_2)$), there exists a unique map $\bar{\varphi} \colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ such that the diagram*

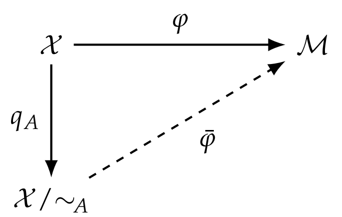

(1)

commutes: $\varphi = \bar{\varphi} \circ q_A$. Moreover, $\bar{\varphi}$ is a morphism $(\mathcal{X}/{\sim_A}, \sim_A^*) \to (\mathcal{M}, \sim_A^{\mathcal{M}})$ in **Adm***, where $\sim_A^*$ is the equivalence on $\mathcal{X}/{\sim_A}$ defined by $[x_1]_A \sim_A^* [x_2]_A \iff x_1 \sim_A x_2$, and $\sim_A^{\mathcal{M}}$ is the induced equivalence of Definition 3.4.*

Proof. Well-definedness. Define $\bar{\varphi}([x]_A) = \varphi(x)$. If $[x_1]_A = [x_2]_A$ then $x_1 \sim_A x_2$, hence $\varphi(x_1) = \varphi(x_2)$ by class-constancy of φ, so $\bar{\varphi}$ is well-defined.

Commutativity. $\bar{\varphi}(q_A(x)) = \bar{\varphi}([x]_A) = \varphi(x)$.

Uniqueness. Any $\psi \colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ with $\psi \circ q_A = \varphi$ must satisfy $\psi([x]_A) = \varphi(x)$ for all x, since every element of $\mathcal{X}/{\sim_A}$ is $[x]_A$ for some x. Hence $\psi = \bar{\varphi}$.

Morphism condition. If $[x_1]_A \sim_A^* [x_2]_A$ in $\mathcal{X}/{\sim_A}$, then $x_1 \sim_A x_2$, so by definition $\varphi(x_1) \sim_A^{\mathcal{M}} \varphi(x_2)$, i.e., $\bar{\varphi}([x_1]_A) \sim_A^{\mathcal{M}} \bar{\varphi}([x_2]_A)$. □

Proposition 3.8 (Admissibility as Injectivity of the Factor)**.** *Let $\varphi \colon \mathcal{X} \to \mathcal{M}$ be class-constant, and let $\bar{\varphi} \colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ be the unique factor of Theorem 3.7. Then φ is admissible (Definition 2.5) if and only if $\bar{\varphi}$ is injective.*

Proof. $\bar{\varphi}$ injective means: $\bar{\varphi}([x_1]_A) = \bar{\varphi}([x_2]_A) \Rightarrow [x_1]_A = [x_2]_A$, i.e., $\varphi(x_1) = \varphi(x_2) \Rightarrow x_1 \sim_A x_2$. This is exactly Definition 2.5. □

The correct categorical picture is therefore:

$$\mathcal{X} \xrightarrow{q_A} \mathcal{X}/{\sim_A} \xrightarrow{\bar{\varphi}} \mathcal{M},$$

where q_A is the universal quotient map, $\bar{\varphi}$ is the unique factor (existing for all class-constant φ), and admissibility is precisely the condition that $\bar{\varphi}$ is injective. The admissibility quotient q_A is not merely the "largest safe equivalence relation"; it is the object through which every class-constant representation factors, and admissible representations are those for which this factoring is injective.

Corollary 3.9 (Initiality of q_A)**.** *In the full subcategory* $\mathbf{Adm}_{\mathcal{X}}$ *of* $\mathbf{Adm}$ *whose objects are class-constant and admissible maps from* $\mathcal{X}$ *(i.e., maps* $\varphi\colon \mathcal{X} \to \mathcal{M}$ *satisfying* $x_1 \sim_A x_2 \iff \varphi(x_1) = \varphi(x_2)$*), the canonical quotient map* $q_A\colon \mathcal{X} \to \mathcal{X}/{\sim_A}$ *is initial: there is a unique morphism from* q_A *to every other object.*

Proof. For any object $\varphi\colon \mathcal{X} \to \mathcal{M}$ in $\mathbf{Adm}_{\mathcal{X}}$, the condition $x_1 \sim_A x_2 \iff \varphi(x_1) = \varphi(x_2)$ ensures class-constancy. Theorem 3.7 gives a unique $\bar{\varphi}\colon \mathcal{X}/{\sim_A} \to \mathcal{M}$ with $\varphi = \bar{\varphi} \circ q_A$. By Proposition 3.8, $\bar{\varphi}$ is injective. This $\bar{\varphi}$ is the unique morphism from q_A to φ in $\mathbf{Adm}_{\mathcal{X}}$. □

3.4. The morphism defect as a measure of non-initiality

For an inadmissible encoder φ, the morphism defect $\mathrm{Def}(\varphi) \neq \emptyset$ and φ does not factor through q_A. The admissibility distortion $D_A(\varphi)$ quantifies *how far* φ is from being a morphism in $\mathbf{Adm}$:

$$D_A(\varphi) = \mathbb{E}_\mu\big[\mathbf{1}_{\mathrm{Def}(\varphi)}(x_1, x_2) \cdot d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2))\big].$$

This is zero iff $\mathrm{Def}(\varphi) = \emptyset$ a.s., iff φ is (strongly) admissible a.s., iff φ factors through q_A (Proposition 2.8).

4. The Quotient Lattice and Compression Morphisms

4.1. The lattice of quotients

Definition 4.1 (Quotient Lattice)**.** For a fixed state space $\mathcal{X}$, the *quotient lattice* $\mathcal{L}(\mathcal{X})$ is the set of equivalence relations on $\mathcal{X}$ ordered by inclusion:

$$\sim_1 \leq \sim_2 \iff \sim_1 \subseteq \sim_2 \iff \forall x_1, x_2\colon x_1 \sim_1 x_2 \Rightarrow x_1 \sim_2 x_2.$$

We say $\sim_2$ is *coarser* than $\sim_1$ (and $\sim_1$ is *finer* than $\sim_2$) when $\sim_1 \leq \sim_2$.

Lemma 4.2 ($\mathcal{L}(\mathcal{X})$ is a Complete Lattice)**.** *$(\mathcal{L}(\mathcal{X}), \leq)$ is a complete lattice with meet given by intersection* $\sim_1 \wedge \sim_2 = \sim_1 \cap \sim_2$ *and join given by the equivalence closure of the union* $\sim_1 \vee \sim_2 = \overline{\sim_1 \cup \sim_2}$*. The bottom element is the diagonal (equality) relation* $\Delta_{\mathcal{X}}$ *and the top element is the total relation* $\mathcal{X} \times \mathcal{X}$*.*

Proof. Standard result in lattice theory. Intersection of equivalence relations is an equivalence relation. Equivalence closure of a union of equivalence relations is the smallest equivalence relation containing both; its existence follows from the fact that the intersection of any family of equivalence relations is an equivalence relation, so the closure is the intersection of all equivalence relations containing the union. □

Every equivalence relation $\sim_\varphi$ on $\mathcal{X}$ induced by an encoder φ (the kernel: $x_1 \sim_\varphi x_2 \iff \varphi(x_1) = \varphi(x_2)$) defines a point in $\mathcal{L}(\mathcal{X})$.

Proposition 4.3 (Lattice Embedding). *The assignment $\sim \mapsto \mathcal{X}/\sim$ defines an order-reversing embedding of $(\mathcal{L}(\mathcal{X}), \leq)$ into* **Adm***: if $\sim_1 \leq \sim_2$, there is a canonical surjective morphism $\pi_{12}\colon (\mathcal{X}/\sim_1, \sim_1^*) \to (\mathcal{X}/\sim_2, \sim_2^*)$ in* **Adm***, and this assignment is functorial.*

Proof. If $\sim_1 \subseteq \sim_2$, then any $\sim_1$-class $[x]_{\sim_1}$ is contained in a unique $\sim_2$-class $[x]_{\sim_2}$. Define $\pi_{12}([x]_{\sim_1}) = [x]_{\sim_2}$. This is well-defined: if $[x]_{\sim_1} = [y]_{\sim_1}$ then $x \sim_1 y$, hence $x \sim_2 y$, hence $[x]_{\sim_2} = [y]_{\sim_2}$. It is surjective since $q_{\sim_2} = \pi_{12} \circ q_{\sim_1}$. The morphism condition in **Adm** holds: $[x]_{\sim_1} \sim_1^* [y]_{\sim_1}$ in $\mathcal{X}/\sim_1$ means $[x]_{\sim_1} = [y]_{\sim_1}$, so $\pi_{12}([x]) = \pi_{12}([y])$, and equal elements are trivially $\sim_2^*$-related. Functoriality ($\pi_{12} = \pi_{23} \circ \pi_{12}$ when appropriate) follows from transitivity of class containment. □

4.2. Compression as upward movement

Definition 4.4 (Compression Morphism). A *compression morphism* from encoder φ_1 to encoder φ_2 is the canonical surjection $\pi\colon \mathcal{X}/\sim_{\varphi_1} \to \mathcal{X}/\sim_{\varphi_2}$ in **Adm** whenever $\sim_{\varphi_1} \leq \sim_{\varphi_2}$. We say φ_2 is a *compression* of φ_1.

Compression morphisms exist exactly when the second encoder merges at least as many pairs as the first. In $\mathcal{L}(\mathcal{X})$, passing from φ_1 to φ_2 is upward movement toward coarser equivalence relations.

Theorem 4.5 (Monotonicity of Distortion). *If $\sim_{\varphi_1} \leq \sim_{\varphi_2}$ (i.e., a compression morphism exists from φ_1 to φ_2), then*

$$D_A(\varphi_1) \leq D_A(\varphi_2).$$

Moreover, the increase is given by the compression cost identity*:*

$$D_A(\varphi_2) - D_A(\varphi_1) = \mathbb{E}_\mu[\mathbf{1}_\mathcal{N}(x_1, x_2) \cdot d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2))],$$

where $\mathcal{N} = \sim_{\varphi_2} \setminus \sim_{\varphi_1}$ is the set of newly collapsed pairs.

Proof. Since $\sim_{\varphi_1} \subseteq \sim_{\varphi_2}$, the support of the indicator $\mathbf{1}[\varphi_1(x_1) = \varphi_1(x_2)]$ is contained in the support of $\mathbf{1}[\varphi_2(x_1) = \varphi_2(x_2)]$. Write $\sim_{\varphi_2} = \sim_{\varphi_1} \cup \mathcal{N}$ (disjoint union over admissibility-inequivalent pairs). Then

$$\begin{aligned} D_A(\varphi_2) &= \mathbb{E}_\mu[\mathbf{1}_{\sim_{\varphi_2}} \cdot d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2))] \\ &= \mathbb{E}_\mu[\mathbf{1}_{\sim_{\varphi_1}} \cdot d(\cdot)] + \mathbb{E}_\mu[\mathbf{1}_\mathcal{N} \cdot d(\cdot)] \\ &= D_A(\varphi_1) + \mathbb{E}_\mu[\mathbf{1}_\mathcal{N} \cdot d(\mathcal{R}_H^\Pi(x_1), \mathcal{R}_H^\Pi(x_2))]. \end{aligned}$$

Since $d \geq 0$, the second term is non-negative, establishing $D_A(\varphi_1) \leq D_A(\varphi_2)$. □

Remark 4.6 (Distortion as a Height Function). Theorem 4.5 gives $D_A(\varphi)$ the structure of a height function on $\mathcal{L}(\mathcal{X})$: it increases weakly as one moves upward toward coarser equivalence relations. Every compression morphism is a step toward higher distortion. The admissibility quotient q_A corresponds to the point $\sim_A \in \mathcal{L}(\mathcal{X})$, and any encoder above $\sim_A$ in the lattice (coarser than $\sim_A$) collapses admissibility-inequivalent pairs, incurring positive distortion.

4.3. The compression–admissibility tradeoff as a functor

Definition 4.7 (Category of Encoders)**.** The *category of encoders* $\mathbf{Comp}(\mathcal{X})$ has encoders $\varphi\colon \mathcal{X} \to \mathcal{M}$ as objects and compression morphisms as morphisms. Composition is composition of the canonical surjections.

Proposition 4.8 (Distortion Functor)**.** *The admissibility distortion defines a functor*

$$D_A\colon \mathbf{Comp}(\mathcal{X}) \longrightarrow (\mathbb{R}_{\geq 0}, \leq),$$

where $(\mathbb{R}_{\geq 0}, \leq)$ *is the poset regarded as a category (one morphism* $r_1 \to r_2$ *whenever* $r_1 \leq r_2$*). This functor sends compression morphisms to inequalities* $D_A(\varphi_1) \leq D_A(\varphi_2)$*.*

Proof. Theorem 4.5 establishes that D_A is order-preserving on the quotient lattice, hence on $\mathbf{Comp}(\mathcal{X})$. Functoriality reduces to this monotonicity, together with the trivial observation that identity morphisms map to equality $D_A(\varphi) \leq D_A(\varphi)$. □

The information bottleneck problem asks: given a mutual information budget I, find the most compressed encoder with $D_A(\varphi) \leq \epsilon$. In $\mathbf{Comp}(\mathcal{X})$, this is: find the initial object in the full subcategory of encoders with distortion at most ϵ. The Lagrangian relaxation

$$\min_{\varphi}\big[I(\mathcal{X};\varphi(\mathcal{X})) + \lambda \cdot D_A(\varphi)\big]$$

traces a path in $\mathbf{Comp}(\mathcal{X})$ parameterised by $\lambda \geq 0$: at $\lambda = 0$ we recover maximal compression (the terminal object, coarsest encoder), and as $\lambda \to \infty$ the path converges to q_A (the initial admissible object of Corollary 3.9).

Example 4.9 (Token Pruning as Controlled Quotienting)**.** Transformer-based vision models operate on sequences of visual tokens. Token pruning reduces the sequence length by discarding tokens judged uninformative. Each pruning operation induces an equivalence relation $\sim_{\text{prune}}$ on the token space: tokens collapsed to a single representation become equivalent under $\sim_{\text{prune}}$. Pruning is therefore movement in the quotient lattice $\mathcal{L}(\mathcal{X})$: each step coarsens the equivalence relation. The performance question—does pruning degrade task accuracy?—is precisely the admissibility question: does the induced compression cross the threshold $\sim_A$ for the downstream task?

Architectures that compute per-token importance scores before pruning are performing an approximation to admissibility-aware quotient construction: they attempt to preserve the equivalence classes that the task (here, keypoint localisation or action recognition) requires to remain distinct, and collapse only those tokens whose reachable "inference futures" are indistinguishable. The compression cost identity of Theorem 4.5 gives a precise account of the distortion incurred by each pruning step: it is the expected reachability separation over the set of newly collapsed token pairs.

5. The Reachability Sheaf

5.1. The Alexandrov topology on $\mathcal{X}$

The fixed-horizon reachability relation $x' \in \mathcal{R}_H^\Pi(x)$ does not in general define a preorder: reflexivity requires every state to reach itself in exactly H steps (which fails for $H \geq 1$ in acyclic environments), and transitivity requires that reachability in H steps followed by reachability in H steps yields reachability in H steps (which holds only if $H + H = H$, i.e., never). We therefore base the topology on the *transitive closure* of reachability, which is the natural planning-relevant notion: a state is accessible if it can be reached in *some* finite number of steps.

Definition 5.1 (Reachability Preorder)**.** The *reachability preorder* $\leq_A$ on $\mathcal{X}$ is defined by

$$x' \leq_A x \iff \exists\, h \in \mathbb{N},\ \exists\, \pi \in \Pi,\ \exists \,\text{path}\ x = x_0, x_1, \ldots, x_h = x' \text{ with } x_{t+1} \in \operatorname{supp}(\pi(x_t))\ \forall t,$$

i.e., x' is at or below x in the order iff x' is reachable from x in some finite number of steps under some policy. We write $\mathcal{R}^\Pi(x) = \{x' : x' \leq_A x\}$ for the full reachability set of x.

Remark 5.2 (Relation to the Fixed-Horizon Reachability Set). The distortion functional $D_A(\varphi)$ of Definition 2.7 uses the fixed-horizon reachability sets $\mathcal{R}_H^\Pi(x)$ as the objects being compared. The preorder $\leq_A$ uses the *cumulative* reachability $\mathcal{R}^\Pi(x) = \bigcup_{h \geq 0} \mathcal{R}_h^\Pi(x)$. These are the appropriate objects for two different purposes: the fixed-horizon set captures the precise planning consequences at horizon H and is the right object for the distortion functional; the cumulative set captures the long-run accessibility structure of the state space and is the right basis for a topology. The admissibility equivalence relation $\sim_A$ is defined in terms of $\mathcal{R}_H^\Pi$ and remains unchanged.

Lemma 5.3 ($\leq_A$ is a Preorder)**.** *The relation $\leq_A$ is reflexive and transitive, hence a preorder on $\mathcal{X}$.*

Proof. Reflexivity. Every state x reaches itself in $h = 0$ steps (the empty path), so $x \leq_A x$.

Transitivity. If $x'' \leq_A x'$ and $x' \leq_A x$, then there exist paths of lengths h_1 and h_2 from x to x' and from x' to x'' respectively. Concatenating these paths gives a path of length $h_1 + h_2$ from x to x'', so $x'' \leq_A x$. □

Definition 5.4 (Reachability Topology)**.** The *reachability topology* τ_A on $\mathcal{X}$ is the *Alexandrov topology* of the preorder $\leq_A$: a set $U \subseteq \mathcal{X}$ is open iff it is an *up-set* of $\leq_A$,

$$x \in U \text{ and } x' \leq_A x \implies x' \in U,$$

i.e., if a state is in U then all states reachable from it are also in U.

Lemma 5.5 (τ_A is a Topology)**.** *The collection τ_A is a topology on $\mathcal{X}$.*

Proof. $\emptyset$ and $\mathcal{X}$ satisfy the up-set condition vacuously and trivially. Arbitrary unions of up-sets are up-sets: if $x \in \bigcup_\alpha U_\alpha$ and $x' \leq_A x$, then $x \in U_\alpha$ for some α, so $x' \in U_\alpha$. Finite intersections of up-sets are up-sets: if $x \in U_1 \cap \cdots \cap U_n$ and $x' \leq_A x$, then $x' \in U_i$ for each i, so $x' \in U_1 \cap \cdots \cap U_n$. □

Remark 5.6 (Basis for τ_A). The principal up-sets $\uparrow x = \{x' \in \mathcal{X} : x' \leq_A x\} = \mathcal{R}^\Pi(x)$ are open in τ_A (by definition of up-sets) and form a basis: every open set is a union of principal up-sets. This is the Alexandrov basis theorem applied to $(\mathcal{X}, \leq_A)$.

Definition 5.7 (Observation Topology)**.** Define the *observation preorder* $\leq_O$ by $x' \leq_O x$ iff $P(O_{1:\infty} \mid x') = P(O_{1:\infty} \mid x)$, where $O_{1:\infty}$ is the passive observation sequence generated from state x. The *observation topology* τ_O is the Alexandrov topology of $\leq_O$.

The Observational–Interventional Separation phenomenon (see Example 2.11) is equivalent to $\tau_O \neq \tau_A$: the two topologies are in general distinct. The observation topology is typically coarser—larger equivalence classes, fewer open sets—since it cannot distinguish states with identical passive distributions but different reachability structures.

5.2. The reachability presheaf and sheaf

Definition 5.8 (Reachability Presheaf)**.** The *reachability presheaf* $\mathcal{F}_A$ on $(\mathcal{X}, \tau_A)$ assigns to each open set $U \in \tau_A$ the set of *reachability sections over* U:

$$\mathcal{F}_A(U) = \{s \colon U \to \mathcal{P}(\mathcal{X}) \mid s(x) \subseteq \mathcal{R}^\Pi_H(x) \text{ for all } x \in U\}.$$

For $V \subseteq U$ open, the restriction map $\rho_{UV} \colon \mathcal{F}_A(U) \to \mathcal{F}_A(V)$ sends $s \mapsto s|_V$ (restriction of the function to V).

A *section* of $\mathcal{F}_A$ over U is an element $s \in \mathcal{F}_A(U)$: a function assigning to each state in U a subset of its reachable futures. The canonical global section is $s^* \colon \mathcal{X} \to \mathcal{P}(\mathcal{X})$, $s^*(x) = \mathcal{R}^\Pi_H(x)$: the complete reachability description.

Lemma 5.9 ($\mathcal{F}_A$ is a Sheaf)**.** *$\mathcal{F}_A$ satisfies the sheaf axioms on $(\mathcal{X}, \tau_A)$.*

Proof. Locality. If $s, t \in \mathcal{F}_A(U)$ and $s|_{U_i} = t|_{U_i}$ for all opens U_i in a cover of U, then $s(x) = t(x)$ for all $x \in U$ (since every x lies in some U_i), so $s = t$.

Gluing. Given a cover $\{U_i\}$ of U and compatible sections $s_i \in \mathcal{F}_A(U_i)$ (compatible meaning $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all i, j), define $s\colon U \to \mathcal{P}(\mathcal{X})$ by $s(x) = s_i(x)$ for any i with $x \in U_i$. Compatibility ensures this is well-defined. Since $s_i(x) \subseteq \mathcal{R}_H^\Pi(x)$ for all i and all $x \in U_i$, we have $s(x) \subseteq \mathcal{R}_H^\Pi(x)$ for all $x \in U$. Hence $s \in \mathcal{F}_A(U)$ and $s|_{U_i} = s_i$. □

5.3. The observation presheaf

Definition 5.10 (Observation Presheaf)**.** The *observation presheaf* $\mathcal{F}_O$ on $(\mathcal{X}, \tau_O)$ assigns to each open $U \in \tau_O$ the set of observation sections over U:

$$\mathcal{F}_O(U) = \{f\colon U \to \mathcal{D} \mid f(x) = P(O_{1:\infty} \mid x) \text{ for all } x \in U\},$$

where $\mathcal{D}$ is the space of probability distributions over observation sequences. Restriction maps are restrictions of functions.

Each section of $\mathcal{F}_O$ over U is (up to the constraint that $f(x)$ must equal the observation distribution at x) essentially the single function $x \mapsto P(O_{1:\infty} \mid x)$ restricted to U. The presheaf records what passive observation reveals about each state.

Remark 5.11. The observation presheaf $\mathcal{F}_O$ is actually a sheaf on $(\mathcal{X}, \tau_O)$ by the same argument as Lemma 5.9: sections are functions into $\mathcal{D}$, and functions glue uniquely. However, $\mathcal{F}_O$ is a sheaf on $(\mathcal{X}, \tau_O)$, not on $(\mathcal{X}, \tau_A)$, because the observation topology is the natural domain for passive observation data.

5.4. The gluing failure and the Observational–Interventional Separation

The canonical global section of $\mathcal{F}_A$ encodes the complete reachability structure. Planning requires this global section. Passive observation provides local sections of $\mathcal{F}_O$. The central question is whether these local observational sections can be assembled to determine the global reachability section.

Theorem 5.12 (Sheaf Formulation of Observational–Interventional Separation)**.** *Let $\mathcal{U} = \{U_i\}$ be a cover of $\mathcal{X}$ by τ_O-open sets. There exist compatible local sections $s_i \in \mathcal{F}_O(U_i)$—satisfying $s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all i, j—that are consistent with more than one distinct global section of $\mathcal{F}_A$. That is, the global reachability section is not uniquely determined by any complete system of local observational sections.*

Proof. By Example 2.11, there exist states $x_1 = s_0^0$ and $x_2 = s_0^1$ with $P(O_{1:\infty} \mid x_1) = P(O_{1:\infty} \mid x_2)$ but $\mathcal{R}_H^\Pi(x_1) \neq \mathcal{R}_H^\Pi(x_2)$. Since the observation distributions are equal, x_1 and x_2 lie in the same τ_O-equivalence class and hence in the same element of any τ_O-cover $\mathcal{U}$: there exists U_{i_0} with $\{x_1, x_2\} \subseteq U_{i_0}$.

A local section $s_{i_0} \in \mathcal{F}_O(U_{i_0})$ assigns $s_{i_0}(x) = P(O_{1:\infty} \mid x)$ for $x \in U_{i_0}$. Since x_1 and x_2 have identical observation distributions, the observational data s_{i_0} is equally consistent with the reachability assignment $x_1 \mapsto \mathcal{R}_H^\Pi(x_1), x_2 \mapsto \mathcal{R}_H^\Pi(x_2)$ and with the swapped assignment $x_1 \mapsto \mathcal{R}_H^\Pi(x_2), x_2 \mapsto \mathcal{R}_H^\Pi(x_1)$.

Both assignments define valid global sections of $\mathcal{F}_A$ (since both assign to each state a subset of its actual reachable futures). The local observational data does not distinguish between them. Hence compatible local sections of $\mathcal{F}_O$ do not uniquely determine the global section of $\mathcal{F}_A$. □

Remark 5.13. Theorem 5.12 reformulates an equivalence-relation statement as a sheaf statement. The equivalence-relation version says: observational equivalence and admissibility equivalence are distinct relations. The sheaf version says: local sections of $\mathcal{F}_O$ do not satisfy the uniqueness part of sheaf gluing when regarded as data for $\mathcal{F}_A$. The sheaf formulation is geometrically richer: it explains *why* the failure occurs (the cover is too coarse to separate states that $\mathcal{F}_A$ must separate) and points toward the cohomological quantification in Section 6.

Example 5.14 (Biological Ageing as Local–Global Sheaf Structure)**.** Medical imaging provides a concrete empirical instance of local sections and global sections in the sense of Theorem 5.12. Separate predictors trained on imaging data from individual organs—brain, liver, lung, heart, spine—each estimate an organ-specific biological age, corresponding to a local section of an ageing sheaf over the anatomical cover of the body. A whole-body predictor attempts to integrate these local estimates into a global biological age, corresponding to a global section. The empirical finding that local organ ages are individually predictable but do not uniquely determine the global ageing trajectory is precisely the sheaf-theoretic non-uniqueness of Theorem 5.12: compatible local sections (organ age estimates) do not determine a unique global section (whole-body biological age) because different combinations of organ ages can produce the same observational signature.

Interventional experiments that modify one organ's apparent age while holding others fixed further illustrate the distinction between local observational access and global reachability structure: a local change can alter the global trajectory, but local observation alone cannot predict by how much. This is the biological analogue of the planning scenario: the global section of $\mathcal{F}_A$ (the reachability structure) is not determined by the local sections of $\mathcal{F}_O$ (the passive observations), and the admissibility distortion of any observation-based predictor measures precisely this gap.

6. Cohomological Perspectives

Theorem 5.12 identifies a failure of gluing: compatible local observational sections do not determine a unique global reachability section. In sheaf cohomology, obstructions to gluing are measured by the first cohomology group of the relevant sheaf. This section sets up the Čech cohomology machinery and states the central conjecture of this paper.

6.1. The Čech complex

Let $\mathcal{U} = \{U_i\}_{i \in I}$ be a τ_O-cover of $\mathcal{X}$, and let $\mathcal{F}_A$ be the reachability sheaf of Section 5. The Čech complex of $\mathcal{F}_A$ with respect to $\mathcal{U}$ is

$$0 \longrightarrow \mathcal{F}_A(\mathcal{X}) \xrightarrow{\delta^0} \prod_{i \in I} \mathcal{F}_A(U_i) \xrightarrow{\delta^1} \prod_{i<j} \mathcal{F}_A(U_i \cap U_j) \longrightarrow \cdots \tag{2}$$

The coboundary maps are:

$$\delta^0(s)_i = s|_{U_i},$$
$$\delta^1(s_i)_{ij} = s_j|_{U_i \cap U_j} - s_i|_{U_i \cap U_j},$$

where the difference in the second line requires $\mathcal{F}_A$ to take values in an abelian group (see Remark 6.1 below). The first Čech cohomology group

$$\check{H}^1(\mathcal{U}; \mathcal{F}_A) = \ker(\delta^1)/\mathrm{im}(\delta^0)$$

classifies the obstructions to lifting compatible local sections to global sections: a class $[\{s_{ij}\}] \in \check{H}^1(\mathcal{U}; \mathcal{F}_A)$ is non-trivial iff the system of local sections $\{s_i\}$ satisfies compatibility on overlaps but does not come from a global section.

Remark 6.1 (Linearisation). As defined, $\mathcal{F}_A(U)$ is a set of functions $U \to \mathcal{P}(\mathcal{X})$, not an abelian group. To define a Čech complex with $\mathbb{R}$-valued coboundaries, one passes to a linearisation: replace $\mathcal{F}_A(U)$ by a sheaf of real vector spaces via the map $s \mapsto \mathbf{1}_s$ (indicator function on $\mathcal{X} \times \mathcal{X}$), or by working with a sheaf of signed measures on $\mathcal{P}(\mathcal{X})$. The appropriate linearisation depends on the metric d used to define admissibility distortion $D_A(\varphi)$. We proceed formally under the assumption that such a linearisation has been fixed, deferring the detailed construction to future work.

6.2. The central conjecture

We are now in a position to state the main conjecture of this paper. The admissibility distortion $D_A(\varphi)$ is a non-negative real number quantifying how far a given encoder φ is from admissibility. The first Čech cohomology $\check{H}^1(\mathcal{U}; \mathcal{F}_A)$ is an algebraic object measuring the obstruction to gluing. The conjecture asserts that these are related.

Conjecture 6.2 (Cohomological Admissibility Obstruction). Let $\mathcal{U}$ be the τ_O-cover of $\mathcal{X}$ by observational indistinguishability classes. After appropriate linearisation of $\mathcal{F}_A$ as in Remark 6.1, the first Čech cohomology group $\check{H}^1(\mathcal{U}; \mathcal{F}_A)$ is non-trivial if and only if there exists an inadmissible encoder φ (i.e., $D_A(\varphi) > 0$) consistent with the observational data provided by $\mathcal{U}$. Moreover, there is a natural map

$$\| \cdot \| \colon \check{H}^1(\mathcal{U}; \mathcal{F}_A) \to \mathbb{R}_{\geq 0}$$

(a norm on the cohomology group with respect to the chosen linearisation) such that

$$\|[\delta s]\| \;\geq\; c \cdot \inf_{\varphi \text{ consistent}} D_A(\varphi)$$

for some constant $c > 0$ depending only on $\mathcal{U}$ and the metric d.

Remark 6.3 (What would constitute a proof). A proof of Conjecture 6.2 would require: (a) a precise linearisation of $\mathcal{F}_A$ as a sheaf of abelian groups; (b) an explicit identification of the cohomology class $[\delta s]$ associated to an inadmissible encoder; (c) a norm on $\check{H}^1$ (or a natural map to $\mathbb{R}$) compatible with the admissibility distortion; and (d) a lower bound on the norm in terms of $D_A(\varphi)$. Steps (a) and (b) are formal but non-trivial; steps (c) and (d) are the substantive mathematical content. We believe this programme is tractable and constitutes the most important open problem in the geometry of admissibility.

Remark 6.4 (Relation to the Separation Theorem). Theorem 5.12 establishes the *existence* of a non-trivial obstruction (there are observational covers that do not determine the global reachability section). Conjecture 6.2 asserts that this obstruction is *measured* by admissibility distortion via a cohomological norm. The former is a theorem; the latter is a quantitative conjecture that would, if proved, make the phrase "admissibility distortion is the cohomological obstruction" into a theorem rather than a heuristic.

7. Temporal Functors and Meta-Admissibility

7.1. The temporal base category

Let $(T, \leq)$ be a poset of time indices. We take $T = \mathbb{N}$ with the natural order in the discrete case, but the construction applies to any poset.

Definition 7.1 (Temporal Category)**.** The *temporal category* $\mathbf{T}$ is the category associated to the poset $(T, \leq)$: objects are elements of T, and there is a unique morphism $t \to t'$ whenever $t \leq t'$. Composition is given by transitivity.

7.2. Temporal admissibility functors

Definition 7.2 (Temporal Admissibility Functor)**.** A *temporal admissibility functor* is a functor $\mathbf{X}\colon \mathbf{T} \to \mathbf{Adm}$. It assigns:

(i) to each time $t \in T$, an admissibility structure $(\mathcal{X}_t, \sim_A^{(t)})$;
(ii) to each morphism $t \to t'$ in $\mathbf{T}$ (i.e., $t \leq t'$), a morphism $F_{tt'}\colon (\mathcal{X}_t, \sim_A^{(t)}) \to (\mathcal{X}_{t'}, \sim_A^{(t')})$ in $\mathbf{Adm}$;

subject to the functoriality conditions: $F_{tt} = \mathrm{id}_{\mathcal{X}_t}$ for all t, and $F_{t't''} \circ F_{tt'} = F_{tt''}$ for all $t \leq t' \leq t''$.

The morphism condition on $F_{tt'}$ is:

$$x_1 \sim_A^{(t)} x_2 \implies F_{tt'}(x_1) \sim_A^{(t')} F_{tt'}(x_2) :$$

admissibility-equivalent states at time t map to admissibility-equivalent states at time t'. This says that the transformation $F_{tt'}$ does not create new admissibility-distinctions by collapsing states that were admissibility-distinct.

7.3. Meta-admissibility as a functor condition

Definition 7.3 (Meta-Admissibility)**.** A sequence of state-space transformations $\{F_{t,t+1} : \mathcal{X}_t \to \mathcal{X}_{t+1}\}_{t \in T}$ is *meta-admissible* if and only if it extends to a temporal admissibility functor $\mathbf{X} : \mathbf{T} \to \mathbf{Adm}$.

Theorem 7.4 (Meta-Admissibility as Naturality)**.** *A sequence $\{F_{t,t+1}\}$ is meta-admissible if and only if for all $t \leq t'$, the diagram*

$$\begin{array}{ccc} \mathcal{X}_t & \xrightarrow{\;F_{tt'}\;} & \mathcal{X}_{t'} \\ {\scriptstyle q_A^{(t)}}\big\downarrow & & \big\downarrow{\scriptstyle q_A^{(t')}} \\ \mathcal{X}_t / \sim_A^{(t)} & \overset{\bar{F}_{tt'}}{\dashrightarrow} & \mathcal{X}_{t'} / \sim_A^{(t')} \end{array} \tag{3}$$

commutes, where $\bar{F}_{tt'}$ is the map on quotients induced by $F_{tt'}$ via the universal property of $q_A^{(t')}$.

Proof. The commutativity of (3) is equivalent to $q_A^{(t')} \circ F_{tt'} = \bar{F}_{tt'} \circ q_A^{(t)}$. The map $\bar{F}_{tt'}$ exists by the universal property (Theorem 3.7) iff $q_A^{(t')} \circ F_{tt'}$ is strongly admissible from $(\mathcal{X}_t, \sim_A^{(t)})$, which holds iff $F_{tt'}$ is a morphism in **Adm**: $x_1 \sim_A^{(t)} x_2 \Rightarrow F_{tt'}(x_1) \sim_A^{(t')} F_{tt'}(x_2)$. The functoriality conditions then follow from identity and composition in **Adm**. □

Corollary 7.5 (Closure Under Composition)**.** *If $F_{t,t+1}$ and $F_{t+1,t+2}$ are both admissibility-preserving, then $F_{t,t+2} = F_{t+1,t+2} \circ F_{t,t+1}$ is admissibility-preserving. Meta-admissibility is closed under composition.*

Proof. Morphisms in **Adm** are closed under composition by Lemma 3.3. □

Remark 7.6 (Non-morphisms Do Not Compose to Non-morphisms). One might expect that a single inadmissible step should propagate forward, "contaminating" all subsequent admissibility structure. This expectation is incorrect. In category theory, the composition of a non-morphism with a morphism is in general neither a morphism nor a non-morphism in any systematic way: a later transformation $F_{t+1,t+2}$ can re-separate states that $F_{t,t+1}$ collapsed, thereby restoring admissibility at the composed level. Formally, $f \notin \mathrm{Mor}(\mathbf{Adm})$ does not imply $g \circ f \notin \mathrm{Mor}(\mathbf{Adm})$.

What is true is the contrapositive of Corollary 7.5: if the composed transformation $F_{t,t+2}$ is not admissibility-preserving, then at least one of $F_{t,t+1}$ or $F_{t+1,t+2}$ is not admissibility-preserving. The diagnostic question is therefore: does the composed map fail to be a morphism? If so, at least one step in the composition is inadmissible, though which one (or ones) requires separate analysis.

7.4. Natural transformations between developmental trajectories

Definition 7.7 (Natural Transformation of Trajectories)**.** Two temporal admissibility functors $\mathbf{X}, \mathbf{Y} \colon \mathbf{T} \to \mathbf{Adm}$ are compared by a *natural transformation* $\eta \colon \mathbf{X} \Rightarrow \mathbf{Y}$: a family of morphisms $\eta_t \colon (\mathcal{X}_t, \sim_A^{X\,(t)}) \to (\mathcal{Y}_t, \sim_A^{Y\,(t)})$ in $\mathbf{Adm}$ such that for all $t \le t'$,

$$\eta_{t'} \circ F^X_{tt'} = F^Y_{tt'} \circ \eta_t.$$

A natural transformation η between two trajectories is a time-indexed family of admissibility-preserving maps between their state spaces that commutes with temporal evolution in both trajectories. It provides a systematic comparison of two developmental paths—two agents, or two stages of the same agent under different conditions—that respects admissibility structure at every time step.

8. The Admissibility Programme as a Geometric Object

The four structures developed in this paper fit together into a coherent geometric picture of the admissibility programme.

At the level of a fixed state space $\mathcal{X}$ with fixed policy class Π and horizon H: the admissibility quotient $q_A \colon \mathcal{X} \to \mathcal{X}/{\sim_A}$ is the initial object in the category of strongly admissible projections (Corollary 3.9); every encoder φ defines a point in the quotient lattice $\mathcal{L}(\mathcal{X})$, and admissibility distortion $D_A(\varphi)$ is a monotone height function on this lattice increasing along compression morphisms (Theorem 4.5); the global section $s^* \colon x \mapsto \mathcal{R}^{\Pi}_H(x)$ of $\mathcal{F}_A$ encodes the complete reachability structure that planning requires; planning-adequate representations are exactly those encoders φ for which $\varphi(x_1) \ne \varphi(x_2)$ whenever $s^*(x_1) \ne s^*(x_2)$—admissibility is the condition that φ separates the fibres of s^*—and passive observation determines only local sections of $\mathcal{F}_O$ over the coarser topology τ_O (Theorem 5.12).

At the level of an evolving state space: temporal evolution is a functor $\mathbf{X} \colon \mathbf{T} \to \mathbf{Adm}$ (Definition 7.2); meta-admissibility is the condition that this functor exists with the required commutativity (Theorem 7.4); comparisons between developmental trajectories are natural transformations between such functors.

The unifying geometric statement is:

> *Observation provides local sections of $\mathcal{F}_O$ over the observation topology. Planning requires a global section of $\mathcal{F}_A$ over the reachability topology. Admissibility*

distortion quantifies the obstruction to assembling the former into the latter; Conjecture 6.2 proposes that this obstruction is measured by the first Čech cohomology of $\mathcal{F}_A$ with respect to observational covers.

This is the geometric form of the Observational–Interventional Separation Theorem: not a logical implication between equivalence classes, but a failure of the sheaf gluing axiom, measurable in principle by a cohomological invariant. The admissibility programme—naming, estimating, and minimising $D_A(\varphi)$—is, in this reading, the programme of computing, bounding, and eventually eliminating a cohomological obstruction.

9. Open Problems

1. Proof of Conjecture 6.2. The central open problem is to prove (or disprove) the cohomological admissibility obstruction conjecture. The first step is a precise linearisation of $\mathcal{F}_A$ as a sheaf of abelian groups. The second is an explicit identification of the cohomology class associated to admissibility distortion. The third is a norm on $\check{H}^1(\mathcal{U}; \mathcal{F}_A)$ with the required lower bound. Whether the natural linearisation by indicator functions on $\mathcal{X} \times \mathcal{X}$ yields a tractable cohomology theory is the first technical question.

2. Čech cohomology of the reachability sheaf. Even independently of Conjecture 6.2, the Čech cohomology $\check{H}^*(\mathcal{U}; \mathcal{F}_A)$ is a new mathematical object. Its properties—whether higher cohomology groups are trivial, what covering conditions minimise the first cohomology, how the cohomology changes under changes in Π and H—are largely unexplored. In particular, the question of whether the reachability sheaf is acyclic (all higher cohomology vanishing) would clarify the scope of obstructions.

3. Derived functors and the quotient sheaf. If $\mathcal{F}_A$ is treated as a sheaf of abelian groups, the right-derived functors of the global sections functor Γ give sheaf cohomology groups $H^k(\mathcal{X}, \mathcal{F}_A)$. Admissibility distortion may be recoverable as a specific class in $H^1(\mathcal{X}, \mathcal{F}_A/\mathcal{F}_O)$, the cohomology of the quotient sheaf. The relationship between Čech cohomology and derived-functor cohomology in this setting—which requires the reachability sheaf to satisfy an acyclicity condition on the reachability topology—is a natural question.

4. The bisimulation tower. Bisimulation equivalence is a refinement of admissibility equivalence: $x_1 \sim_{\text{bis}} x_2$ iff the transition distributions are equal, which is a strictly stronger condition than $\mathcal{R}_H^\Pi(x_1) = \mathcal{R}_H^\Pi(x_2)$. The bisimulation sheaf $\mathcal{F}_{\text{bis}}$ sits between $\mathcal{F}_O$ and $\mathcal{F}_A$ in a tower of sheaves over $\mathcal{X}$. The obstruction to lifting sections of $\mathcal{F}_O$ to $\mathcal{F}_{\text{bis}}$ may be computable in practice (since bisimulation metrics have efficient approximations), providing a computable lower bound on the obstruction to lifting to $\mathcal{F}_A$.

5. Topos-theoretic internalisation. The category **Adm** with the Alexandrov topology on each object suggests a site $(\mathbf{Adm}, \tau_A)$ obtained by equipping **Adm** with a Grothendieck topology. The topos of sheaves $\mathbf{Sh}(\mathbf{Adm}, \tau_A)$ would provide a generalised geometric universe in which admissibility arguments can be conducted internally. Whether the internal logic of this topos has a natural interpretation in terms of planning-theoretic reasoning—and in particular, whether the universal property of q_A has a clean internal formulation—is an open question. The definition of the Grothendieck topology on **Adm** itself (as opposed to on individual objects) requires separate development.

6. Persistent homology of reachability. For metric state spaces with $\mathcal{X} \subseteq \mathbb{R}^n$, the reachability sets $\mathcal{R}_H^\Pi(x)$ vary with x. The filtration

$$\{\mathcal{R}_h^\Pi(x)\}_{h=0}^{H}$$

of reachability sets as the horizon h grows from 0 to H defines a filtered topological space for each x. The persistent homology of this filtration captures how topological features of the reachability geometry appear and disappear as the planning horizon grows. The birth and death times of homological features are intrinsic invariants of the policy class Π and may provide a topological analogue of admissibility depth.

7. Enriched and $(\infty, 1)$-categorical generalisations. The category **Adm** as defined is an ordinary category. Two natural enrichments are worth exploring: an *enrichment over metric spaces* (with the admissibility distortion $D_A(\varphi)$ as a distance between morphisms) and an $(\infty, 1)$-*categorical* version (replacing sets of morphisms by spaces of morphisms, with higher morphisms encoding homotopies of admissibility-preserving maps). The latter would connect the admissibility programme to the emerging theory of reachability in homotopy type theory.

10. Conclusion

We have developed the geometry of admissibility through four interlocking structures. The category **Adm** provides the natural language for speaking about admissibility-preserving maps: admissible representations are morphisms, the admissibility quotient is an initial object with a universal property, and admissibility distortion measures deviation from morphismhood. The quotient lattice organises all possible compressions of a state space into a partial order on which distortion is monotone, making the compression–admissibility tradeoff a functor. The reachability sheaf over the Alexandrov topology shows that planning requires global sections and observation provides only local ones, with the separation between the two sheaves being a failure of sheaf gluing. Temporal admissibility functors capture evolving state spaces and characterise meta-admissibility as the functor condition.

The central mathematical object that emerges is the reachability sheaf $\mathcal{F}_A$ over the Alexandrov topology τ_A. Its canonical global section $s^*\colon x \mapsto \mathcal{R}^\Pi(x)$ encodes the complete reachability structure. A representation $\varphi\colon \mathcal{X} \to \mathcal{M}$ is planning-adequate precisely when it separates the fibres of s^*: admissibility is the condition $s^*(x_1) \neq s^*(x_2) \Rightarrow \varphi(x_1) \neq \varphi(x_2)$. The distance from the observation sheaf $\mathcal{F}_O$ to the reachability sheaf is the geometric form of the Observational–Interventional Separation Theorem. The quantitative measure of this distance, which we conjecture to be the first Čech cohomology class of $\mathcal{F}_A$ relative to observational covers, is admissibility distortion. The admissibility programme—developing estimators for $D_A(\varphi)$, designing admissible encoders, characterising the information-theoretic constraints on planning—is, in this geometric reading, the programme of computing, bounding, and eventually eliminating a cohomological obstruction.

Acknowledgements. The author has no institutional affiliation and no conflicts of interest.

References

[1] Kashiwara, M., and Schapira, P. (2006). *Categories and Sheaves*. Springer.

[2] Mac Lane, S. (1998). *Categories for the Working Mathematician*, 2nd ed. Springer.

[3] Mac Lane, S., and Moerdijk, I. (1992). *Sheaves in Geometry and Logic: A First Introduction to Topos Theory*. Springer.

[4] Ferns, N., Panangaden, P., and Precup, D. (2004). Metrics for finite Markov decision processes. *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence (UAI)*, 162–169.

[5] Pearl, J. (2009). *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge University Press.

[6] Quillen, D. (1967). *Homotopical Algebra*. Lecture Notes in Mathematics, vol. 43. Springer.

[7] Bredon, G. E. (1997). *Sheaf Theory*, 2nd ed. Graduate Texts in Mathematics, vol. 170. Springer.

[8] Grothendieck, A. (1957). Sur quelques points d'algèbre homologique. *Tôhoku Mathematical Journal*, 9(2), 119–221.

[9] Chen, Z., Wang, Y., Bai, Y., et al. (2026). MAE-YOLO improves small object detection for intelligent inspection. *Scientific Reports*.

[10] Mueller, T.T., Starck, S., Llalloshi, R., et al. (2026). Local and global patterns support medical imaging as a biomarker of ageing. *Communications Medicine*, 6, 335.

[11] Zhang, J., Deng, Z., Li, X., et al. (2026). Enhanced human pose estimation via self-distilled and token-pruned transformer. *Scientific Reports*.